

Durham E-Theses

Valuing Farm Animal Welfare in a Market Economy: A Philosophical Study of Market Failure

VENTIN, MARCUS

How to cite:

VENTIN, MARCUS (2020) *Valuing Farm Animal Welfare in a Market Economy: A Philosophical Study of Market Failure*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/13661/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Valuing Farm Animal Welfare in a Market Economy

A Philosophical Study of Market Failure

Marcus Ventin

Abstract:

Do people in the UK care about farm animal welfare? How well are their concerns represented in society? How can we use economic policy to build a society that reflects public attitudes towards farm animals and their welfare?

In this thesis, I contend that markets in animal products – which facilitate many people’s quotidian interactions with farm animal welfare – are susceptible to four forms of market failure: externalities, public good problems, information asymmetries and uncompetitive consumer behaviour. I analyse how these market failures can subvert the expression of altruistic preferences and prevent markets from reflecting the public’s concern for farm animal welfare, before considering how policymakers can address these market failures.

I conclude that preference satisfaction theories of utility and welfare do not provide a suitable grounding for economic farm animal welfare policy, which should instead seek to ensure that public values are appropriately represented in society. I develop a policy framework that draws upon public values and facts about farm animal welfare in society to assist policymakers in this work.

Where the public is almost universally opposed to certain husbandry practices, government intervention to directly protect farm animal welfare is likely to be in almost everyone’s interests, and animal welfare should be viewed as a public good. Where significant groups of people are opposed to the use of certain practices, market intervention through externality policies may be justified as a means of affording greater representation to the public’s altruistic concern for farm animal welfare: farm animal welfare should be viewed as a merit good in these cases. Where public values are largely represented in society, farm animal welfare should be treated as a private good: policymakers should empower consumers to express any further dissatisfaction towards animal agriculture through market mechanisms.

Valuing Farm Animal Welfare in a Market Economy

A Philosophical Study of Market Failure

Marcus Ventin

Submitted for the Degree of Doctor of Philosophy

Department of Philosophy

Durham University

2020

Contents

List of Figures	6
Statement of Copyright.....	7
Acknowledgements.....	8
Introduction	10
1. The Lives (and Deaths) of Others	10
2. Talking About Talking About Animal Welfare.....	14
3. Overview.....	16
Chapter 1 – The Fundamentals: Utility and Efficiency	18
1. Introduction.....	18
2. Classical Utilitarianism.....	19
3. Preference Utilitarianism.....	21
a. History	21
b. Preferences	22
c. Types of Preference.....	23
d. Utility Functions	33
e. Welfare.....	34
f. Social Utility.....	37
4. Efficiency	40
a. Pareto Efficiency.....	41
b. Kaldor-Hicks Efficiency	42
Chapter 2 – The Actors: Consumers, Producers and Farm Animals	48
1. Introduction.....	48
2. Price Motivations and Farm Animal Welfare	49
a. Pareto Efficiency and Farm Animal Welfare.....	50
b. Kaldor-Hicks Efficiency and Farm Animal Welfare.....	57
3. Consumer Preferences for Farm Animal Welfare.....	60
4. Willingness to Pay for Farm Animal Welfare	61
5. Summary.....	62
Chapter 3 – The Imperfect Stage: Markets and their Failures.....	64
1. Introduction.....	64
2. Free Markets	64
3. Perfectly Competitive Markets.....	66
4. The Problem with Perfectly Competitive Markets	68
5. Summary.....	69
Chapter 4 – Externalities	71
1. Introduction.....	71
2. Externalities	71

3.	Externalities in Animal Agriculture	74
4.	Externality Policy Aims	77
5.	Taxes.....	78
6.	Consumer Subsidies.....	85
a.	Product Subsidies	85
b.	Charity Subsidies	90
7.	Cap and Trade.....	92
8.	Market Creation	94
9.	Summary.....	102
Chapter 5 – Public Goods		104
1.	Introduction.....	104
2.	Public Goods.....	105
a.	Excludability and Rivalry	105
b.	The Free-Rider Problem	106
c.	Optimal Supply	108
3.	Farm Animal Welfare as a Public Good	111
4.	Farm Animal Welfare as a Public Good?	112
5.	Legislation.....	118
6.	Taxes.....	124
7.	Summary.....	127
Chapter 6 – Information Asymmetries		129
1.	Introduction.....	129
2.	Farm Animal Welfare as a Product Attribute	130
3.	Welfare Labelling.....	132
4.	The RSPCA Assured Scheme	134
5.	Textual Labels	141
6.	An Effective Welfare Label?	145
a.	Tiered Labels	145
b.	Access to Information about Labelling Standards	147
c.	Limitations.....	148
7.	Summary.....	150
Chapter 7 – A Trip to the Nudge Supermarket.....		152
1.	Introduction.....	152
2.	Irrational Preferences.....	153
3.	Thinking, Fast: Choice Environments and System 1 Processes	156
4.	An Introduction to Nudge.....	157
5.	System 1 Nudges	160
6.	Thinking, Slow: Choice Environments and System 2 Processes	164
a.	Schema Theory.....	165

b.	Economic and Citizen Schemas	169
c.	Schemas and Preferences for Farm Animal Welfare	171
7.	System 2 Nudges	173
a.	Social Norms.....	174
8.	Summary.....	181
Chapter 8 – Conclusion: Building an Economic Farm Animal Welfare Policy Framework		183
1.	Introduction.....	183
2.	The Need for a Policy Framework	184
3.	Cracks in the Foundations	187
a.	Social Utility.....	187
b.	Efficiency	189
c.	Welfare.....	190
4.	Creating a Values-Based Policy Framework	191
a.	Public Goods.....	193
b.	Merit Goods.....	194
c.	Private Goods	196
5.	Conclusion	197
Appendices.....		199
Appendix I – A Comparison of Prices of Free Range and Caged or Barn Eggs.....		200
Appendix II – A Graphical Explanation of Externalities		201
Appendix III – Externality Taxes and Subsidies.....		205
Appendix IV – Nudges in the People’s Supermarket.....		209
Bibliography		214

List of Figures

Figure 1 – Ordinal Utility Function.....	34
Figure 2 – Comparison of Utility Levels.....	37
Figure 3 – The Role of the Price Mechanism at Non-Equilibrium Prices.....	65
Figure 4 – Aggregating Demand for Private Goods.....	109
Figure 5 – Aggregating Demand for Public Goods.....	110
Figure 6 – The RSPCA Assured Label.....	135
Figure 7 – The GAP Animal Welfare Certified Label.....	146
Figure 8 – Explaining the Animal Welfare Certified Label.....	149
Figure 9 – The Dreaded Norman Door.....	158
Figure 10 – Using Discounts to Grab Online Shoppers’ Attention.....	163

Statement of Copyright

The copyright of this thesis rests with the author. No quotation from it should be published without the author's prior written consent and information derived from it should be acknowledged.

Acknowledgements

Despite there only being one name on the cover, this thesis could not have been completed without the assistance of so many others. First, I cannot express enough gratitude to my primary supervisor, Simon P. James: you helped me develop the idea for this thesis from the very beginning, and put up with me even as I veered wildly off-course and the project turned into an entirely different beast. Your patience, advice, feedback and superhuman eye for a stray hyphen have been invaluable, and I am forever grateful.

I would also like to thank my secondary supervisors, Julian Reiss and Nancy Cartwright, whose feedback and insight have helped refine and shape this thesis into something that (hopefully) resembles a coherent body of work.

I gratefully acknowledge the generous funding provided by the AHRC Northern Bridge Doctoral Training Partnership, which enabled me to complete the project while living on a diet more varied than beans on toast. Their support has provided opportunities for development and travel that I had never imagined possible while doing a philosophy PhD.

Several people kindly offered their unique perspectives which informed the analysis in various areas of the thesis. I would like to thank John Malone at the Food Standards Agency, who generously agreed to host me at a slaughterhouse and was a fount of knowledge on animal health-related issues. In addition, many thanks to the staff at the People's Supermarket in Holborn, who allowed me to explore their innovations in supermarket choice architecture and provided insights into the effectiveness of their changes.

Last, but by no means least, I cannot offer enough thanks to the many, many friends from all walks of life with whom I have had the pleasure of sharing my time in Durham. Without your kindness and support, I no doubt would have finished this thesis eighteen months ago.

To my parents, who made it all possible.

Introduction

1. The Lives (and Deaths) of Others

Animal agriculture in the UK is big business. Despite the recent rise in self-reported animal welfare-conscious lifestyles such as vegetarianism, veganism, pescatarianism and flexitarianism, consumption of animal products remains the norm for the overwhelming majority of people in Britain: in one recent survey, 97% of respondents claimed to have eaten animal products at some point in 2018 (Waitrose and Partners 2018: 6-7). The animal product industry is correspondingly large: the processed poultry and red meat (£5.648bn), unprocessed poultry and red meat (£6.635bn), milk, milk drinks and cream (£4.118bn), cheese (£2.561bn), egg (£1.035bn) and fish and seafood (£3.684bn) sectors produced sales of an estimated £23.681 billion in 2017 (Mintel 2017a: 10; Mintel 2017b: 11; Mintel 2018: 11; Mintel 2017c: 10; Egg Info 2019; Mintel 2017d: 11).

This combination of widespread animal product consumption and a large animal agriculture industry entails that animal products have a significant impact upon the UK. Animal agriculture was responsible for 10% of the UK's greenhouse gas emissions in 2017 (DBEIS 2019), and many writers have proclaimed that, as a society, we need to reduce our consumption of animal products, in particular red meat, if we are to avoid a climate catastrophe (EAT-Lancet Commission 2019: 21; WRI 2019: 2; Harwatt et al. 2019; Gulland and Wilson 2019; Milman 2018; Carrington 2018). Additionally, the regular consumption of some animal products, namely red and processed meats, has been linked with chronic illnesses such as heart disease and various cancers. Treating these illnesses comes at a cost to society – one study estimated these costs could be as large as £5 billion per year by 2020 (Springmann et al. 2018a).

Then there are the animals themselves. A large animal agriculture industry must raise and slaughter huge numbers of animals to satisfy consumer demand for animal products. In 2018, producers in the UK raised approximately 10,004,000 cattle, 4,969,000 swine, 181,818,000 poultry and 34,832,000 sheep, as well as millions of farmed fish, the numbers of which are not counted (DEFRA 2019a: 6).

Many consumers want more from their animal products than just something to eat, however: they care about farm animal welfare and want the animals that produce their food to die good deaths and live good lives, or at the very least not bad ones (Eurobarometer 2016: 10, 13; Brook Lyndhurst 2010: 40; Harper and Henson 2001: 19). After a cursory examination of modern

husbandry practices, though, you could be forgiven for concluding that producers are failing to meet this demand: mutilations, deprivations of the ability to perform natural behaviours and selective breeding-induced health issues are the norm in intensive production rather than the exception (FAWC 2009: 8; World Animal Protection 2019; Wasley et al. 2017; Safran Foer 2010: 126).

I first became aware of the apparent rift between the demand and supply of animal welfare as a fresh-faced eighteen-year-old on the brink of starting university. I was the first of three children to move any meaningful distance away from home, and my mother's response to her impending empty nest was to make the metaphor reality and replace her departing child with actual birds – three chickens.

These were no ordinary chickens, however: these were hens with a history. When hens used in industrial egg production reach about two years of age, the rate at which they lay eggs begins to decline. They become less productive to the extent that it is no longer economically viable for producers to maintain them. These hens – who, in the UK, number more than 40 million per year – are sent to slaughter (HSA 2014). Charities such as the British Hen Welfare Trust (BHWT) rescue some of these hens (the BHWT takes in over 60,000 per year (BHWT 2019)) and find homes for them as pets. We adopted three of these ex-layers.

We received the hens shortly after they had been collected from the farm. The condition they arrived in was therefore likely typical of the millions of hens who have reach the ends of their productive lives in these systems. The hens were in a bad way, to put it mildly. One was almost entirely bereft of feathers. Another sat dazed in the garden for days, simply staring at the sky as if the sight of the sun was a novelty to her.

But they improved. With time, the feathers returned, the sight of the sky became familiar and freedom to roam became the norm. As the scars of their previous lives began to heal, the hens' personalities emerged. They could be by turns affectionate, curious and bossy, with unique preferences for people and parts of the garden. They would happily jump up onto a knee for a stroke, or the opportunity to peck at an unsuspecting victim's ice cream. If the back door was left open, they would migrate to the house and surprise everyone by turning up to dinner unannounced with an indignant cluck for having not been invited. In short, it was clear to me that there was a spark behind the eyes, that these birds, who had suffered such hardships in order to satisfy human demand for their eggs, were both mentally and emotionally complex and deserved better than they had been given.

These hens were part of a lucky minority whose stories end with a happily ever after, but the overwhelming majority of egg-laying hens, and other farmed animals, are not so fortunate. They are sent to slaughter. That is the point: their deaths are an essential part of the production process. As part of a public policy internship carried out in summer 2019, I visited a cattle slaughterhouse to witness how the lives of most farmed animals are brought to an end. This slaughterhouse was something of an anomaly in the sense that each step in the production process was carried out by hand – industrial-scale slaughterhouses tend to rely more on machines than manual labour. The slaughterhouse prided itself on the quality of care shown to its animals: its commitment to manual labour allowed greater attention to be paid to the animals' welfare, while they also required cattle to be bussed in from local farms to reduce transport-related stress. In addition, the slaughterhouse made sure to adhere to the many regulations surrounding meat production: it employed Official Veterinarians and Meat Hygiene Inspectors from the Food Standards Agency to ensure that animal welfare legislation was being followed, that the animals were healthy and good food hygiene was being practised.

The slaughter process begins with a cow being herded into a small metal pen, in which it has little space to move, and none to turn around. A worker stood above the pen then applies what is referred to as a 'permanent stun', a bolt gun shot between the eyes. The term 'stun' is something of a misnomer: when applied correctly, part of the cow's brain is destroyed, instantly rendering it unconscious and causing a cessation of breathing. It is no longer capable of thought or feeling. Legally, the animal cannot yet be declared dead, although for all intents and purposes it is. It is only after the next stage, where the cow is hoisted up by a hind leg, its throat slit and left to drain for half a minute, that it officially dies. At this point, a swarm of industrious hands get to work. The skin is removed, the ribs sawed apart, the innards carefully removed, and the animal carved into two to be hung out to dry. The whole process, from living creature to something that very closely resembles food, takes perhaps three minutes per animal.

The cattle die quickly. I would be reluctant to say, however, that they die well. Some met their end with stoic indifference, but others put up a fight. They were vocal in their resistance and refused to enter the slaughter pen for as long as they could. When finally forced in, they would kick out at the door to prevent its closure. But the slaughterman always wins. Eventually, the cow tires and is stunned, slaughtered and butchered like the millions across the country that came before and the millions that will inevitably come after.

When casually browsing the homogenous packaged steaks, burgers and mince in the supermarket's freezer counters, the consumer learns nothing of the beings that produced them.

They do not know whether the animal that will provide their dinner clung to its life for as long as it possibly could, that it staved off the butcher's cleaver with a fiery passion – a quality which, in humans, is lauded and has historically been turned into the stuff of legend. In fact, these little packets of meat provide practically no indication that they were once parts of living, breathing creatures: we even use different terms to refer to the flesh, beef, and the animal itself, cow, further facilitating consumer disengagement from the lived experience of the animal.

Many consumers would be horrified, or at least severely discomfited, to learn about the lives and deaths of farmed animals. Surveys repeatedly indicate that there is substantial public demand for farm animals to live good lives and to die good deaths, but my own experiences make me think that this demand is not being satisfied. In technical terms, I contend that markets are providing a socially suboptimal level of farm animal welfare.

The majority of this thesis is dedicated to supporting this claim. My argument proceeds in several stages. I begin by considering the circumstances in which improving farm animal welfare might be beneficial to society. I argue that, for the economist, welfare improvements are most justifiable when there is sufficient consumer demand for them, and I note that there appears to be considerable concern for farm animal welfare in the UK. I then propose that four factors might be limiting the effectiveness of animal product markets in accommodating consumers' preferences for farm animal welfare, leading to a socially suboptimal provision of farm animal welfare. I analyse these factors in greater detail and explore the effectiveness and appropriateness of numerous policies in rectifying these market failures.

In conducting this analysis, I reveal some fundamental problems with preference satisfaction theories of welfare, which have played a central role in many of the leading economic studies in this field (see Lusk 2011, McInerney 2004, and Harvey and Hubbard 2013 for instance). These problems render many of the tools of conventional economic analysis largely inappropriate and ill-suited to accounting for concern for farm animal welfare in a market environment. To rectify these issues, I propose a new way of thinking about the objectives of economic farm animal welfare policy that moves away from the prevailing and problematic preference satisfaction foundations of conventional economic analysis and instead draws upon public *values* and facts about animal welfare in a society. Such a shift allows economic farm animal welfare policy to better recognise different ways of valuing farm animal welfare, and therefore promotes policies that better reflect people's attitudes towards farm animals.

Although this project is interdisciplinary in the sense that it engages with concepts and theories from disciplines as various as philosophy, economics, animal ethology, marketing, psychology,

and public policy, the analysis is primarily philosophical in nature. Food policy is a sprawling and multidisciplinary area of study, and so engaging with a range of fields is necessary to build support for my conclusion that a values-based framework is a more appropriate way of approaching economic farm animal welfare policy than traditional preference-based approaches.

2. Talking About Talking About Animal Welfare

Given that this is a thesis about animal welfare, it is necessary to provide a definition of this concept. In line with the definition of human welfare that I adopt in Chapter 1 (p.34), I understand animal welfare as referring to how well (or poorly) an animal's life is going for it. The welfare of farm animals is likely to be influenced by factors including production inputs (feed, shelter, etc.), production outcomes (health outcomes, mental states, opportunities to express certain natural behaviours, etc.) and the husbandry practices that an animal is exposed to.

When we talk about animal welfare, though, much of what we say contains dimensions that are not captured by this definition: we might, for example, want to make normative claims about what constitutes good and bad welfare, or relative claims about what constitutes higher and lower welfare. Normative and relative claims about animal welfare are of central importance to this thesis and so it is important to explain what they refer to in my analysis. Normative claims about farm animals and their welfare can be found in abundance in the field of animal ethics, and I will discuss the works of several animal ethicists in this thesis. I am, however, primarily interested in how economic policy can be used to create a society that reflects public concern for farm animal welfare. Whether this is a morally desirable end, as it can promote improvements in farm animal welfare, or a morally undesirable end, as it implicitly condones both the continued treatment of farm animals as human property and their slaughter to satisfy human appetites, is a question that is not addressed here. My own views on the permissibility of treating animals as property, slaughtering them for human consumption, and inflicting suffering upon them in pursuit of greater productivity are irrelevant to this project and thus remain unstated in what follows. For the purposes of this thesis, I will instead adopt a stance popularised by the retailer Harry Gordon Selfridge, and hold that the consumer is always right¹ – normative claims about animal welfare will be understood in terms of consumers' beliefs.

¹ Well, maybe not *always*.

I adopt a parallel stance towards relative claims about animal welfare. Making relative claims about farm animal welfare is not always straightforward, as elements of welfare can conflict with each other, casting uncertainty over what is best for an animal. Things that, in isolation, we might think are good for an animal's welfare might come at a cost in other areas of welfare (for instance, outdoor access and freedom to roam can expose an animal to harmful pathogens and leave it at risk of attack and predation). The reverse is also true: things that, in isolation, we might consider detrimental to an animal's welfare could make a positive contribution in other areas of welfare (mutilations such as beak trimming and tail docking can inhibit cannibalistic behaviours, for example). In an attempt to enable us to make relative claims about farm animal welfare, researchers have sought to synthesise various factors into models that produce an overall welfare score that captures these trade-offs (see De Mol et al. 2006's FOWEL model of laying hen welfare, for instance), but these models have been subject to criticism (Cotra 2017; Bollard 2016). Because there is not currently a generally-accepted model upon which to base relative claims about animal welfare, and given that this thesis is chiefly interested in how *markets* influence farm animals' welfare, I will primarily understand relative claims in terms of consumer *beliefs* about what constitutes higher and lower welfare: these beliefs inform consumers' preferences for particular welfare attributes, and these preferences can in turn motivate their actions in the market.

In line with these consumer-centric definitions of normative and relative claims about farm animal welfare, I hold that the satisfaction of public preferences for farm animal welfare is an appropriate policy aim. There can be other appropriate policy aims in this area, namely ethical concern for the lives of farm animals, but, for pragmatic reasons, my focus in this thesis is on economic farm animal welfare policy. Ethically motivated calls for a broad proscription on the consumption of animal products have fallen on deaf ears for decades. Engaging with the policy debate from an economic perspective, however, and highlighting areas where consumers' voices are not being heard in the market, might provide a more viable way of influencing policy and improving the lives of farm animals.

Consequently, I will assess the desirability of a state of the world, and the appropriateness of the provision of farm animal welfare, according to how far it reflects what matters to people. For the majority of this thesis, this will be understood in terms of preference satisfaction, although, due to the limitations of this approach – which will be outlined in my final chapter – I will contend that policymakers should instead seek to engage with public values.

3. Overview

The structure of this thesis is as follows. In **Chapter 1**, I introduce two key concepts that underpin welfare economics: utility as preference satisfaction and economic efficiency. I unpack the different types of preference that can motivate demand for farm animal welfare, and contend that, because many consumers dissociate themselves from the ultimate fate of farm animals, moral concern can stimulate purchases of higher-welfare products.

In **Chapter 2**, I assess the extent of consumer demand for farm animal welfare in the UK. I argue that, within the production space bounded by the UK's farm animal welfare legislation, welfare improvements tend to be detrimental to productivity and thus lead to price increases, meaning that there is no universal reason for UK consumers to prefer higher welfare standards. The question of whether there is demand for higher farm animal welfare should therefore be treated as an empirical one, and I argue that recent surveys provide good evidence that UK consumers are prepared to pay more for higher-welfare products and production practices.

In **Chapter 3**, I explore the market, the institution which theoretically reconciles consumer demand with supply. I explain why ideal markets are a highly efficient means of distributing resources, before outlining how many real-world markets, including markets for animal products, fail to live up to this perfectly competitive ideal.

In the following four chapters, I focus on the forms of market failure that are especially relevant to demand for farm animal welfare, and consider how policy can be, and in many cases has been, used to rectify market inefficiencies. **Chapter 4** concerns externalities, **Chapter 5** looks at public goods and **Chapter 6** examines information asymmetries. In **Chapter 7**, I draw on insights from behavioural economics and consumer psychology to make the claim that the environment in which consumers most frequently make decisions affecting farm animal welfare – the supermarket – primes them to discount other-regarding preferences for animal welfare in favour of more self-interested ones. I then discuss how consumers can be nudged to express other-regarding concerns for farm animal welfare.

Finally, in **Chapter 8**, I consider what economic farm animal welfare policy should seek to achieve, and propose a policy framework to facilitate this. I contend that preference satisfaction accounts of welfare are not suitable guides with which to assess the efficacy of interventions in animal product markets: what consumers prefer in a supermarket context may not be what they prefer in other types of choice environment. Instead, I propose that policy in this area should be guided by people's values, rather than their preferences, with the aim of ensuring that different ways of valuing farm animal welfare are afforded appropriate representation in society.

I proceed to argue that policymakers should consider facts about both farm animal welfare in society and public values to determine whether farm animal welfare should be viewed primarily as a public, merit or private good. Consequently, policy aims evolve according to the extent of public concern for farm animal welfare. I link these different classifications of farm animal welfare to the market failures developed in Chapters 4-7, and argue that the form of market failure that policymakers should be most concerned with, and thus the animal welfare policies that they should concentrate their energies upon, is primarily determined by whether farm animal welfare is viewed as a public, merit or private good. I thus place public values and facts about animal welfare in society at the heart of economic farm animal welfare policy.

Chapter 1

The Fundamentals: Utility and Efficiency

In this Chapter

- A brief discussion of classical utilitarianism, and why it is ill-suited to welfare economic aims.
- A detailed account of preference utilitarianism.
- A discussion of the kinds of preferences which will be relevant to this thesis, including support for the claim that ethical concern can motivate purchases of higher-welfare animal products.
- An outline of the problem of interpersonal utility comparisons.
- An introduction to the concept of efficiency as a means of making interpersonal utility comparisons.
- An overview of the Pareto and Kaldor-Hicks accounts of efficiency.
- A defence of the claim that the Kaldor-Hicks account can identify improvements in efficiency only when compensation is actually paid.

1. Introduction

Before examining the numerous ways that markets affect the expression of preferences for farm animal welfare, it is useful to introduce the economic concepts which underpin this thesis. This chapter concerns utility and efficiency, two concepts which, by enabling economists to determine whether policies are beneficial to society (Blackburn 2008: 375), are essential to welfare economic policy analysis.

I begin with an outline of the multi-disciplinary concept of utility. In line with mainstream welfare economics, I will adopt a preference satisfaction interpretation of utility, and will briefly outline why classical or hedonistic utilitarianism is ill-suited to the discipline's needs. I then develop my account of preference utilitarianism by exploring the assumptions that underpin preferences. Next, I list the types of preferences that are relevant to this thesis' project, and argue that ethical concern can motivate purchases of animal products. After this, I discuss whether preference utilitarianism provides an acceptable account of human welfare and explore its limitations in enabling interpersonal utility comparisons.

In order to overcome the problem of interpersonal utility comparisons, I introduce the second fundamental concept that underpins this thesis, efficiency. I explore the Pareto and Kaldor-Hicks accounts of efficiency, and consider how far they enable economists to determine whether a policy benefits society. I conclude that the Pareto account is ill-suited to dealing with the interests of large numbers of people, while the Kaldor-Hicks account offers a viable means of measuring improvements in social utility, provided that compensation is actually paid to people left worse-off by policies.

2. Classical Utilitarianism

The concept of utility originated in philosophy, and lies at the heart of the ethical theory of utilitarianism. The doctrine was first explicitly formulated by Jeremy Bentham in his 1789 *An Introduction to the Principles of Morals and Legislation*, although elements of the theory can be traced at least as far back as Epicurus in the late 4th century BCE (Russell 2004: 233). Bentham defined utility as the usefulness of a thing in bringing about pleasurable consequences, describing it (2004: 66) as:

[T]hat property in any object, whereby it tends to produce benefit, advantage, pleasure, good, or happiness ... or ... to prevent the happening of mischief, pain, evil, or unhappiness to the party whose interest is considered.

Bentham (2004: 65) further linked utility to the foundations of morality through the claim that what is pleasurable is morally good and actions that bring about pleasure are morally right, while what is painful is morally bad and actions that cause pain are morally wrong. This interpretation of utility forms the basis of *classical utilitarianism*, and, when linked to welfare, theories of *hedonism*.

Classical utilitarianism also has a long history in the field of animal ethics, largely thanks to Bentham, who famously stated in 1789 that the relevant question regarding the moral considerability of beings is not ‘Can they *reason*? nor Can they *talk*? but, Can they *suffer*?’ (quoted in Singer 2015a: 7). Bentham was not, however, the first to justify compassion towards animals in terms of the capacity to experience pleasures and pains, referred to in the modern literature as *sentience*. As early as 300CE, the Roman philosopher Porphyry claimed that animals deserve moral consideration because of their capacity to suffer (Fraser 2012: 190). Even in Britain, Bentham was not the first to argue that animal sentience warrants moral concern. More than 60 years before Bentham, the scholar William Wollaston (1725: 139) wrote: ‘[i]t is grievous to see or hear (and almost to hear of) any man, or even any animal whatever, in *torment*’. Writing 13

years before Bentham, British author Humphry Primatt declared (1776: 7-8) '[p]ain is pain, whether it be inflicted on man or beast; and the creature that suffers it... being sensible of the misery of it whilst it lasts, suffers *Evil*'. Further instances of the argument that animal sentience demands moral consideration are found in the works of other 18th century thinkers, including the poet Christopher Smart, author John Lawrence, philosophers Jean-Jacques Rousseau and John Oswald, and clergyman James Granger, all of whom predate Bentham (Thomas 1984: 175-76).

Classical utilitarianism was also a key component of the theories advanced by 19th century economists such as F.Y. Edgeworth, W.S. Jevons, Léon Walras and Alfred Marshall (Edwards 1954: 383; Heilbroner 2000: 172-77). Today, however, classical utilitarianism plays a much-reduced role in the discipline, where it is primarily used to explain decision-making in situations of risk and uncertainty (Edwards 1954: 391-2). The decline in classical utilitarianism's influence in economics can be attributed to two factors. First, there are occasions where we seek to bring about certain states of affairs independent of the pleasures and pains they bring us (Harsanyi 1977a: 644-45): we might be motivated by factors such as a concern for others, a sense of duty or respect for tradition. Consequently, classical utilitarianism cannot account for many common human motivations. Second, the theory has struggled to support the claim that we can measure utility (High and Bloch 1989: 354; Schumpeter 2006: 1028). Bentham and many other classical utilitarians treated utility as a calculable quality that can be measured in *hedons* or *utils* (Henson 1971: 323; Blair 1974: 124). This 'enabled' them to determine the right action by summing each option's anticipated pleasures and pains and choosing the one that maximised pleasure. The measurement of hedonistic utility remains, however, a purely theoretical enterprise: we lack both an objective unit of pleasure and pain, and a 'hedonimeter'-like device with which to measure it (Edgeworth 1881: 101; Colander 2007: 216; Feldman and Serrano 2010: 13). Without an objective and measurable unit of pleasure, statements of the form 'policy *y* yields *x* hedons' cannot be justified (Köbberling 2004: 375).² This issue is problematic for this thesis, as I will be considering whether society can be made better-off by improving farm animal welfare. As classical utilitarianism does not provide a practicable way of measuring a policy's impact, it is not suitable for my purposes, despite its long history in animal ethics and economics. Fortunately, the limited applicability of classical utilitarianism to policy discussions does not

² Other measures, notably best-worst scaling, may overcome these difficulties by allowing researchers to calculate the extent to which one option is preferred relative to another. As best-worst scaling examines the strength of preferences relative to other preferences, however, it does not rely upon the existence of an objective unit of pleasure. Despite the promise of this technique, I will nonetheless move away from classical utilitarianism, in line with much of economics.

derail this project. Over the centuries, utilitarianism has been adapted into new forms, one of which, preference utilitarianism, has assumed an influential role in both economics and animal ethics and is better suited to my aims.

3. Preference Utilitarianism

a. History

Preference utilitarianism rather self-explanatorily defines utility in terms of the satisfaction of preferences or desires (Qizilbash 1998: 58; Hare 1981: 91). In short, what satisfies a person's preferences is utility increasing. This theory is better-suited to economic and policy analysis than classical utilitarianism, because its basic unit – the preference – can generally be identified either by observing people's behaviour or by asking them. As with its classical counterpart, preference utilitarianism also maintains a respect for individuals by defining the good in terms of what matters to each person rather than in objective terms.

Preference utilitarianism appears in economics as early as 1821, when J.B. Say (1971: 62) defined utility as the 'inherent fitness or capability of certain things to satisfy the various wants of mankind'. In 1906, Vilfredo Pareto demonstrated that preference utilitarianism could provide a viable foundation for economics when he showed that indifference curves, a cornerstone of classical economics – economics predicated on classical utilitarianism – could be replicated using preferences (Pareto: 2008). His work was developed further by J.R. Hicks and R.G.D. Allen (1934a; 1934b), who showed that, from this preference utilitarian foundation of indifference curves, we can derive all of the claims that classical economics makes about consumer demand (Edwards 1954: 384-85). With preference utilitarianism shown to be capable of facilitating the same level of economic analysis as classical utilitarianism, while bypassing its difficulties, the theory was well-placed to provide a grounding for economics.

In contrast to its place in economics, preference utilitarianism is a relatively recent addition to the field of ethics. It was first formally proposed in J.C. Harsanyi's paper 'Morality and the Theory of Rational Behaviour' (1977a). Harsanyi connects economic theories of rational behaviour with the works of three moral philosophers. From Bentham, he takes the idea that the basic criterion of moral action is to maximise utility; Harsanyi's concept of utility diverges from Bentham's, however, by drawing from economics and defining utility in terms of preference satisfaction. From Adam Smith's *Theory of Moral Sentiments*, he takes the idea that adopting a moral viewpoint involves acting as an impartial but humane and sympathetic

observer. Last, Harsanyi draws from Immanuel Kant's *Groundwork of the Metaphysics of Morals* when he recognises that moral laws ought to be universalizable.

Preference utilitarians have also made large contributions to animal ethics. No discussion of the theory in this field can ignore Peter Singer, whose seminal 1975 book *Animal Liberation* brought animal welfare issues, explained through the lens of preference utilitarianism, to the public's attention. Although Singer's work predates Harsanyi's by two years, he did not yet explicitly identify as a preference utilitarian. The theory clearly underpins Singer's thought, however: he argues that sentience is sufficient for possessing interests, and we have a duty of equal consideration to all beings that hold interests (Singer 2015a: 7-8; Rowlands 2009: 39-40). Singer's work has inspired others to approach animal ethics from a preference utilitarian perspective (see Matheny 2006 for example) and has laid the foundations for the application of the effective altruism movement to the issue of animal suffering (see Singer 2015b: 137-47).

As a leading theory in both economics and animal ethics, preference utilitarianism is therefore well-placed to ground the analysis in this thesis. Before this analysis takes place, however, it is useful to consider four aspects of preference utilitarianism in greater detail: the nature of preferences, utility functions and preference utilitarian interpretations of both human welfare and social utility.

b. Preferences

Preferences are subjective comparative evaluations of bundles of goods (Hansson and Grüne-Yanoff 2018) which, in welfare economics, are assumed to satisfy three assumptions. First, they are *transitive*: if x is preferred to y and y is preferred to z , x will also be preferred to z (Fishburn 1968: 336). Second, preferences are *continuous*: if a is preferred to b , then situations appropriately similar to a are also preferred to b (Snyder and Nicholson 2012: 85). Third, preferences are *independent of context*: a preference for x over y is not affected by the presence or absence of alternatives (Hausman and McPherson 2009: 5).

Preferences can be collected in *preference sets*, which represent a person's attitudes towards various states of affairs. For instance, we could construct a preference set containing my preferences for the different sandwich fillings available at my local supermarket, which, when a , b , c and d represent different fillings, might look like this:

$$a > b > c > d$$

Preference sets are assumed to be *complete*, meaning that agents have ranked, or at least *can* rank, all possible alternatives (Edwards 1954: 381).

c. *Types of Preference*

i. Self-Interested Preferences

Preference utilitarians recognise that human behaviour can be motivated by many types of preferences, one of which is self-interest. Self-interest refers to what a person believes will benefit them the most and has long been viewed as a particularly strong motivator of economic behaviours, with Adam Smith (1999: 119) famously proclaiming:

It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest. We address ourselves, not to their humanity but to their self-love, and never talk to them of our own necessities but of their advantages.

It is self-interest that motivates producers to make goods, as doing so provides an income with which they can purchase goods to satisfy their own preferences. Similarly, as consumers, many of our purchases are motivated by considerations of how a good will benefit us.

Preferences for farm animal welfare can be self-interested in nature, as many consumers associate it with a range of quality attributes (Brook Lyndhurst 2012: 21; Matthews 1996: 41-42; Jago et al. 2000: 163-64; Harper and Henson 2001: 21-22; Verbeke et al. 2010: 285; Heng et al. 2013: 425; RSPCA 2006: 23). One such attribute is physical quality: some consumers believe that animals raised in higher-welfare systems will produce goods with greater physical integrity or other desirable physical attributes (RSPCA 2006: 23). This association does not always hold, however: a low-welfare broiler system will likely produce more higher-grade carcasses than higher welfare systems (Proudfoot et al. 1979), and the historical production of high-quality veal involved depriving calves of iron and severely restricting their movement (Singer 2015a: 133; Rowlands 2002: 108; Mason and Singer 1990: 12-13; Marcus 2005: 37).

Another quality attribute that consumers associate with farm animal welfare is greater food safety (RSPCA 2006: 23; Brook Lyndhurst 2012: 22). In many cases, this association does appear to hold true. Giving animals the space and conditions to express natural behaviours can reduce stress, which in turn increases the shedding of harmful bacteria and improves animals' resistance to pathogens, reducing the risk of products being contaminated by pathogens such as *E. coli*, salmonella and campylobacter (Norwood and Lusk 2013; Callaway et al. 2006: 65; Haslam et al. 2004: 293-94). The opposite also holds true in some circumstances: conditions that produce higher levels of stress, such as extreme temperatures or anxiety-inducing

conditions before slaughter, can facilitate bacterial growth and thus produce goods that carry a higher risk of causing foodborne illness (Voogd 2009).

The connection between higher welfare and product safety does not always hold, however. Chickens raised in free-range systems can experience higher rates of aggression than in cage systems. This increases animal stress, which can facilitate the growth of harmful bacteria in the final product. Raising animals in outdoor environments may enable the expression of natural behaviours, but it also exposes animals to potentially harmful pathogens from wild animals. Pathogens such as those responsible for avian influenza can mutate and spread rapidly in large flocks and herds, decreasing food safety and contributing to the spread of zoonoses. In some cases, therefore, food safety can be improved by confining animals, at the expense of welfare (Norwood and Lusk 2013; Norwood and Lusk 2011: 135-36).^{3,4}

Many consumers also associate higher-welfare goods with a superior nutritional profile, typically because they believe that more 'natural' production methods produce healthier goods. Some think that healthy feed produces healthier animals, which in turn produce healthier products; others think that higher-welfare goods are less likely than conventional goods to contain harmful and unnatural additives. Still others believe that access to free range enables animals to exercise more, which produces more muscle tissue and higher-protein products (Moran and McVittie 2008: 49; IGD 2007: 39-40; Brook Lyndhurst 2012: 22).

In certain instances, higher-welfare goods, both organic and otherwise, are nutritionally superior to their conventional counterparts. Organic, free-range and other forms of higher-welfare beef, lamb, pork, eggs, milk and chicken possess a healthier ratio of omega-3 to omega-6 polyunsaturated fatty acids than conventionally-produced alternatives, and higher-welfare beef and chicken contain less fat and more protein than conventional counterparts. Despite public perceptions, however, protein content, along with the ratio of saturated to unsaturated fatty acids, is largely unaffected by whether production is higher-welfare or conventional (Średnicka-Tober et al. 2016a, 2016b; CIWF 2012: 3-5; Lawrence 2016).

Some consumers also identify higher welfare with superior taste (Eurobarometer 2007: 78-79; Harper and Henson 2001: 21; IGD 2007: 40; Heng et al. 2013: 425; Brook Lyndhurst 2010: 26;

³ Indeed, the use of indoor systems in US pig production has been linked to a decline in swine dysentery and trichinosis (Fraser 2012: 200).

⁴ By some accounts, the effect of confinement upon overall welfare is unclear, as it requires us to understand the consequences of the trade-off between an animal's freedom to express natural behaviours and freedom from harm or disease. De Mol et al.'s FOWEL account of hen welfare (2006) indicates that typical enriched cage systems provide lower welfare than non-cage systems, but for other species and other systems, models may favour greater confinement, or, more likely, the utilisation of lower flock or herd numbers.

RSPCA 2006: 22-23; Moran and McVittie 2008: 49). This may stem from the belief that an animal with a more natural diet will produce better-tasting goods (Fearne and Lavelle 2016: 11), but a general association between higher-welfare production methods and taste is also reported (Pettersson et al. 2016: 2009). As with the connections between welfare and quality and safety, the link between welfare and taste is not consistent: blind taste tests have shown that consumers cannot taste the difference between free-range and caged eggs (Haspel 2010; López-Alt 2010), as well as pork produced at different levels of welfare (Dransfield et al. 2005: 68; Jonsäll et al. 2002: 77). This supports the already-established idea that expectations play a major role in affecting both our perceived sense of taste and neural response to taste stimuli (Just 2011: 111; McClure et al. 2004; Nitschke et al. 2006). Given the perceived connection between higher welfare and taste, it may therefore be enough for consumers to simply believe they are eating higher-welfare goods for them to enjoy a more pleasurable gustatory experience.

In short, the perceived associations between animal welfare and product quality, food safety, nutritional profile and taste provide consumers with self-interested reasons to hold preferences for higher animal welfare.

ii. Altruistic Preferences

By adopting a theory of preference satisfaction that *only* recognised self-interested preferences, we run the risk of creating a warped understanding of human behaviour. Even Adam Smith, who is often cast as believing that human beings are entirely self-interested, recognised that many actions are motivated by concern for others. His moral treatise, *The Theory of Moral Sentiments* (2002), begins:

How selfish soever man may be supposed, there are evidently some principles in his nature, which interest him in the fortune of others, and render their happiness necessary to him, though he derives nothing from it except the pleasure of seeing it.

Of particular relevance to this thesis are two models of altruism developed by K.E. McConnell (1997). His model of *pure paternalism* recognises that someone can value the quantity of services that another receives from a particular resource, even if the ‘beneficiary’ in fact feels that they have been made worse-off by this. McConnell illustrates this point with the example of parents insisting on children eating their vegetables: the parents are better-off in the knowledge that their children have eaten their greens, even if the children themselves feel worse-off. In contrast

to pure paternalism, models of *mixed paternalism* acknowledge that people can derive value not only from the provision of the resource itself, but also from the beneficiary's gain.⁵

Preferences for farm animal welfare can be altruistic in nature. Below, I will consider ethical concern for farm animal welfare as a subset of altruistic preferences. Beyond this, some consumers hold altruistic preferences for the welfare of certain types of animals, such as ones perceived as cute (Zickfeld et al. 2018), that may be motivated by mere liking.

Of the two models of altruism outlined by McConnell, the model of pure paternalism will likely be of most use to my analysis, as many consumers of higher-welfare goods are driven by a desire for farm animals to experience certain goods, such as access to free range, without understanding how these goods actually affect welfare. In fact, many consumers will deliberately avoid or otherwise ignore detailed information about the impact of husbandry practices upon animal welfare, as this can cause feelings of guilt and discomfort (Harper and Henson 2001: 17; Miele 2010: 4).

Instead, engagement with animal welfare generally occurs on an emotional level. This claim is supported by a range of studies which assess consumer willingness to pay for animal welfare improvements by asking consumers how much they would be prepared to pay to move 1, 100 or 1,000 animals into higher-welfare housing. As the number of animals increases, willingness to pay per animal declines precipitously. In one study, participants were willing to pay an average of \$1.08 to move one laying hen from a cage system to a free-range one, but only \$0.19 per animal to move 100 hens and \$0.08 to move 1,000 hens. Similarly, participants were willing to pay \$2.85 to move one sow from a confinement crate to a shelter-pasture system, but only \$0.08 per animal to move 100 sows and \$0.02 to move 1,000 sows (Lusk and Norwood 2012: 207). Another study that examined willingness to pay to save seabirds from an oil spill found that average total contributions were largely unaffected by the number of birds affected: willingness to pay to save 2,000 birds was \$80, to save 20,000 birds was \$78 and 200,000 birds was \$88 (Desvousges et al. 1993). While diminishing marginal utility may partly explain this phenomenon – just as the pleasure gained from eating a second slice of cake is lower than that of the first, the pleasure of improving the life of the 2,000th bird is lower than that of the first – the precipitous rate of decline could be taken to support the claim that people engage with animal welfare emotionally, rather than rationally (Norwood and Lusk 2011: 208, 298). This means that they react to an

⁵ These models show that a person's utility can increase as a result of another's consumption. This is particularly relevant to the issue of farm animal welfare, where consumers' altruistic preferences may be satisfied by an animal's consumption of a bundle of goods with certain perceived welfare attributes.

image, whether it be of a hen in a cage, a sow in a stall or a seabird drowning in an oil slick rather than the number of animals involved (Kahneman 2012: 93).

The emotional nature of engagement with animal welfare leads consumers to seek cues that broadly indicate that an animal has lived a happy life (Brook Lyndhurst 2010: 40; Harper and Henson 2001: 19). One such cue is naturalness (Harper and Henson 2001: 15; Brook Lyndhurst 2010: 39), which may explain why sheep and cattle are generally viewed as enjoying some of the highest welfare standards (IGD 2007: 20): these animals can often be seen grazing on pasture. Similarly, free-range egg production is generally perceived to provide high welfare (Pettersson et al. 2016), perhaps because access to an outdoor space evokes images of natural production.

Naturalness cues do not guarantee high levels of animal welfare, however, with access to free range providing a prime example. One measure of laying hen welfare ranks access to free range as the 19th most significant element of welfare, out of 25 (De Mol et al. 2006: 160), and free-range production is compatible with harmful practices such as beak trimming and the use of high stocking densities.

Purchases motivated by a desire for animals to live more natural lives are therefore best understood through models of pure paternalism. Consumers with these motivations take satisfaction from the idea of animals living in conditions believed to be natural, irrespective of how these conditions affect animal welfare, information about which many will actively avoid (Harper and Henson 2001: 17; Miele 2010: 4). This is not to say that models of mixed paternalism are not also applicable to farm animal welfare issues. People who care about animal welfare, engage with it on a rational level and make welfare-informed purchases of animal products are likely to be mixed paternalists, as they care about the benefits received by animals rather than the mere provision of certain resources.

iii. Ethical Preferences

One type of altruistic preference that warrants further consideration is the ethical preference. Many consumers buy goods because doing so supports environmental or social causes. Some internet users use the search engine *Ecosia*, for instance, because the firm spends 80% of its profits on tree-planting initiatives in biodiversity hotspots around the world (Ecosia 2019), while other consumers actively seek out the Fairtrade logo on products such as chocolate and coffee to ensure that farmers receive a fair wage.

The claim that ethical preferences can motivate purchases of higher-welfare animal products requires detailed justification. It is clear that people can hold ethical preferences for farm animal welfare: some people care about farm animals because they believe that their welfare is an appropriate object for ethical concern. Although there are many ethical theories that can ground this belief, the philosophical literature is dominated by two. Some people have adopted a *welfarist* position, which draws heavily from utilitarianism by attributing moral significance to animal interests, which in turn are predicated upon the capacity for sentience. One version of this position is articulated by Peter Singer in *Animal Liberation*. The welfarist position states that, while we ought to consider the preferences of all sentient beings, the capacity of sentience does not endow its holder with inviolable rights.

Other philosophers disagree with Singer and the welfarists on this point and propose a *rights-based* approach to animal ethics. This position was perhaps most famously advanced by Tom Regan, who argued that a being is morally significant if it is likely to be a subject-of-a-life, i.e. whether it is a being with:

[B]eliefs and desires; perception, memory, and a sense of the future, including their own future; an emotional life together with feelings of pleasure and pain; preference- and welfare-interests; a psychophysical identity over time; and an individual welfare in the sense that their experiential life fares well or ill for them, logically independently of their utility for others and logically independently of their being the object of anyone else's interests (Regan 2004: 243).

Those beings that fulfil the criteria for a subject-of-a-life possess a right to life that humans are morally obliged to respect, and, Regan argues, we ought to err on the side of caution with respect to those beings whose capacities we are uncertain about.

A rights-based position does not have to be grounded upon the subject-of-a-life criterion, however. Gary Francione and Anna Charlton (2015) also hold a rights-based position, although, in contrast to Regan, they base it upon the quality of sentience. Their position is frequently referred to as *abolitionism*, due to their claim that sentient beings possess a basic right not to be the property of humans, which demands the abolition of animal ownership.

There are several forms of animal product production that threaten to expose divisions between welfarist and rights-based theories. Some producers achieve higher levels of welfare by giving their animals more natural diets, using lower stocking densities, providing access to natural environments and refraining from noncurative operations such as beak trimming and tail docking. Many of these producers also farm in a sustainable manner that benefits both the local and global environment (Pollan 2011: 123-238; Solotaroff 2013). Such production methods may be amenable to welfarists. Indeed, Peter Singer and author and attorney Jim Mason remain open

to the possibility that – in certain, very limited circumstances – the consumption of some animal products may be morally acceptable. In a 2006 interview with *The Vegan* magazine, Singer indicates that, if suffering is our primary concern, the occasional consumption of meat produced by animals that have lived natural lives and are humanely killed on the farm may not be objectionable (Raha 2006: 19). Singer and Mason point out that land in the Welsh hills is unsuitable for crop production, but has been used to raise lambs using traditional forms of husbandry. If the sheep live good lives, and would not have existed if they were not to be killed and eaten, Singer and Mason (2006: 255) conclude that the production and consumption of such goods may be beneficial to both humans and animals.

While meat produced in conditions of extremely high welfare may be acceptable to the welfarist, the ultimate killing of the animal will make it unacceptable to most rights-based thinkers, as, although the animals are given a much higher quality of life than in conventional production, their inevitable slaughter still violates their right to life and thus requires them to be treated as mere means to human ends. Abolitionists will also oppose extremely high-welfare products, as their production still requires sentient animals to be treated as property.

A second form of animal product production that threatens to expose divisions both within the animal rights camp and between rights-based and welfarist positions is that of cultivated (lab-grown) meat. This emerging technology grows meat from animal cells that can be collected without causing harm to the donor – the US firm JUST, for instance, can cultivate animal-free chicken nuggets using cells taken from a chicken feather (Morris and Cook 2018). The cells are then placed in a liquid called a culture medium, which induces them to multiply (proliferate) before differentiating into muscle cells (van Mensvoort and Grievink 2014: 10-11).⁶ These cells are placed around a scaffold, which encourages them to merge into myotubes and grow into muscle fibres. These fibres are then harvested and processed into a final product. Proponents of cultivated meat claim that the technology has the potential to resolve many of the welfare, health, food safety and environmental issues that stem from animal agriculture (Reese 2018: 73-94; Datar et al. 2016: 125; Mattick and Allenby 2012; Mattick et al. 2015; Shapiro 2018). Many also anticipate that cultivated meat will one day be cheaper to produce than its conventional counterpart, as the production of cultivated meat does not require energy to be expended on metabolisms, movement, digestion, bones and internal organs, as is the case in conventional animal agriculture (Post 2012: 299) – instead, resources are expended almost entirely on the growth of consumable products. For our purposes, the technology may offer a way of producing

⁶ The technology can also be used to produce fat cells.

meat without inflicting suffering or death upon animals. Despite the potential benefits of the technology, it may create divisions among animal ethicists. A welfarist is unlikely to take issue with cultivated meat, provided that any animals used to supply the cells are treated well (Singer and Mason 2006: 263). Given that cultivated meat does not require animal deaths, the technology may also win support from those supporters of animal rights who believe that farm animals possess a right to life. Cultivated meat is likely, however, to remain unacceptable to abolitionists, as the most viable forms of cell-cultivating technology require the maintenance of host animals from which tissue samples are harvested (Francione 2018; Ventin forthcoming).

Despite their disagreements on these issues, however, both welfarists and supporters of animal rights are likely to agree that the overwhelmingly majority of animal products are produced in ways that are morally unacceptable. Seemingly intrinsic elements of modern production such as high stocking densities, large flock or herd sizes and mutilations are incompatible with welfarist concerns, while the ownership of sentient beings and use of subjects-of-lives for human ends renders these forms of production incompatible with abolitionist and animal rights positions. Given that modern intensive farming is incompatible with the welfarist, animal rights and abolitionist positions, the outputs of these systems, i.e. their products, must surely also be incompatible. Adopting a rights-based or abolitionist position currently seems to require people to become vegans or vegetarians, and, given the animal suffering that permeates modern farming – including in systems perceived as being higher welfare, such as free-range eggs, where mutilations and high stocking densities are common – welfarists also seem obliged to abstain from most animal product markets. The suffering and killing of animals in modern farming thus seems to preclude the possibility that ethical concern for animal welfare can justify purchases of higher-welfare goods, and therefore gives us reason to be sceptical of the claim that ethical concern can motivate higher-welfare purchases.

Nevertheless, many consumers of animal products do claim to assign moral significance to farm animals. This apparent incompatibility between their values and behaviour is often referred to as the *meat paradox* (Loughnan et al. 2010: 156). Consumers respond to the meat paradox in many ways.⁷ Many people hold consistency to be valuable for its own sake (Cialdini 2007: 59), and awareness of inconsistency between one's behaviour and values can be a source of anxiety, known as *cognitive dissonance* (Festinger 1957: 1-2; Rothgerber 2014: 32). Because consistency is often viewed as valuable in itself and inconsistency can be a source of anxiety, some people will

⁷ This section is not intended to provide a comprehensive account of consumer responses to the meat paradox. A more detailed discussion can be found in Rothgerber 2014.

respond to perceived inconsistencies by becoming more consistent (Cialdini 2007: 73). One way that some people do this is by adjusting their values to reflect their behaviour, leading them to reduce their ethical concern for farm animals. Loughnan et al. (2010) found that omnivores who had just eaten beef jerky were more likely to attribute lower moral significance to cows and other animals than a control group that was given cashews, a finding that can be explained by the subjects wanting their attitudes to be consistent with their behaviour. Alternatively, people may adjust their actions to reflect their ethical concern for animal welfare by becoming vegetarian or vegan, or reducing their consumption of animal products: exposing test subjects to issues of animal welfare has been found to reduce willingness to eat meat (Hoogland et al. 2005; Berndsen and van der Pligt 2004).

While these findings demonstrate one way that consumers can reduce the cognitive dissonance associated with becoming aware of their personal meat paradox, it does not show that ethical concern is capable of motivating purchases of higher-welfare goods. Instead, the consistency response suggests that people must in fact reduce their ethical concern if they are to continue consuming animal products.

Another way that consumers respond to the meat paradox is by avoiding engaging with their ethical values at points of interaction, often by dissociating animal products from living animals (Harper and Henson 2001: 11). Although some consumers care about farm animals, and can experience feelings of guilt and even disgust when they think about the fact that an animal has died in order to produce a good (Kunst and Hohle 2016; Harper and Henson 1999: 17; McEachern and Seaman 2005: 582), their ethical concern may not be at the forefront of their thought at points of purchase and consumption (Safran Foer 2010: 213). Dissociation may be encouraged by environmental cues: I argue in Chapter 7 that supermarket choice architecture assists dissociative processes by nudging consumers to subdue their ethical preferences at points of purchase.

Dissociation can also have internal causes: some people engage in 'voluntary ignorance' by avoiding or ignoring information that links animal products to the living farm animal, and thus evokes feelings of guilt or disgust (Harper and Henson 1999: 17; Brook Lyndhurst 2010: 32-33; Pettersson et al. 2016: 2000; Singer 2015a: 217). Consumer engagement in voluntary ignorance is assisted by the fact that animal products are frequently processed and prepared in a way that hides their animal origins (a chicken nugget has very little in common visually with a live chicken) (Gruzalski 2004: 126). Additionally, dissociation is aided by the fact that English speakers tend to use Norman words such as beef and pork to refer to the flesh of animals that

they call by their Saxon names, cows and pigs (Ryder 2000: 93). This enables consumers to put further distance between the food on their plate and the animal that produced it. Even in the case of chicken and lamb, where the flesh and the animal share the same name, product names such as mince, burger, nugget and drumstick help consumers distance the product from its animal origins (Hoogland et al. 2005: 16; Kunst and Hohle 2016: 770-71; Rothgerber 2014: 33).

Despite initial appearances, dissociative strategies can enable consumers' ethical concerns to motivate higher-welfare purchases. This may appear counterintuitive: after all, if a consumer engaged in dissociation fails to consider his ethical beliefs at points of purchase and consumption, how can his actions be ethically-motivated?

This claim can be supported with reference to the nature of consumer engagement with farm animal welfare. We saw earlier (p.26) that many people engage with farm animal welfare on an emotional level, preferring to focus on the idea and imagery of 'happy animals' rather than the discomfort that often follows engaging with the detail of modern animal agriculture (IGD 2007: 16; Brook Lyndhurst 2010: 40). Consequently, consumers' ethical concern for farm animal welfare is often manifested through purchases of goods labelled as free-range, grass-fed, etc. that imply higher welfare and happy animals. Dissociative strategies, on the other hand, serve to distance consumers only from the elements of animal agriculture that they find discomfiting, namely suffering and the animal's ultimate fate.

It seems, therefore, that, despite the apparent incompatibility between ethical values and conventional animal agriculture, ethical concern can still motivate higher-welfare purchases. Consumers may dissociate themselves from the parts of animal agriculture – such as animal deaths, animal suffering and low-welfare production – that cause them to feel discomfort or guilt, and instead selectively focus their attention, and ethical concern, on elements they are more comfortable with, in particular the idea of happy animals living natural lives in higher-welfare systems. Their ethical preferences can thus motivate them to seek out cues that farm animals have lived good lives, while ignoring cues that remind them of suffering and death.

This kind of ethical preference can be contained within the altruistic frameworks discussed earlier. Ethical preferences for farm animal welfare may be purely paternalistic (McConnell 1997: 32), where consumers take pleasure in the knowledge that farm animals have been provided a particular set of resources, such as access to free range.⁸ Ethical motivations can also be explained with reference to a third theory of altruism called *impure* or *warm glow altruism*,

⁸ Models of mixed paternalism do not appear relevant because many consumers appear uninterested in learning about the real impact of certain resources upon animal welfare.

which recognises that many consumers experience a feeling of intrinsic satisfaction that results from a belief that their actions have made a positive difference (IGD 2007: 16; Andreoni 1990). Irrespective of whether consumers' ethical preferences are purely paternalistic, impurely altruistic or a combination of the two, it is likely that ethical motivations for purchases of higher-welfare goods can in fact operate within the bounds created by dissociation.

In sum, I have argued that ethical concern for farm animal welfare can serve as a motivator for higher-welfare purchases, provided consumers disengage from animal suffering and death in the production process. More broadly, I have proposed that preferences for welfare can be altruistic and self-interested in nature. It is important to recognise that welfare economics affords no special value to any of these preferences: seemingly arbitrary altruistic preferences are afforded as much significance as preferences motivated by ethical concern, which are in turn as significant as self-interested preferences. All are regarded simply as matters of taste. This is unlikely to sit well with ethicists, who tend to enshrine matters of ethics with a special significance that extends beyond mere taste. Nevertheless, given that this thesis is situated within the discipline of welfare economics, I will continue to subscribe to the idea that ethical preferences for farm animal welfare are held on the same level as matters of taste, although I recognise that this assumption is highly contentious and tends not to reflect the way people think about ethical concern.

Beyond self-interest, altruism and ethical concern, other writers recognise that economic agents can also be concerned with purposes, personal and social values, beliefs about what others are like and the actions that they are likely to take, natural events that may occur, as well as psychological attitudes, peer group pressures, personal experiences and the general cultural environment (Dasgupta 2005: 231; Snyder and Nicholson 2012: 86). Having a theory that is capable of recognising a range of different types of preferences allows economists to develop models that more realistically-explain and predict behaviour, and also allows preference utilitarianism to better capture what matters to economic agents.

d. Utility Functions

When we gather preferences (of all types) together in a preference set and assign them numerical values, we create a *utility function*. Utility functions are *ordinal* in nature, meaning that the numerical values assigned to preferences only indicate their positions in the ranking, rather than their relative strengths (Pindyck and Rubinfeld 2013: 80; Dasgupta 2005: 229). Because utility functions are ordinal, the assumptions of transitivity, continuity and

independence all hold under monotonically increasing transformations, i.e. transformations that preserve the order of numbers rather than the distance between them (Mandler 2006: 1114). Figure 1 provides an example of an ordinal utility function containing preferences for drinks.

Figure 1 – Ordinal Utility Function
(adapted from Reiss 2013: 38)

		Utility Ranking			
		A	B	C	D
Drink	Big Blue	2	1,002	-11.8	10,000
	Tripod	1	1,001	-11.9	1,000
	Tickler	0	1,000	-12	1

Despite the difference between utility values in Figure 1 varying between each ranking, they all indicate the same thing: the Big Blue is preferred to the Tripod, which in turn is preferred to the Tickler. Ordinal utility rankings therefore do not reveal the extent to which one option is preferred to another (Köbberling 2004: 375; Strotz 1953: 384; High and Bloch 1989: 354).

e. Welfare

When writers use preference utilitarianism to ground the claim that a policy is beneficial to an individual or a society, they implicitly link utility to *welfare*, which refers to ‘how good people’s lives are *for them*, ... how well their lives are going’ (Hausman 2012: 78; Hausman and McPherson 2009: 2).⁹ The implication is that if a policy increases a person’s utility, i.e. if it satisfies their preferences, it must therefore be to their advantage to see this policy enacted, i.e. the policy must be welfare-increasing. The identification of preference satisfaction with welfare is contentious, however. Detractors have claimed that there are at least four types of case where the satisfaction of preferences is either irrelevant or even detrimental to how well a person’s life goes for them, implying that preference satisfaction is not a good basis for welfare.

First, most people will hold at least some preferences that are not concerned with their own experience or life. Additionally, some people might hold preferences for states of the world that they will never knowingly experience (Harsanyi 1986: 3; Hausman and McPherson 2009: 6; Parfit

⁹ There are other definitions of welfare, such as Stiglitz et al.’s claim that welfare concerns what is ultimately good for people, i.e. what contributes to their quality of life (Stiglitz et al. 2010: 67). Accounts of this type tend to claim that there are objective constituents to welfare that matter independently of whether an individual cares about them. I will not explore these accounts in this thesis, as they are not relevant to the project. For a brief overview of objective list theories of welfare, see Reiss (2013: 222-25). For an example of an objective list theory of welfare, see Stiglitz et al. (2010: 67-90).

1987: 484): Aimee sees in the news that a summer heatwave has contributed to the deaths of thousands of chickens at a local farm (see Creaghan 2019 for one instance in the UK). She develops a preference that the farm installs appropriate ventilation and temperature control to prevent such an incident reoccurring. The farm in fact does take these measures, although, because this is not reported in the news, Aimee remains unaware of this. Despite Aimee's preference being satisfied, it does not appear that her welfare – how well her life goes for her – has been affected (Fletcher 2016: 13, 34; Adler and Posner 2006: 36). Instead, it is arguably the feeling of satisfaction that arises after a preference is satisfied, as opposed to the satisfaction of the preference itself, that increases a person's welfare (Parfit 1987: 494).¹⁰

Second, many people hold preferences, often based on false or incomplete information, whose satisfaction can be harmful to their welfare (Harsanyi 1977a: 645-46; Hausman 2010: 323; Griffin 1986: 12). As discussed earlier (p.23), some people hold preferences for higher welfare because they associate it with greater food safety. This association does not always hold true, however: access to outdoor spaces can increase an animal's exposure to pathogens, which may compromise food safety. People who buy higher-welfare goods for reasons of food safety may be acting in a way that is in fact harmful to their welfare.¹¹ This poses a problem for the preference utilitarian, who must seemingly hold that the satisfaction of these preferences is welfare-increasing.

Some have suggested that this criticism can be circumvented, either by only regarding the satisfaction of fully-informed preferences as being relevant to welfare, or by laundering uninformed preferences to compensate for errors in reasoning and knowledge. Laundering 'fixes' those preferences derived from false information by asking what someone who is both free from cognitive errors and fully-informed about herself and her circumstances would prefer (Railton 1986: 16; Hare 1981: 104-5; Harsanyi 1977a: 646). It may not be clear, however, when someone possesses an irrational preference (Hausman 2012: 84). As Julian Reiss writes (2013: 216), '[w]e can observe only how people actually choose, not how they would choose if they were fully informed, rational and possess infinite amounts of strength of will'. Without knowing what a rational and fully-informed version of someone would prefer and how she would behave, we

¹⁰ The same point can be made for some self-interested preferences. Behzad might prefer to be descended from Charlemagne. It turns out that Behzad is in fact a descendant of Charlemagne, yet he is unable to trace his ancestry and so remains ignorant of the fact (Hausman 2012: 85-86). Behzad's satisfied preference does not seem to affect his welfare.

¹¹ Other examples of preferences based on false or incomplete information include preferences for the consumption of large quantities of alcohol, drugs and cigarettes, or listening to loud music when not aware of the health risks.

cannot identify her ‘true’ preferences, and thus also cannot identify her ‘false’ preferences (Hausman and McPherson 2010: 129). While this is a problem regarding the measurability of laundered preferences and thus may not be philosophically fatal, it is surely relevant to the usefulness of preference satisfaction theories of welfare in applied contexts such as public policy.

Third, some preferences conflict. This often occurs when people hold second-order desires, preferences about the kinds of preferences they wish to have (Frankfurt 1971: 6-7). Upon learning about the lives of hens in cage systems, Caroline may prefer not to prefer caged eggs. However, if Caroline is weak-willed or highly price-motivated, her first-order desire for cheap caged eggs may conflict with her second-order desire to not hold a preference for caged eggs. When preferences conflict, it is not clear which should be taken to be representative of a person’s welfare, and preference satisfaction theories appear unable to provide an answer to this problem. Without an answer, it is difficult to determine which course of action – in this case, whether or not Caroline should buy caged eggs – will improve a person’s welfare (Hausman and McPherson 2009: 7).

Fourth, people can adapt their preferences to fit their circumstances. As noted earlier, preferences for farm animal welfare and animal products can be influenced by recent experience. Loughnan et al. (2010) found that people who had just eaten beef jerky were more likely to attribute lower moral significance to cows and other animals than a control group of omnivores that were given cashews, a finding that can be explained by the subjects having a desire for their attitudes to be consistent with their behaviour. Other studies have found the inverse to also be true: exposing subjects to issues of animal welfare can reduce willingness to eat meat (Hoogland et al. 2005; Berndsen and van der Pligt 2004).¹² In short, some people will adapt their preferences to support their recent behaviour and experience. This raises a problem for preference theories of welfare: perceptions of how a person’s life is going for them may be arbitrarily or inappropriately influenced by environmental factors.

In short, although the satisfaction of preferences may constitute a significant part of welfare, it does not appear to capture everything that affects how well a person’s life is going for them (Broome 1991: 4). Even though the satisfaction of preferences may not always be welfare-

¹² More concerningly, many people have grown accustomed to, and satisfied with, living in challenging conditions, often where their rights are violated and their lives are of low quality (Sen 2001: 62-63; Nussbaum 2001). Incredibly, some adjust to environments of oppression by preferring not to possess things such as liberties and protection. A preference utilitarian account of welfare could be forced to conclude that these people enjoy a high level of welfare, despite their lack of rights and liberties. In fact, a preference satisfaction account of welfare may even identify a harm to welfare if rights and liberties were provided, if people prefer not to have these things (Hausman and McPherson 2009: 8).

increasing, and some things may be welfare-increasing irrespective of how well they satisfy preferences, there is surely a strong evidential connection between preference satisfaction and welfare: the satisfaction of preferences is a good indicator that welfare has also increased (Hausman et al. 2017: 128-29). Adopting this position is useful because it does not require us to commit to any particular account of welfare; recognising that preferences and preference-satisfaction are an *indicator* of someone's welfare does not require us to claim that welfare can be entirely captured by the mere satisfaction of someone's preferences (Hausman and McPherson 2009: 16-18).

f. Social Utility

Preference utilitarianism allows us to make claims about how policies affect a person's utility: when more preferences are satisfied, utility increases, while utility decreases when fewer preferences are satisfied or preferences are thwarted. This thesis, however, is less concerned about individuals' preferences and more concerned about the preferences of large groups of people. To be suitable for this end, preference utilitarianism must therefore be able to ground meaningful claims about *social utility*.

By themselves, preference satisfaction theories only support some very basic claims about social utility. We see this when we examine the *social utility function*. The social utility function is derived from the preference rankings of every individual in a society, which are brought together using some sort of rule to form the basis of an overall social ranking, which in turn allows us to make claims about how policies affect social utility (Arrow 1963: 23).

To create a social utility function, it is therefore necessary to show that utility functions can be aggregated, i.e. that they are commensurable. If commensurable, we should be able to make comparisons between individual utility functions. One proposed way of making these interpersonal utility comparisons is to look at *utility levels*, the position that people occupy in their preference sets. Figure 2 illustrates the principle.

Figure 2 – Comparison of Utility Levels

David	Emily
a	x
b	y
c	z

Assume that David and Emily each possess three preferences about states of the world. Assume $a > b > c$ and $x > y > z$. David's a preference is satisfied, while Emily's y preference is satisfied. We can therefore say that David's preferences are better-satisfied than Emily's as he occupies a higher position in his preference set, and so possesses a higher level of welfare (Hammond 1991: 207-8).

Utility level comparisons, however, are only possible when preference sets are finite in size. Unfortunately for the preference utilitarian, there are many cases, such as when dealing with preferences about states of the world, where preference sets are infinitely large: the addition of a token sum of money or another good will surely always be preferred to a state of the world where these things are absent. When preference sets are infinite in size, we cannot identify where someone stands in their preference set,¹³ seemingly rendering interpersonal comparisons impossible (Hausman 1995: 476).

Several writers have argued that this objection is not fatal. One proposed way of overcoming the problem of infinite preference sets is to use what A.F. MacKay calls the 'mental shoehorn manoeuvre' to make judgements of *extended sympathy or empathy*. The mental shoehorn manoeuvre changes the nature of our judgements from *interpersonal* comparisons to *intrapersonal* ones (MacKay 1986: 310). Rather than asking whether Finley would be better off than Gowri in a given situation, we instead ask whether we would prefer *to be* Finley or Gowri. To make this assessment, we must distance ourselves from our personal characteristics, which would introduce non-empathetic influences to any preference for being Finley in situation x over being Gowri in situation y (Binmore 2009: 554; Hausman and McPherson 2010: 106). We must instead imagine ourselves entirely in the circumstances of the other, with his or her tastes, education, social background, cultural values and psychological make-up (MacKay 1986: 305, 308; Harsanyi 1977a: 638). If the empathetic process is successful, there will be no difference between the preferences held by the person performing the mental shoehorn manoeuvre and the subject of that process (Binmore 2009: 555). Only once we have truly adopted the perspective of another and can treat his or her experiences as if they were our own can we reintroduce our own preferences to determine whether it is better to be Finley in situation x or Gowri in situation y (Harsanyi 1977b: 59; Sen 1979: 196). From this, we can construct what Hausman and McPherson dub a 'universal extended preference ranking', a complete ranking of the preferences of different people that serves as a social utility function.

¹³ Except in those rare cases where they stand at the very top or bottom of their preference set.

Hausman and McPherson argue (2010: 107), however, that the universal extended preference ranking provides an answer to a question that is, for our purposes, irrelevant. When making interpersonal utility comparisons, the question is not ‘would I rather be Finley in situation x or Gowri in situation y ’, but rather ‘whose preferences are better satisfied?’. I might prefer to be Gowri in situation y , even though Finley’s preferences are better satisfied in situation x , simply because, for instance, I have greater admiration for Gowri. Consequently, preferences for being Finley in situation x over being Gowri in situation y may be influenced by *our own* personal preferences, which have no bearing on how well-satisfied Finley and Gowri’s preferences are. It therefore seems that the mental shoehorn manoeuvre fails to overcome the problem of infinitely large preference sets and thus does not enable us to make interpersonal utility comparisons: although it may help us to work out who we would prefer to be, it does not enable us to determine whose preferences are better-satisfied.

Even if this objection should not prove fatal, the mental shoehorn manoeuvre is surely too demanding to be useful in a policy context, where it would require us to put ourselves in the shoes of every person in a society. To do this, we would need to possess incredibly high levels of empathy and have access to prohibitively large quantities of information about every member in a society. Consequently, even if we can overcome the theoretical issues that threaten our ability to make interpersonal utility comparisons, it appears that the mental shoehorn manoeuvre does not represent a serious option for policymakers looking to assess the impacts of policy upon social utility.

Our apparent inability to make interpersonal utility comparisons makes it difficult to develop a rule that allows us to move from a collection of individual utilities to a social utility function. Without an effective universal extended preference ranking, it seems that attempts to create a social utility function must instead rely upon judgements about how much weight to afford the interests of one person relative to another. Making such judgements is problematic because any statement about the relative weighting of an individual’s utility in a social utility function requires us to make subjective value judgements that extend beyond the core assumptions of welfare economics (Robbins 1932: 123; Scanlon 1991: 17). In a two-person society, the statement ‘Hassan prefers policy x to the status quo, and the status quo to policy y , while Ida prefers policy y to the status quo, and the status quo to policy x ’ is one that can be made objectively: we can derive the preferences of Hassan and Ida by observing their behaviour or by simply asking them. Similarly, we can objectively ascertain that if x is chosen, Ida will feel worse-off than before, while if y is chosen, Hassan will experience a decrease in utility.

Preference utilitarianism does not, however, allow us to make statements such as ‘choosing policy x over policy y will increase social utility’ in an objective manner. This is because the ordinality of preference sets cannot recognise the magnitude of utility changes; the only thing we can know for sure is the direction in which Hassan and Ida’s utilities move. Should policy x be chosen over policy y , our knowledge of Hassan and Ida’s preferences only lets us say that Hassan will experience a utility increase, while Ida will experience a utility decrease – because of the difficulties we face in making interpersonal comparisons, we cannot say whether Hassan’s gain outweighs Ida’s loss. Accordingly, we cannot objectively state that choosing x over y will increase social utility.

Because we cannot identify the size of utility changes, we can only justify choosing x over y if we assign more weight to Hassan’s preferences than Ida’s. However, making this claim requires us to make a normative judgement, to say that Hassan’s gain is more valuable than Ida’s loss. Judgements about how much weight to assign to people’s interests cannot be made using preference utilitarianism alone; consequently, we cannot determine how a policy affects social utility when some people experience utility increases and others experience utility decreases.

In sum, it appears that, by itself, preference utilitarianism cannot provide a foundation for social utility functions. A basic theory of preference utilitarianism cannot justify assigning different weights to different people’s preferences, while attempts to construct a universal extended preference ranking by imagining what it is like to be other people in different circumstances seem destined to fail, as the mental shoehorn manoeuvre asks us to consider who *we* would rather be in these cases, rather than whose preferences are best-satisfied.

Without a viable social utility function, we are severely limited in the kinds of policy analysis we can conduct. When, for instance, a policy makes someone worse-off and someone better-off, we appear unable to assess how this affects social utility. Despite this, we can still make some, admittedly limited, claims about social utility. These limited claims provide the foundations of another fundamental economic concept, efficiency.

4. Efficiency

The concept of efficiency allows policymakers to assess the desirability of a resource distribution relative to alternatives, and thus provides a means of evaluating policies. Given the difficulties we have encountered in producing a viable social utility function, however, it would be easy to assume that a concept of efficiency derived from preference utilitarianism will be of little use: if

a change in resource distribution makes one person better-off and another worse-off, we appear to be unable to determine whether the new distribution is more efficient, i.e. whether it has improved social utility. This is not the case, however: the work of Vilfredo Pareto has turned a preference utilitarian-grounded concept of efficiency into a powerful economic tool.

a. Pareto Efficiency

Vilfredo Pareto was an Italian sociologist and economist whose 1906 *Manuale di Economia Politica* (Manual of Political Economy) has been hugely influential in the development of microeconomics. In this book, Pareto proposed his *Pareto criterion*,¹⁴ which has enabled economists working in the preference utilitarian tradition to identify, in certain situations, the effects of policy upon social utility, while avoiding the problems raised by interpersonal comparisons that hinder attempts to build a social utility function.

The Pareto criterion avoids the difficulties of interpersonal utility comparisons by focusing on cases where individual utilities do not move in opposing directions. It states that, when a change in resource distribution is preferred by at least one person and nobody prefers the old distribution, social utility has increased. The distribution of resources is now more efficient, and has undergone what is known as a *Pareto improvement* (Hausman and McPherson 2010: 136; Kreps 1990: 153).¹⁵ A distribution is *Pareto efficient* or *Pareto optimal* when it is no longer possible to change it in a way that constitutes a Pareto improvement (Hausman and McPherson 2010: 136). Consequently, from a Pareto-optimal distribution, nobody can be made better-off without making someone else worse-off (Stiglitz 1991: 2). It is only in those cases where some people prefer the new distribution while others prefer the old that the overall effect upon social utility is uncertain. This is because of the problem of interpersonal utility comparisons in preference utilitarianism, as discussed above, and neatly summarised by Pareto, who explains that ‘when we reach the point at which some [people’s changes in utility] are positive and others negative, we shall not be able to proceed further because we shall be favouring some individuals at the expense of others’ (Pareto 2008: 397). Consequently, the Pareto criterion cannot identify the overall effect upon social utility when some people are made better-off and others worse-off,

¹⁴ Not to be confused with the *Pareto principle*, which was first identified in Pareto’s 1896-7 *Cours d’Économie Politique* (Course of Political Economy), and states that, in many cases, roughly 80% of the effects come from 20% of the causes.

¹⁵ If, on the other hand, the old distribution of goods is preferred by at least one person, and nobody prefers the new distribution, there has been a decrease in social utility: nobody has been made better-off and somebody has been made worse-off, therefore the system has become less efficient in its resource allocation.

and so can only be used to make claims about social utility in very limited circumstances (Begg et al. 2000: 258).

This limitation surely hinders the usefulness of the Pareto criterion. While it may be easy to identify policies which benefit everyone, or at least do not harm anyone, in a two-person society, the same cannot be said in a 66 million-person society such as the UK. The more people in a society, the greater the number of interests that need to be considered when drafting policies, and the larger the likelihood of there existing conflicting interests. It is highly unlikely that policies could be developed in large societies that make some people better off without making anyone worse off, which severely limits the usefulness of the Pareto criterion as a means of assessing the impact of policy (Trumbull 1990: 203; Reiss 2013: 259). A more useful measure of efficiency must allow us to identify the overall effect upon social utility when a policy makes some people better-off and others worse-off. It was in recognition of this need that Nicholas Kaldor and John Hicks adapted the Pareto criterion to create a concept of efficiency that may be better suited for discussing social utility in larger societies.

b. Kaldor-Hicks Efficiency

The Kaldor-Hicks principle¹⁶ recognises that when people are made better-off by a change in resource allocation, making them *policy winners*, they will be willing to pay for this reallocation. This payment could be used to compensate those made worse-off by the reallocation, the *policy losers*, who would be willing to accept a certain sum of money for the reallocation to occur.

A similar idea can be found in Pareto's *Manuale*, although it was largely overshadowed by his Pareto criterion. Although Pareto argued that individual utilities are incommensurable and thus a direct transfer of utility to compensate those negatively affected by changes in allocation would be impossible, he remained open to the possibility that utility could be indirectly transferred by being converted into a resource, namely money. If utility can be converted into money, policy winners will be able to compensate policy losers. If such a transfer is sufficient to cover the utility losses, while preserving at least some gain for the policy winners, the reallocation of resources will have increased social utility: at least someone has been made better-off and compensation ensures that nobody has been made worse-off (Pareto 2008: 398; Kemp and Penzani-Christou 1999: 443).

¹⁶ Also referred to as the 'potential Pareto principle' and the 'test of hypothetical overcompensation' (Mishan 1988: xxix).

More than thirty years after Pareto's *Manuale*, Nicholas Kaldor produced his Kaldor criterion, which has much in common with Pareto's ideas. The Kaldor criterion states that a policy increases social utility if the gains of the policy winners are large enough to compensate the policy losers so that nobody is left worse-off by the change and at least somebody benefits (Kaldor 1939: 551). By itself, however, this criterion may not be a useful policymaking tool. Tibor de Scitovszky (1941: 88)¹⁷ identifies cases where introducing and then reversing a policy are both Kaldor improvements. The Scitovsky paradox occurs when the policy winner can compensate the policy loser for a change in allocation from a to b , but the policy loser can then compensate the policy winner to reverse this.

To avoid such outcomes, we must combine the Kaldor criterion with a criterion developed by John Hicks. The Hicks criterion recognises that, if a policy makes people worse-off, they will be willing to pay to avoid its enactment, up to their subjective valuation of their impending loss. It states that a policy will be advantageous to social utility if the amount that policy losers are willing to pay to prevent the change is smaller than the gains of the policy winners. When the losses of the policy losers stand to be larger than the gains of the policy winners, it is possible to redistribute resources in a way that leaves some people better-off and nobody worse-off than they would be under the policy (Hicks 1940: 111).

These criteria are combined into the Kaldor-Hicks account, which identifies increases in social utility where a policy creates a net gain for society: the gains of policy winners outweigh the losses of policy losers and the policy losers' willingness-to-pay to prevent the reallocation is insufficient to compensate the policy winners for their foregone gains. When these criteria are met, policy winners can compensate policy losers so that nobody is made worse-off by the change (Lusk and Norwood 2011: 470; Dasgupta and Pearce 1972: 57; de Scitovszky 1941: 86-87).

Four points need to be clarified. First, the Kaldor-Hicks account has been charged with unfairly favouring the status quo.¹⁸ The account compares a proposed distribution against the current state of the world, and therefore treats the current distribution as the gold standard without questioning whether it is desirable in the first place (De Scitovszky 1941: 88). Part of why the Kaldor-Hicks account favours the status quo stems from problems similar to those we encountered when trying to develop a social utility function. To say that a given distribution is desirable requires us to make value judgements about how society should be organised, which cannot be supported by basic preference utilitarian assumptions and are considered

¹⁷ Sometimes called Tibor Scitovsky in the literature.

¹⁸ This criticism can also be levied at the Pareto account.

inappropriate by many economists (Robbins 1932: 123). By comparing to the status quo, we are not describing distributions as desirable or undesirable in themselves, but rather as *more* or *less* desirable than the current distribution. The Kaldor-Hicks account allows such claims to be made independent of contentious value judgements about which kinds of distributions are desirable. It is thus only by comparing potential states of the world to our current circumstances that we are able to meaningfully speak of changes in social utility.

Second, the use of compensatory payments in the Kaldor-Hicks account appears to couch the discussion in monetary terms, and so the role of utility may seem unclear (Snyder and Nicholson 2012: 157). Kaldor-Hicks efficiency presents utility in monetary terms, through the concepts of *willingness to pay* (WTP) and its counterpart *willingness to accept* (WTA). WTP relates to the sum of money a person is willing and able to spend in order to enjoy another unit of a good or to avoid a unit of something undesirable (Lusk and Norwood 2011: 480; Varian 1992: 161; Norwood and Lusk 2011: 206; Dupuit 1969: 257-58). It is equivalent to the sum of money that someone places on the utility either received from enjoying a good or saved by avoiding a bad (Bannock et al. 1984: 89; Dupuit 1969: 256). For example, if I value my enjoyment of a sandwich at £5, I will not be willing to pay more than £5 for it.¹⁹ WTA, on the other hand, refers to the sum of money that a person is willing to accept as compensation to forego consuming an extra unit of a good, or to consume a unit of something undesirable (Snyder and Nicholson 2012: 158; Mas-Colell et al. 1995: 82).²⁰

Kaldor-Hicks efficiency can thus be defined in terms of WTP and WTA. The Kaldor criterion states that a policy increases efficiency if the policy winners' WTP is greater than the policy losers' WTA. The Hicks criterion states that a policy increases efficiency if the policy losers' WTP to *avoid* the change is less than the policy winners' WTA to maintain the status quo.

Third, it may appear that the interpretation of utility used by the Kaldor-Hicks account has more in common with Bentham's classical utilitarianism than with preference utilitarianism. By determining the effects of a policy on social utility by summing the gains and losses of policy winners and losers – measured in terms of WTP and WTA – the Kaldor-Hicks account could be charged with treating utility changes as commensurable, which implies that we can talk meaningfully about different magnitudes of preference. It must be stressed, however, that the

¹⁹ If I am only asked to pay £3 for the sandwich, then I enjoy a *consumer surplus* of £2, as my subjective monetary valuation of my utility gain is £2 greater than the price charged for the good (Pindyck and Rubinfeld 2013: 132).

²⁰ Due to features of human psychology such as loss aversion, endowment effect and status quo bias, it should not be assumed that WTP and WTA are symmetrical, i.e. what a person is willing to pay for a good is often less than they are willing to accept to give it up (Kahneman et al. 1991).

Kaldor-Hicks account does in fact still depend upon a preference satisfaction interpretation of utility. Both WTP and WTA are couched in terms of preference satisfaction, as they refer to the sum of money that makes someone indifferent between a policy and the status quo. They are not measures of preference strength, and are distinctly unreliable tools for this purpose. This is because the concept of willingness to pay incorporates both willingness *and ability* to pay. A billionaire may have a higher WTP to ransom a close family member who has been kidnapped than a pauper, but this does not mean that the billionaire has a stronger preference to save her loved one – she may simply possess a greater ability to pay. The same applies to WTA. The billionaire may turn down £1 million from a co-worker to switch parking spaces to a less desirable one slightly further from the office, while most people would gladly accept the offer. Accepting the offer does not mean that we must have a weaker preference for a desirable parking space, however: we may simply possess a higher *marginal utility of wealth*, meaning that an increase in wealth provides greater utility to people with less wealth than to people with more wealth (Norwood and Lusk 2011: 206; Dasgupta and Pearce 1972: 44). Essentially, because of his lower level of wealth, average Joe may have a stronger preference for £1 million than for a desirable parking space; he thus assigns a higher level of utility to increases in wealth than a billionaire.²¹

Despite being an unreliable tool for comparing the strengths of preferences between people, using WTP and WTA to approximate utility in monetary terms can be useful, as doing so reveals how much compensation policy losers must be paid to leave them as well-off as if they had consumed a certain bundle of goods (Varian 1992: 109). It needs to be recognised, however, that when using the Kaldor-Hicks account for this purpose, these subjective monetary valuations cannot be converted back into utility terms (Mishan 1988: 27), meaning that WTP and WTA cannot be used to make interpersonal utility comparisons.²²

Fourth, claims about the effects of policies upon social utility can only be reliably made if compensation is actually paid. In contrast, Kaldor (1939: 551) writes: '[o]nly if the increase in

²¹ Variations can also exist between people's preferences for wealth independent of their current levels of wealth. Although a miser such as *A Christmas Carol's* Ebenezer Scrooge and a spendthrift reality television star may possess similar levels of wealth, their different attitudes towards money may support our belief that, when both are willing to pay up to £10,000 for poltergeist-removal services, Scrooge's preference to be rid of apparitions is the stronger. The concept of WTP does not enable us to make this claim, however: all it recognises is that, in financial terms, both Scrooge and the reality star value the ghostbusting services equally.

²² On the level of policy assessment, where a large number of interests are being considered, it may be possible to compensate for different marginal utilities of wealth. Modelling techniques can assign greater weight to the preferences of people on lower incomes, results could be drawn from the median, rather than the total, WTP, or cost-benefit analysis could simply be treated as just one input in the decision-making process, to be considered alongside issues such as distributional effects (Rigby 2019).

total income is sufficient to compensate for such losses, and still leaves something over to the rest of the community, can it be said to be “justified” without resort to interpersonal comparisons’; Hicks (1940: 111) writes that ‘the real income of society is higher in Situation II than in Situation I if it is impossible to make everyone as well off as he is Situation II by any redistribution of the actual quantities acquired in Situation I’. Neither say that these redistributions need to be made, but instead imply that an improvement in social utility can be identified as long as it is *possible* to make them (Dasgupta and Pearce 1972: 58).

It is not obvious, however, that the overall effect of a policy upon social utility can be identified when compensatory payments are not actually made, but exist only as a theoretical possibility. When compensation is not paid, policy changes will likely have a positive impact on the utility of some people and a negative impact upon the utility of others. It was shown before that, when some individuals are made better-off and others are made worse-off, the overall effect upon social utility is uncertain. This uncertainty can be attributed to differences between individuals in their marginal utilities of wealth. A wealthy factory owner could earn an extra £100,000 per year by bringing in a new machine. This machine makes the positions of two workers, each on a salary of £40,000, redundant. The Kaldor-Hicks account states that buying the machine is welfare-increasing: the gains of the factory owner are sufficiently large to compensate the now-redundant workers. If compensation is not paid, however, it is possible that the benefit enjoyed by the already-wealthy factory owner may be less than the misery inflicted upon the newly unemployed and financially insecure labourers (Edwards 1954: 390). Even though the workers’ financial losses do not outweigh the factory owner’s financial gains, the difference in wealth between the factory owner and the labourers means that the utility loss might exceed the utility gains: subjective monetary valuations cannot be converted back into utility. Consequently, when compensation is not paid, the overall effect of a policy upon social utility is likely to be uncertain (Trumbull 1990: 203). If the Kaldor-Hicks account is to reliably identify the effects of policies upon social utility, compensation must be paid rather than exist as a mere theoretical possibility.²³

In short, the Kaldor-Hicks account can enable the preference utilitarian to determine how a policy affects social utility, even in situations where some people experience utility increases and others experience utility decreases. Moreover, the theory avoids interpersonal utility comparisons; all we need to know is how much each policy winner is willing to either pay for a policy to be enacted or accept to maintain the status quo, and how much each policy loser is

²³ This was also Pareto’s view (Pareto 2008: 398; Kemp and Penzani-Christou 1999: 443).

willing to either accept for a policy to be enacted or pay to maintain the status quo. We must, however, recognise that claims about the effects of a policy upon social utility can only be reliably made if policy losers are actually compensated for their losses.

With the two related concepts of Pareto and Kaldor-Hicks efficiency outlined, we are now well-placed to apply these ideas to issues of farm animal welfare and consumer demand, and to consider when improvements to farm animal welfare can be justified with reference to theories of efficiency.

Chapter 2

The Actors: Consumers, Producers and Farm Animals

In this Chapter

- A defence of the claim that, within the UK's legal frameworks, improvements to farm animal welfare tend to come at the expense of productivity, and thus lead to cost increases.
- An assessment of the extent of consumer preferences for farm animal welfare.
- An assessment of consumer willingness-to-pay for improved farm animal welfare.

1. Introduction

In the previous chapter, I outlined two concepts that will underpin much of the analysis in this thesis: utility as preference satisfaction and efficiency. In this chapter, I apply these concepts to issues of farm animal welfare and consumer demand, by asking whether efficiency can justify improvements in farm animal welfare. I begin this chapter by discussing the relationship between product prices and animal welfare. *Ceteris paribus*, lower prices are more desirable to consumers than higher ones,²⁴ and so, if improvements in animal welfare correspond with lower prices, we can expect to find both a universal consumer preference for animal welfare and opportunities for efficiency improvements. I argue, however, that positive associations between price and animal welfare tend to be found only at extremely low and high levels of animal management, rather than at the levels of management currently permitted in the UK. Without a universal reason to support animal welfare improvements, efficiency improvements in this area will require consumers to be willing to pay for improved welfare.

In the second half of this chapter, I assess the extent of demand for farm animal welfare in the UK. I discuss the reliability of the survey evidence, in particular whether factors such as virtue signalling distort responses. Ultimately, I find that survey data provides good reason – but not certainty – to believe that preferences for improved farm animal welfare are widespread in the UK and that there is substantial willingness-to-pay for welfare improvements.

²⁴ There appear to be rare exceptions to this statement. In November 2013, the makers of the popular party game *Cards Against Humanity* recognised the Black Friday shopping event by announcing that, for one day only, they would be raising the price of their product: sales increased slightly (Cards Against Humanity 2013; Dicker 2013).

2. Price Motivations and Farm Animal Welfare

Price is an influential driver of purchases of almost all goods, including foods such as animal products (Schroeder and Tonsor 2011: 801; Ince 2010; Darnton 2016: 5). Holding all else equal, lower prices are nearly always preferred to higher prices as they enable consumers to consume more on the same budget. The consumer preference for lower prices could therefore be said to be universal. The price of animal products is influenced by production processes, and so any discussion of the relationship between consumer price and animal welfare necessarily involves an examination of production.

The producer's role is to supply goods in response to consumer demand. A producer's primary motivation is profit. There are two main ways of increasing profits. When the producer's price²⁵ is greater than production costs, finding ways to increase sales, such as marketing campaigns, can increase profits. Alternatively, producers may reduce costs by improving productivity, which allows them to either sell more units at a lower price while maintaining the same per unit profit margin, or sell at an unchanged price, but with a higher profit margin. Profit may not be the only thing that matters to a producer: irrespective of the economic gains or losses, many producers of animal products also care about the welfare of their animals and will take steps to give them a higher quality of life. This compassion, while laudable, is not of primary concern in this thesis as it is contingent on the producer being the kind of person who feels compassion towards animals. In contrast, *all* producers of animal products must be concerned with productivity as they need to be responsive to the almost universal consumer preference for lower prices.

Where productivity aligns with animal welfare, there is a universal reason for consumers to prefer greater farm animal welfare: higher welfare leads to greater productivity, which can lead to lower prices for consumers. When this happens, improvements in farm animal welfare can leave all consumers better-off and so might amount to efficiency improvements. To assess whether such efficiency improvements are possible, my analysis will be framed in terms of the Pareto and Kaldor-Hicks theories of efficiency. For each theory of efficiency, I will conduct both

²⁵ Consumer price refers to the sum paid for the final good, while the producer price refers to the price received by producers. Consumer prices tend to be higher than producer prices due to the involvement of a range of actors, such as distributors, retailers and processors, in the supply chain. In what follows, price will refer to consumer price unless otherwise stated. It will, however, be assumed that, when price changes are driven by productivity changes, producer prices will move in tandem with consumer prices, or at least will not move against them.

a *speciesist* analysis,²⁶ which focuses solely on human preferences, and a *non-speciesist* analysis that considers both human and animal preferences. I begin with the Pareto account.

a. Pareto Efficiency and Farm Animal Welfare

The Pareto criterion states that efficiency improvements, and thus improvements in social utility, exist where a change in the allocation of resources leaves someone better-off and nobody worse-off. To identify a non-speciesist Pareto improvement, we need to find a reason why all consumers should hold preferences for farm animal welfare. The preference for lower prices is a plausible candidate for a universal preference, but more work needs to be done to determine how far this preference aligns with farm animal welfare. This proposition arouses scepticism; if such an alignment existed, surely low welfare production systems such as battery cages would never have found use or would have been quickly withdrawn? Nevertheless, S.E. Curtis, an animal scientist friendly to the industry, has attempted to prove through his account of animal welfare that an alignment between consumer and animal interests exists (Safran Foer 2010: 64).

Curtis is critical of attempts to define animal welfare in terms of unobservable indicators such as subjective animal mental states, or animal behaviour, which is often taken to be indicative of these unobservable mental states (Curtis 2007: 577). He proposes that an account of farm animal welfare should instead be derived from directly measurable characteristics, namely the productivity of an animal relative to its species' production potential. He writes (2007: 573):

[T]he best single set of measurable ... indicators of [an] animal's state of being will be its rates of productive and reproductive performance relative to its predicted potential to perform. Feed-conversion efficiency, interindividual variation in performance, body condition index, and rates of culling, morbidity, and mortality also will provide valuable information on animal state of being.

According to Curtis, if a farm animal experiences a reduction in performance relative to what its species is capable of, its welfare has declined. Curtis thus argues that the most productive animals – where productivity is measured in terms of weight gain, egg or milk production – will have the highest levels of welfare, while less-productive animals will have lower levels of welfare (Curtis 1989: 172). Importantly, Curtis recognises that the relationship between performance and welfare can be disrupted by changes in rates of morbidity and mortality.

By linking higher productivity with higher welfare, Curtis' theory might be used to propose an alignment between consumer and animal interests. It is in a farm animal's interests for rates of

²⁶ 'Speciesist' is used in a purely descriptive sense to delineate an analysis that does not directly consider animal interests. It is not intended to make claims about the moral propriety of such an analysis.

morbidity and mortality to decline and for its body condition to improve, so it has a corresponding interest in possessing a higher level of some of Curtis' constituents of welfare. Consumers also have an interest in farm animals enjoying a higher level of welfare, as the greater productivity that this entails can contribute to lower prices (Curtis 1989: 172).²⁷ Curtis' theory of welfare therefore implies that non-speciesist Pareto improvements are possible: both consumers and farm animals are potential beneficiaries of some welfare improvements.

This alignment of consumer and animal interests is most consistent at the lowest and highest levels of animal management. When no effort is made to manage farm animals they may suffer from starvation, predation, exposure and disease. Providing nourishment, protection, shelter and medical care will improve both animal health and productivity (McInerney 2004: 18; WSPA 2014), while making efforts to manage animal stress at slaughter can prevent instances of PSE (pale, soft and exudative) flesh in pig meat and DFD (dark, firm and dry) flesh in cattle and sheep meat (FAO 2001: 3-4). At the other end of the scale, where production is extremely intensive, efforts to improve animal health can improve productivity. The most intensively farmed animals are often kept in environments with little stimulation. In pig farms, these environments contribute to vice behaviours such as tail-biting, which cause health issues such as spinal column abscesses (Malone 2019). Addressing tail-biting, through providing more stimulating environments, utilising lower stocking densities or docking tails, produces healthier and more productive animals. Similarly, producers of mink fur found that introducing novelty items such as toys or plastic pipes into previously barren cages reduced the animals' stress levels and contributed to higher-quality pelts (Conniff 2016: 106).

These findings may be of little use to UK producers, however: animal welfare legislation and guidelines already contain stipulations about housing, healthcare and protection, as well as stocking densities and opportunities to engage in natural behaviours (European Council Directive 1999/74/EC, for instance, covers welfare requirements for laying hens), while approved meat premises such as slaughterhouses are monitored by Official Veterinarians from the Food Standards Agency to ensure compliance with welfare regulations detailed in Council Regulation (EC) No 1099/2009. Producers are thus already prohibited from providing levels of welfare so low that they harm productivity.²⁸

²⁷ As welfare improvements may be costly, prices can only decrease when the costs of improving animal welfare are offset by the corresponding increase in animal productivity (Dawkins 1980: 27). The structure of the market and supply chains may also prevent consumers from receiving the full benefit of reduced production costs.

²⁸ Given that I am defining relative statements about farm animal welfare in terms of consumer beliefs about what constitutes higher and lower welfare, this statement is true insofar as consumers believe that these

Although there likely exists an alignment between animal productivity and animal welfare at the lowest and highest levels of animal management, it is not apparent that this alignment persists within the production space created by the UK's legal constraints. Bernard Rollin (2008: 16) claims that the productivity-welfare alignment applies mainly to traditional agricultural societies where animals could only grow properly, or provide milk or eggs, when their physical and psychological needs were met. In such societies, animal health and productivity were extensionally equivalent. In modern industrial agricultural systems, however, Rollin argues that technological advances have driven a wedge between health and productivity to the extent that productivity is no longer a good indicator of welfare.²⁹ Rollin presents two cases where productivity improvements have come at the expense of animal health, and, in the eyes of many consumers, animal welfare.

First, producers have made 'progressive efforts to exploit biological potential through the application of higher levels of inputs' (McInerney 2004: 18), frequently at the expense of animal health: a prime example of this comes from the use of high stocking densities (Rollin 2008: 11-12). When chickens are kept in cages at higher stocking densities, vice behaviours such as cannibalism become more prevalent. Additionally, close confinement provides ideal conditions for pathogens to rapidly mutate and spread, to the detriment of productivity. Despite this, high-density systems are more productive than lower-density (higher-welfare) alternatives³⁰ – California's 2015 ban of caged eggs caused a 9% price increase (Mullally and Lusk 2018: 650), while, in the UK, eggs produced in free-range systems cost about 62% more than their enriched cage counterparts (see Appendix I).

A second way that producers have historically made productivity gains by exploiting animals' biological potential is through confinement, which limits an animal's ability to move. The use of smaller stalls and cages represents a more intensive use of housing inputs by enabling more animals to be contained in the same space, and, by limiting an animals' energy expenditure on movement, maximises its ability to turn feed into body mass. Confinement was historically used to produce veal: calves would be housed in a manner that left them unable to turn around and made it difficult to stand up and lie down.³¹ Giving calves space to roam and exercise would

extremely low and high levels of animal management, which are detrimental to productivity, are also detrimental to an animal's welfare.

²⁹ Again, this assumes that consumers believe animal health to be relevant to animal welfare.

³⁰ This does not mean that animals in high-density systems are more productive than their lower-density counterparts: cage systems are more productive than free-range ones because their higher population densities more than compensate for reductions in individual animal productivity (Proudfoot et al. 1979).

³¹ The use of veal crates was prohibited in the UK in 1990 and is currently covered by The Welfare of Farmed Animals (England) Regulations 2007. The EU banned the practice in 2006.

cause them to develop firmer muscles, burn energy and consume grass, which would cause their flesh to lose its much-vaunted paleness. Consequently, intensive housing maintains the softness and paleness of the flesh and increases the rate of weight gain. This came at a cost to health though, as, beyond having their movement severely curtailed, mortality rates of 10-15% were commonly observed across the 15 weeks of a calf's confinement, although admittedly we cannot know how much of this is a result of confinement and how much is a result of factors such as inadequate diet and deprivation of social contact (Singer 2015a: 129-35).

Last, producers can exploit animals' biological potential for productivity gains by manipulating animal diets to maximise weight gain. This is common in beef production, where cattle are frequently raised on diets largely based upon cheap and energy-rich corn, which leads to faster growth than a more natural grass diet (Norwood and Lusk 2011: 152-53; Singer 2015a: 139-40). Similar to force feeding, the use of higher levels of feed inputs exploits an animal's ability to turn energy into body mass. A corn-rich diet enables cattle to grow fat and promotes marbling of the flesh; when this diet is combined with protein and fat supplements and nontherapeutic medications, a steer can grow from 80 pounds to 1,100 pounds in just 14 months. Once more, however, it seems that this productivity comes at a cost to health, as ruminant digestive systems have evolved to digest grass, not grain. Consequently, a grain-rich diet with little grass can cause health issues such as bloat, acidosis, liver damage and abscesses (Pollan 2011: 71-79; Singer 2015a: 140).

In addition to a more intensive use of inputs, the connection between productivity and animal health has also been severed by the use of selective breeding. Selective breeding has typically been used to produce farm animals that have higher rates of weight gain or produce greater quantities of eggs or milk.³² An early instance can be found in America in 1946, where the Great Atlantic and Pacific Tea Company poultry retailer ran a nation-wide 'Chicken of Tomorrow' competition, a five-year saga that encouraged participants to breed a new chicken that provided more breast, thigh and drumstick meat with less feed (Mason and Finelli 2006: 105). The \$5,000 prize, awarded by US Vice-President Alben Barkley, was won by Charles Vantress, who combined a New Hampshire female with a California Cornish male to produce a broad-breasted chicken that was approximately twice as heavy as a typical barnyard chicken of the time (Lawler 2014). The Chicken of Tomorrow competition kickstarted the use of selective breeding in intensive broiler production, and the subsequent increase in productivity has been startling:

³² More recently, selective breeding has been used for environmentally-friendly ends: Dr Eileen Wall from Scotland Rural College has used the technique to produce cows that produce lower methane emissions (BBC Newsday 2019).

between 1935 and 1995 the average weight of broilers increased by 65% and the time required to reach market weight fell by 60% – all using 57% less feed (Safran Foer 2010: 106). At the end of the Second World War, broilers took more than 13 weeks to grow to 2kg; today, selective breeding means that birds can reach 2kg in under six weeks, and at a cost to consumers of less than £5 (Appleby 1999: 96). Productivity similarly increased in the egg industry, where, today, selectively-bred layers produce twice as many eggs per year as chickens in the 1940s (Mason and Finelli 2006: 105). Modern dairy cows also produce three to four times more milk than their counterparts did 70 years ago (Rollin 2008: 12). Technological advances have also allowed producers to use genetic testing to introduce desirable traits or remove undesirable ones. The pork industry used this technique to remove carriers of a gene connected to pig stress, with the intention of reducing stress-related deaths, particularly from transportation (Safran Foer 2010: 154).

Although selective breeding has significantly increased yields, it frequently has a detrimental effect upon animal health, particularly in species that are modified to grow faster or develop quicker. Breeding animals for increased growth rates or yields without regard for health can lead to weakened immune systems, damaged cardiovascular health, reproductive failure, lameness, metabolic disorders and joint and bone problems (Fearing and Matheny 2007: 162; Almeida Pez et al. 2008: 103; Marcus 2005: 12-13; Rowlands 2002: 103; Voogd 2009; Appleby 1999: 97; Fraser 2012: 194; Norwood and Lusk 2011: 148). This is evident in battery cage systems, banned in the European Union since 2012: the demands placed on hens' bodies by the intensity of egg production, combined with the constricted movement afforded by their environment, results in 30% of live battery hens arriving at the slaughterhouse with at least one freshly-broken bone (Knowles 1994: 60).³³

Broiler chickens also endure serious health problems. The aim of selective breeding in this sector has been to develop birds with faster-growing muscle and fat tissue. Consequently, these parts of the bird grow faster than its bones, which often leads to deformity and disease.³⁴ Three-quarters of broilers suffer from some form of walking impairment, and one-quarter of them suffer a *significant* walking impairment (Safran Foer 2010: 126). Similarly, dairy cows raised in

³³ Knowles acknowledges that the quality of the removal team tasked with getting laying hens from their cages to the slaughterhouse plays a role in this statistic, but also notes that bone strength is significantly reduced in battery hens compared to hens from other production systems (Knowles 1994: 61). The wing strength of battery hens, who live in conditions which do not allow wing exercise, can be just 54% of birds living in higher-welfare systems (Knowles and Broom 1990: 78).

³⁴ Breeder birds must therefore be semi-starved so they can live long enough to produce eggs for hatching; if fed *ad libitum*, these birds would likely collapse and die under their own weight (D'Silva 2008: 36).

traditional agricultural systems could remain productive for as long as 15 years, while it is common for selectively-bred cows in intensive systems to suffer metabolic burnout and become spent after little more than two lactations (Rollin 2008: 13).³⁵

There remain perhaps two further areas where the productivity-welfare alignment could hold true. Producers could identify costless welfare-enhancing changes to husbandry practices (Harvey and Hubbard 2013: 107). This option may be of limited use, however, as leading accounts of hen, pregnant sow and dairy cattle welfare indicate that most constituents of welfare are not linked to costless husbandry practices (De Mol et al. 2006: 160; Bracke et al. 2002: 1822; Ursinus et al. 2009: 547). Instead, the majority of welfare constituents are linked to an animal's living environment. Seeking welfare improvements through costless changes to living environments is also unlikely to be a fruitful approach: welfare-increasing changes to an animal's living environment tend to be costly, and, at more intensive levels of production, are likely to have a negative impact upon total unit productivity.

In short, attempts to identify an alignment between animal welfare and consumer preferences for lower prices – a non-speciesist Pareto improvement – appear destined to fail. Beyond very low and very high levels of animal management, there is a conflict between animal health and the consumer preference for higher productivity and lower prices. Efforts to improve animal health will likely lead to higher prices, while efforts to reduce prices will come at the expense of animal health.³⁶ Without a clear alignment between consumer and farm animal interests, a non-speciesist Pareto analysis cannot identify efficiency improvements.

A speciesist Pareto analysis is unlikely to yield positive results either, and in fact may be more likely to advocate harms to animal welfare. The existence of concern for farm animal welfare is a contingent fact about society, while the consumer preference to pay a lower price for the same good is nearly universal. Successful producers must be responsive to consumer preferences in order to remain economically viable in highly competitive markets (McMullen 2016a: 21-22). A speciesist Pareto analysis conducted in a society with little concern for farm animals could advocate lower animal welfare as a way of satisfying preferences for lower prices. The universal

³⁵ Although the scale of selective breeding, and its associated detriments to health, has expanded in recent decades as animal products have become perceived as staples rather than luxury items in western diets, the practice is not a new phenomenon. In the 16th century, dwarf fowl was perceived as a luxury food, and an analysis of chicken bones found in a feast deposit in Chester concluded that it was likely that chickens were being selectively bred to promote dwarfism in a manner that was contributing to skeletal disorders (Gordon et al. 2015).

³⁶ This conclusion is in keeping with economists including Blandford and Fulponi 1999: 416 and McInerney 2004: 32.

preference for lower prices could therefore create a downward pressure on producers to reduce welfare standards and engage in a race to the bottom, with the bottom being marked as either the point where consumer preferences for welfare are satisfied, where further decreases in welfare standards will cause a decrease in overall productivity, or where welfare reaches legal minimum standards.

There are, however, two ways that producers might seek to avoid this race to the bottom. First, they may focus on producing high-quality goods. Proudfoot et al. (1979) found that the quality of broiler chickens declines when stocking densities increase. The percentage of high-quality Grade A male broiler carcasses declined from 58.4% at a low stocking density of 0.927m² per bird to 27.6% at a high stocking density of 0.372m² per bird, while the percentage of Grade A female broiler carcasses declined from 80.5% at the low stocking density to 58.6% at the high stocking density. Producers could therefore try to avoid lowering welfare standards by focusing on producing higher-quality carcasses, which sell for a higher price than lower-quality ones.

This strategy is destined to end in failure, however, should the figures produced by Proudfoot et al. hold true across the industry. Assume that a broiler producer has a 1000m² barn, and must select a stocking density. She could choose a low density of 0.927m² per chicken, which would fit roughly 1079 birds in the barn, or a high density of 0.327m² per chicken, which would fit 3058 birds in the barn. If the chickens in the low-density scenario are male, she can expect 58.4% (about 630) to produce Grade A carcasses. If they are female, 80.5% (869) of carcasses should be Grade A. In the high-density barn, however, the producer can expect 27.6% of male broilers (844) and 58.6% of female broilers (1792) to be Grade A carcasses. When even those carcasses that are too damaged to be Grade A are still profitable, it seems that producers looking to maximise their output of Grade A carcasses subject to space constraints ought to use higher stocking densities. While the rate of production of Grade A carcasses declines, this is more than offset by the increase in population size, leading to an overall increase in the production of Grade A carcasses, as well as an increase in the still-profitable lower-quality carcasses.

Alternatively, producers could overcome the incentive to reduce animal welfare through product differentiation. Rather than reducing animal welfare to satisfy consumer preferences for lower prices, producers could instead charge a premium for not engaging in low-welfare practices. If at least some people care about the welfare of the animals they consume, higher-welfare producers will avoid competing with low-welfare ones in terms of price or quality, as

they will instead compete in a separate market for higher-welfare animal products.³⁷ This can only be an effective strategy, however, when consumers hold preferences for farm animal welfare. A reliance upon contingent preferences for welfare will therefore likely preclude the possibility of animal welfare improvements amounting to a Pareto improvement, which would require almost everyone in a society to prefer a higher level of farm animal welfare.

In contrast, if an adequate mode of reparation can be established to compensate those who are left worse-off by improvements to farm animal welfare, Kaldor-Hicks efficiency improvements may be achievable.

b. Kaldor-Hicks Efficiency and Farm Animal Welfare

The Kaldor-Hicks criterion defines an improvement in efficiency, and thus in social utility, as a change where the gains of the policy winners are greater than the losses of the policy losers. In technical terms, policy winners' willingness to pay (WTP) for the change is greater than policy losers' willingness to accept (WTA), and policy losers' WTP to maintain the status quo is less than policy winners' WTA. Although Kaldor and Hicks seemed to think that the mere possibility of compensation was sufficient for improvements in social utility, in the previous chapter (p.45) I argued that, due to differences in marginal utility of wealth, we cannot be sure that a Kaldor-Hicks 'improvement' actually increases social utility until compensation is paid. As with the Pareto criterion, we can conduct both a non-speciesist and speciesist Kaldor-Hicks analysis. I begin with a non-speciesist analysis.

One precondition of a non-speciesist analysis is that the concepts of WTP and WTA must be applicable to farm animals. The concepts are essential to the compensatory payment mechanism that lies at the heart of Kaldor-Hicks efficiency and so, if they are not applicable to animals, a non-speciesist Kaldor-Hicks analysis will not be possible. Both WTP and WTA require us to be at least theoretically capable of determining the subjective monetary value of changes in utility. While farm animals cannot tell us the monetary value they attach to utility changes, it may be possible to determine these values by observing their behaviour: in fact, experimental evidence sheds light on the trade-offs animals make between goods. L.R. Matthews and Jan Ladewig (1994) compared the value that pigs place on access to feed and access to social contact. They trained pigs to push levers in order to access either feed or social contact, and then counted how many times a pig was willing to press the lever, i.e. how much

³⁷ Indeed, if at least some people hold preferences for farm animal welfare, cost-cutting measures that are detrimental to animal welfare cannot amount to a speciesist Pareto improvement.

effort it was prepared to exert, for these things. The authors calculated a demand elasticity of 0.02 for feed and 0.49 for social contact, meaning that, when the number of lever presses required for a reward increased by 1%, the pigs' willingness to press the lever decreased by 0.02% for feed and 0.49% for social contact: this also indicates a preference for feed over social contact. The study's results can also be used to roughly estimate the extent to which a pig would be willing to trade food for social contact. Through this, we can calculate a pig's WTP for an increase in social contact in terms of the feed it is prepared to sacrifice. A pig's WTP for an extra unit of social contact can be calculated by dividing the demand elasticities for social contact and feed:

$$\frac{0.49}{0.02} = 24.5 \text{ units of feed}$$

If we assume, in line with F.B. Norwood and J.L. Lusk, that one unit of feed refers to a pound of corn, and the market value of a bushel of corn is £4 (£0.0714/lb), then 24.5 units of feed amounts to a WTP of £1.75 for a unit increase in social contact (Norwood and Lusk 2011: 218).

This is clearly a rudimentary method, and the above results should not be taken for granted, especially considering that Matthews and Ladewig used a sample size of just eight pigs. These calculations also do not consider that pigs possessing a higher level of a good might attach a lower value to an additional unit relative to pigs with less of it. Additionally, preferences for goods might vary between individuals: some pigs may value social contact more highly than others, irrespective of provision. Last, the issue of different marginal utilities of wealth seems especially salient here: a WTP of £100 is likely to indicate a much stronger preference when it comes from a farm animal, which has very few (if any) resources of its own, compared to a human.

Despite these limitations, the Matthews and Ladewig experiment shows that, like humans, farm animals are willing and able to make trade-offs between goods and that we can capture the relative value they place on each good. Comparing these relative valuations to a good that has a clearly defined financial value, such as feed, allows us to calculate an animal's WTP and WTA. Consequently, the application of WTP and WTA to non-humans is not obviously incoherent, and a non-speciesist Kaldor-Hicks test could in theory be used to identify efficiency-increasing welfare improvements.

For farm animals to have a WTP for something, however, they need to also be *able* to pay, which requires the possession of property. This raises both theoretical problems, as it requires support for the contentious claim that animals possess property rights, and practical problems, because

farm animals possess no property: it is the animal's owner who possesses any feed, shelter etc. that farm animals consume and make use of.

Consequently, a non-speciesist Kaldor-Hicks analysis is not well-suited to identifying efficiency improvements in this area. Although farm animals may have preferences for constituents of welfare, their lack of property means that they have zero willingness to pay for welfare improvements.³⁸ A speciesist Kaldor-Hicks analysis that focuses on human altruism, however, may yield more success.

I recognised in Chapter 1 (p.25) that people can hold altruistic preferences for farm animal welfare. These people may merely like certain kinds of animals, or they may attach ethical significance to farm animals and their welfare. A group of animal lovers that is sufficiently large and has a sufficiently high level of altruism may find that, collectively, they value their utility increase from an animal welfare improvement enough to pay for it. If so, animal lovers may be able to compensate cost-bearing policy losers³⁹ to the extent that nobody is made worse-off by the welfare improvement. Because compensation is human, rather than animal, in origin, there is no need to directly consider animal interests, and thus the analysis is speciesist in nature.

A speciesist Kaldor-Hicks analysis appears to be the most plausible way of identifying efficiency-increasing animal welfare improvements. In contrast to a non-speciesist analysis, animals are not required to somehow provide compensation for their own welfare improvements; instead, animal welfare is treated as a subset of human welfare (McInerney 1991). This removes the direct consideration of farm animal welfare from the analysis; I will proceed holding the assumption that farm animal welfare is significant only insofar as consumers are willing to pay for it.

Such a stance is rightly susceptible to criticism, and there exists an entire literature that defends the idea that animal interests and lives are important in and of themselves, rather than insofar as they satisfy human preferences. My analysis in this chapter, however, finds that economic measures of efficiency are ill-suited to recognising this fact in a way that supports animal welfare improvements: consumer preferences for lower prices clash with animal preferences for greater welfare, which renders the Pareto criterion impotent, while farm animals' lack of WTP means that the direct consideration of their interests contributes nothing to a Kaldor-Hicks analysis.

³⁸ WTA will not help here either, as compensation restores policy losers to their pre-change level of welfare, rather than bringing about welfare improvements.

³⁹ Those who bear the costs of welfare improvements are primarily the animals' owners, i.e. producers of animal products. In essence, consumers must have a sufficient WTP for a welfare improvement to compensate the animals' owner for implementing the improvement.

Some would argue that the most appropriate response is to change the entire economic paradigm to recognise the moral worth of animals and to value farm animals and their welfare in themselves (see Francione and Charlton 2015 for example), but this is unlikely to happen any time soon. Working within the existing paradigm to justify welfare improvements in the short-to-medium term is therefore the path that I take in this thesis, and considering animal welfare indirectly, through self-interested and altruistic human preferences, may be the most effective way of doing this. Improvements to farm animal welfare will count as efficiency improvements only where consumers are willing to pay for them.

Given the contingent nature of consumer preferences for farm animal welfare, I must first explore the extent of these preferences. This requires an empirical examination of preferences for welfare through the consideration of survey evidence.

3. Consumer Preferences for Farm Animal Welfare

Perhaps the most convenient way of determining whether people hold preferences for farm animal welfare is to simply ask them. Numerous surveys pose this question, and all indicate that there exists significant support for farm animal welfare. In the 2016 Eurobarometer survey,⁴⁹ for instance, 98% of British respondents answered that it was either very important or somewhat important to protect the welfare of farm animals, and 86% stated that the welfare of British farm animals ought to be better-protected (Eurobarometer 2016: 10, 13). In the Food Standards Agency's biannual public attitudes tracker (2019: 11), 43% of respondents reported that farm animal welfare was an issue of concern with respect to food, while, in another survey, 24% of UK meat-eaters identified concerns about animal welfare as a reason why they would be interested in reducing the amount of red meat and poultry in their diets (Mintel 2017e: 12).

More work needs to be done, however, to confirm that these results are a reliable indicator of attitudes towards animal welfare. At least three confounding factors could be undermining the validity of these surveys, and we must be able to confidently dismiss them before accepting their results. First, the surveys' questions could have been worded in a way that led respondents to overestimate their concern for animal welfare. Second, respondents may have been virtue signalling; talk is cheap in surveys, as respondents do not have to put their money where their

⁴⁹ This survey, funded by the European Union, used the responses from 27,672 citizens from all member states, including 3,461 from the UK, to determine European attitudes towards farm animal welfare.

mouths are. Consequently, instead of honestly reporting their preferences, respondents may have been trying to build or support a reputation as a virtuous person.⁴¹

The international focus of the 2016 Eurobarometer survey gives us good reason to dismiss the first of these concerns. The survey was carried out across the European Union and the country-by-country results reveal clear national divisions. In the UK, 78% of respondents stated a belief that the protection of farm animal welfare is very important, in contrast to 33% in Hungary and 34% in Poland (Eurobarometer 2016: 10). Assuming that the question was faithfully translated into other languages, the low reported concern for farm animal welfare in these countries gives us reason to believe that the question's wording did not prompt respondents to provide pro-animal welfare answers.

The international differences in reported concern for farm animal welfare cannot, however, be used to dismiss the virtue signalling objection, as there may exist cultural or social reasons that motivate Britons but not Hungarians or Poles to engage in virtue signalling in this area.

A third reason to be cautious of accepting the Eurobarometer results is that respondents might have failed to consider that higher animal welfare might come at a personal cost in the form of higher prices. This concern is perhaps best addressed by examining willingness to pay studies for animal welfare improvements.

4. Willingness to Pay for Farm Animal Welfare

Although there are few recent studies of UK consumers' willingness to pay for farm animal welfare, the 2016 Eurobarometer survey indicates that many people are prepared to pay more for higher welfare, or at least say they are. The survey asked respondents whether they would be willing to pay more for higher-welfare products, to which 72% of UK respondents answered in the affirmative (Eurobarometer 2016: 50) – 47% were prepared to pay up to 5% more for higher-welfare goods, 16% between 5% and 10%, 6% between 11% and 20%, and 3% over 20%.⁴²

At face value, these results lend credence to respondents' earlier claim that farm animal welfare is important to them, and implies that many people factor the possibility of higher prices into their support for greater animal welfare. We must be careful, however, not to take these results at face value: they are stated preferences, and so once again may be influenced by virtue

⁴¹ Concerns such as these are often referred to as a *hypothetical bias*.

⁴² Additional studies have identified a substantial WTP in British consumers for a range of perceived welfare improvements, including: legislation to prohibit battery cages (Bennett and Blaney 2003); legislation to ban the export and import of live animals; and broiler legislation (Moran and McVittie 2008).

signalling. As there is no cost to stating a pro-animal welfare attitude in a survey, respondents may inflate their WTP simply to portray themselves as virtuous. As before, though, the international differences in WTP give us reason to reject this possibility: in Portugal and Bulgaria, for instance, only 22% and 28% of respondents stated a positive WTP for higher-welfare products. Although there may be UK-specific cultural or social reasons that encourage respondents to engage in virtue signalling that do not apply to Portuguese or Bulgarian respondents, these international discrepancies lend credence to the claim that many people in the UK genuinely value farm animal welfare and are willing to pay for it.

To further support the claim that stated WTP for farm animal welfare is genuine, one meta-analysis of willingness to pay studies found that, contrary to expectations, stated preference studies tended to provide a lower valuation of animal welfare than revealed preference studies, where willingness to pay is derived from consumer behaviour (Clark et al. 2017: 119).⁴³ This suggests that fears of virtue signalling and cheap talk in WTP studies may be exaggerated, although perhaps cannot be dismissed entirely. As a result, although we cannot be certain, the available research provides evidence for the claim that UK consumers hold preferences for higher-welfare goods and are willing to pay more for improved animal welfare.

5. Summary

In this chapter, I set out to assess how far the concept of efficiency could justify improvements in farm animal welfare. I began by considering the almost universal consumer preference for lower prices and considered the extent to which this preference could contribute to a Pareto improvement by aligning with animal welfare improvements. On the production side, incentives to adopt higher-welfare practices are primarily found either where husbandry is undeveloped and animals have little by way of shelter, adequate feed and protection from predators, or where husbandry is extremely intensive and animals are provided with little space or freedom to express natural behaviours. Given that the UK has legislated against both of these outcomes, there is little scope for further productivity-increasing welfare improvements. In fact, a speciesist Pareto analysis is more likely to justify harms to animal welfare. Instead, some form of compensation is likely necessary if welfare improvements are to amount to efficiency improvements.

⁴³ This may be a result of extremely effective ‘cheap talk scripts’, short statements presented at the start of stated preference surveys intended to reduce the gap between stated and revealed preference studies.

A Kaldor-Hicks analysis reveals that animal preferences add little to our assessment, as animals are incapable of contributing to welfare improvements and so possess zero WTP. Instead, the most plausible overlap between welfare and efficiency improvements appears to be found in a speciesist Kaldor-Hicks analysis. If consumers are prepared to compensate policy losers for improving farm animal welfare, then producers have an incentive to shift to higher-welfare forms of production.

In short, whether an animal welfare improvement can be a Kaldor-Hicks improvement is an empirical matter, dependent on people's preferences for farm animal welfare and how much they are willing to pay for these preferences. Survey evidence indicates that UK consumers hold preferences for farm animal welfare and are willing to pay for welfare improvements. The evidence is limited, however: we cannot be certain of the reliability of these survey findings because they elicit stated preferences, which means that respondents are not required to put their money where their mouths are. Consequently, cheap talk and virtue signalling may undermine survey data. I have sought to alleviate these concerns by showing that attitudes towards farm animal welfare vary across the EU, meaning that, if virtue signalling affects UK results, it must be motivated by factors unique to the UK. The finding that stated preference willingness to pay studies tend to produce lower valuations than revealed preference studies provides further evidence for the existence of genuine consumer demand for improved animal welfare, even if it cannot provide certainty.

At this stage in the thesis, all appears to be well. Although producers have little motive to utilise high-welfare systems for reasons of productivity, they do have an incentive to respond to consumer demand. In the UK, many consumers hold preferences for farm animal welfare, and there appears to be a willingness to pay for higher welfare. For this reason, we may be forgiven for thinking that there is no problem: consumers demand a certain level of animal welfare and it is in producers' interests to respond. The reality is rather different, however: this simple mechanism of demand and supply can be disrupted in a plethora of ways to frustrate the expression of consumers' preferences for farm animal welfare. To understand why this happens, we first need to consider the stage upon which consumers and producers interact: the market. By understanding how markets work, and where they go wrong in the case of animal product markets, we achieve a greater understanding of what the role of policy will be in this domain.

Chapter 3

The Imperfect Stage: Markets and their Failures

In this Chapter

- An introduction to markets, both theoretical and real-world.
- An outline of the various ways that markets can fail to ensure an efficient allocation of resources.

1. Introduction


I concluded the previous chapter with the claim that animal welfare improvements will amount to efficiency improvements – and thus improvements in social utility – only when consumers are willing to pay for them, i.e. if there is sufficient demand. This chapter extends this analysis by introducing the market, the institution where consumers express their preferences. I begin this chapter by examining what perfectly competitive markets look like and why they are viewed as desirable institutions. I then argue that almost all real-world markets fail to live up to this perfectly competitive ideal. In some cases, the factors preventing markets from achieving perfect competition also contribute to an inefficient allocation of resources: these inefficiencies are known as market failures. I finish the chapter by outlining the four forms of market failure that are especially salient to the market provision of farm animal welfare, which will be the focus of the following four chapters.

2. Free Markets

An *economic market* is an institution where buyers and sellers can exchange goods. Properly-functioning markets allow people who wish to buy or sell goods and services to contact each other for the purpose of making exchanges, enable households to make consumption choices, let workers decide how much and for whom to work, and permit firms to choose what to produce and how to produce it (Begg et al. 2000: 9). The market reconciles all of these decisions through the *price mechanism*, where the price of a good moves towards an equilibrium where supply equals demand. The market clears when this equilibrium is achieved, meaning that there is no surplus or deficit of a good (Pindyck and Rubinfeld 2013: 7). Figure 3 provides a demonstration of the price mechanism in action when a good is priced too high at P_0 . As the

good's price is higher than many potential buyers are willing to pay, demand is low (Q_{0D}). In contrast, sellers are willing to supply more of the good (Q_{0S}), as they receive more money for selling it at current prices. The result is that supply outstrips demand and there is a surplus of the good, measured by $(Q_{0S} - Q_{0D})$. To reduce the surplus, or to prevent its further increase, the good's price must decrease. As the price declines, supplying the good becomes less attractive to sellers, causing a reduction in supply. The cheaper good becomes more appealing to buyers, causing an increase in demand. This process continues until the price of the good is at a level (P_1) where demand is equal to supply and the market clears.

Figure 3 – The Role of the Price Mechanism at Non-Equilibrium Prices
(Pindyck and Rubinfeld 2013: 25)



Disrupting the price mechanism can affect supply and demand for a good. Governments can impose taxes or establish a price ceiling or a price floor to regulate markets. The price mechanism may also be disrupted by the actions of a price-setting monopoly. Markets where the price mechanism is not limited in these ways, and thus operates freely, are called *free markets* (Bannock et al. 1984: 180).

3. Perfectly Competitive Markets

The free operation of the price mechanism, although sufficient to create a free market, is just one of the characteristics of a *perfectly competitive market*. Perfectly competitive markets also require the presence of many buyers and sellers, each of whom is responsible for only a very small portion of the total quantity of the good traded on the market. These buyers and sellers must act as *price-takers*, as if their actions have no effect upon prices (Kreps 1990: 263; Begg et al. 2000: 127; Snyder and Nicholson 2012: 371).

Third, goods in a perfectly competitive market are homogenous, meaning that competing firms produce identical, or nearly identical, goods (Pindyck and Rubinfeld 2013: 280; Bannock et al. 1984: 337). This means that the goods or services sold in each market are *perfect substitutes*; each producer thus competes in the market solely on the basis of the product's price, rather than quality or brand or another factor. We could not speak of a perfectly competitive market for apples, for instance, as there are different varieties of apple, but perhaps a market for braeburn apples could be perfectly competitive, as consumers looking to buy braeburns will likely compare them to other braeburns solely on the basis of price.⁴⁴

Fourth, participants in a perfectly competitive market will possess perfect information. Buyers must be aware of and understand the attributes of the goods they are buying – without this knowledge, they may be unable to recognise that each firm in the market is offering the same product, and so may make consumption choices based on non-price factors, such as mistaken beliefs about product quality (Pindyck and Rubinfeld 2013: 127-28). In addition, buyers must hold perfect information about how sellers price their goods (Kreps 1990: 264). This allows them to take their custom to other vendors if they find a lower price, and also provides sellers with an incentive to price their products competitively. Sellers are also assumed to have perfect

⁴⁴ This is perhaps an over-simplification – some consumers care about other attributes, such as country of origin and whether a product is organic.

information about how rival vendors price their goods, allowing them to adjust their prices to remain competitive (Bannock et al. 1984: 337; Kreps 1990: 264).

Fifth, entry to, exit from and transactions in a perfectly competitive market are costless (Kreps 1990: 264; Snyder and Nicholson 2012: 371). This means that new sellers can freely enter the market and sell goods under the same conditions as more established sellers, should these vendors not be pricing their goods competitively (Bannock et al. 1984: 337; Kreps 1990: 264). Similarly, there are no costs to a firm leaving the market if it is no longer able to compete.

Sixth, the price of goods in perfectly competitive markets will capture all costs and benefits; only the buyer and seller should be affected by a good's production and consumption (Nechyba 2017: 520; Kreps 1990: 202). This means that there are no unpriced externalities from the production of a good. Although externalities will be discussed in greater detail in the next chapter, it is worth providing an example here. Suppose that a good's production creates air pollution. This pollution affects people who are not involved in the transaction between consumer and producer and imposes a cost upon them. The cost of this pollution is unlikely to be captured by the good's market price (Pindyck and Rubinfeld 2012: 663). As a consequence, those who are affected by the air pollution are unlikely to be compensated for the ill effects they suffer. The full social cost of the good is therefore not reflected in its market price, contributing to an inefficient level of supply; in this instance, the market will oversupply the good. Externalities are not confined solely to a good's negative effects: goods can also benefit those who are not involved in the transaction, and these benefits are not always reflected in the good's price. When this happens, the market will supply an insufficient quantity of the good (Nechyba 2017: 753); those paying for the good do not receive its full benefit, and are therefore not willing to pay for a level of supply that is socially-optimal (Pindyck and Rubinfeld 2013: 665). Instead, they will consume the good up to the extent that their own personal benefit is equal to the price they are charged, irrespective of the benefit received by others (Begg et al. 2000: 265). Consequently, markets in goods with externalities, both positive and negative, cannot be perfectly competitive.

Seventh, and last, both buyers and sellers behave competitively in a perfectly competitive market. This means that buyers will not pay more for a good when they can purchase the same good elsewhere at a lower price, and sellers seek to receive as high a price as possible; they are utility-maximisers and profit-maximisers respectively (Kreps 1990: 264; Kaneko and Wooders 1994: 73). In short, both buyers and sellers are self-interested. This implies that perfectly competitive markets will tend towards just a single price for a good, with sellers who charge

above the market price losing their customers to rivals and being run out of the market (Begg et al. 2000: 127; Kreps 1990: 264). Those charging below the market price will either make a loss on each unit sold and eventually leave the market, or will draw customers away from other sellers, forcing less-competitive rivals out of business. Their lower price will then be established as the new market price.

These seven characteristics (free operation of the price mechanism, consumers and producers who act as price takers, homogenous goods, perfect information, free entry to and exit from the market, absence of externalities and competitive behaviour) must all be present for a market to be perfectly competitive. When these conditions are met, markets will lead to a Pareto-efficient distribution of resources.⁴⁵ This is largely a result of decision-making being decentralised, which provides two key benefits. First, consumers can pursue their own advantage by exploiting local information. This in turn transmits information to the rest of the economy through price changes, which encourages producers to respond to changes in demand and satisfy consumer preferences. Second, competitive pressures encourage innovation and experimentation, meaning that goods will be produced at the lowest possible costs (McMullen 2016b: 35-36).

4. The Problem with Perfectly Competitive Markets

Although real markets may possess some of the seven characteristics of a perfectly competitive market, it is rare to find a market that possesses them all. Commodities that we encounter on a day-to-day basis are often differentiated using brand names (Hill and Myatt 2007: 63); some markets naturally tend towards monopoly (Posner 1969: 548); market entry can be prohibitively expensive for newcomers (Snyder and Nicholson 2012: 449-51); we often possess incomplete information about the goods that we buy (Nelson 1970: 311-12; Darby and Karni 1973: 68-69); and

⁴⁵ The connection between perfectly competitive markets and Pareto efficiency is made by the first fundamental theorem of welfare economics, which states that ‘any Walrasian equilibrium is Pareto-efficient’ (Reiss 2013: 212). The concept of Walrasian, or general, equilibrium was named after 19th century economist Léon Walras, who noted that changes in one market have knock-on effects in others. A product surplus in one market indicates that a good’s price is too high, and means that there must exist another market where a good is priced too low. These knock-on effects can be a result of one good being a substitute or a complement to another, or alternatively because a good serves as a production input in another market (Pindyck and Rubinfeld 2013: 595). A Walrasian equilibrium exists when all markets within an economy are in equilibrium, i.e. when all markets are perfectly competitive. In a state of Walrasian equilibrium, all consumers maximise their utility subject to their budget constraints, all firms maximise profits subject to production constraints, and supply equals demand in all markets; all markets clear (Reiss 2013: 212). A consequence of this is that there is no way of changing the distribution of resources in such a way that the utility of at least one individual is improved without making another person worse-off – a general equilibrium is therefore Pareto-efficient across all markets. Contrary to many economists’ claims (see Mas-Colell et al. 1995: 524; Pindyck and Rubinfeld 2013: 609; Snyder and Nicholson 2012: 428; Begg et al. 2000: 259; Kreps 1990: 200), the first fundamental theorem is not a formal statement of Adam Smith’s invisible hand theory (Blaug 2007: 188-89).

habit often causes consumers and producers to behave in ways that are not competitive (IGD 2007: 35). The theory of perfect competition is thus rarely a descriptive account of everyday markets, but rather an ideal system of resource allocation that real markets tend only to approximate to a greater or lesser degree (Bannock et al. 1984: 338).

In some cases, a market's failure to achieve the perfectly competitive ideal can result in an inefficient distribution of resources. Such markets suffer from a malaise known as *market failure*, a dramatic-sounding term used to denote any degree of distributive inefficiency, no matter how small (Begg et al. 2000: 264). These market failures could be structural in nature, as in those markets that naturally lead towards monopoly or monopolistic competition, or in those with high costs of entry. Other markets suffer failure due to the nature of the goods being exchanged: a good might bear externalities, or it might be a public good, one which is nonexcludable and non-rival in nature.⁴⁶ Alternatively, market failure can stem from the participants themselves: to give two examples, agents may not possess complete information about a good, a transaction or the market as a whole, or they may not behave in an entirely rational manner. There may be scope for governments to intervene in markets suffering from some forms of market failure, if these interventions improve efficiency by leaving at least some people better-off and nobody worse-off.

5. Summary

The ideal form of market is often held to be the perfectly competitive market, one that requires the possession of a range of characteristics that real-world markets simply cannot achieve. Real-world markets fail to live up to the perfectly competitive ideal in many different ways, some of which contribute to an inefficient distribution of resources. Four forms of market failure – externalities, public goods, information asymmetries and uncompetitive consumer behaviour – appear to be especially salient to the market provision of farm animal welfare.

Over the next four chapters, I address these market failures in greater detail, and discuss how policy can be used to rectify them. In **Chapter 4**, I explore externalities, and consider the extent to which taxes, subsidies, cap and trade and market creation can encourage markets to reach a socially-optimal level of supply. In **Chapter 5**, I discuss public goods, and discuss how legislation and taxes can be used to directly supply these goods. In **Chapter 6**, I consider the impact of informational asymmetries and how policies such as labelling schemes can empower consumers

⁴⁶ Public goods will be addressed in greater detail in Chapter 5.

to make purchase decisions based on preferences for farm animal welfare. In **Chapter 7**, I investigate uncompetitive consumer behaviour with reference to findings from behavioural economics and discuss how policy can be used to change decision-making environments in ways that encourage pro-animal welfare behaviours.

In **Chapter 8**, I pull these findings together to propose a new way of thinking about the objectives of economic farm animal welfare policy, and develop a policy platform that determines policy aims by considering facts about both consumer concern for animal welfare and farm animal welfare in a society.

Chapter 4

Externalities

In this Chapter

- An introduction to externalities.
- An outline of externalities created by animal agriculture.
- An assessment of the effectiveness of externality policies in internalising the externalities created by animal agriculture.
- A discussion of how health and environmental policies could be detrimental to animal welfare.

1. Introduction

Your neighbour buys an expensive and powerful new sound system and pumps out bass-heavy music until the early hours of the morning, depriving you of your much-needed beauty sleep. After several disturbed nights, the sound system has provoked so much stress and rage that you become ill and are forced to take a day off work. The costs of the sound system thus extend beyond what your neighbour paid for it: it has imposed costs upon you as well. The sound system has undoubtedly adversely affected your quality of life and yet you were neither consulted before its purchase, nor compensated for the disruption it has caused you. In addition to its thumping bass, the sound system also produces what is referred to in economics as an externality, the focus of this chapter.

I begin this chapter by developing the technical concept of externalities, before discussing the extent to which they exist in animal agriculture. I then consider the role of policy in this area, and in particular the extent to which four policies – taxes, subsidies, cap and trade, and market creation – can resolve externality effects in animal product markets.

2. Externalities

Externalities are costs or benefits that are imposed upon individuals and firms not directly involved in the production or consumption of a good or service (Nechyba 2017: 744). Externalities can be *technological*, *pecuniary* or *psychological* in nature. Technological externalities relate to physical effects such as pollution. Many manufacturers emit pollution, the

harmful effects of which are rarely captured by their production costs. Pecuniary externalities concern financial effects such as losses arising from increased competition. In a hypothetical city dominated by a hypothetical university, an increase in student numbers leads to a rise in property prices as residential houses are turned into student houses and land is purchased to build student accommodation. This imposes a pecuniary externality upon local residents, who find themselves no longer able to afford property in the city. Another example of a pecuniary externality can be found in the decline of the local high street as a consequence of the rise of online shopping and out-of-town superstores (Halpern 2015: 223). Last, psychological externalities are feelings of pleasure or displeasure that stem from a transaction but do not affect prices (Mann 2005a: 371). Many people get into the festive spirit by purchasing and displaying extravagant Christmas decorations, which provide festive cheer to passers-by. These passers-by, however, make no financial contribution to the Christmas lights, despite the pleasure they draw from them.⁴⁷

Economists do not believe that governments should intervene to rectify all externalities, as poorly informed interventions can be detrimental to a market's efficiency. Additionally, it is widely held that pecuniary externalities do not warrant intervention as they do not affect the efficiency of a market. Even though they may harm some people's welfare, these externalities may in fact be necessary for markets to operate efficiently (Holcombe and Sobel 2000: 157; Prest and Turvey 1965: 688).⁴⁸ In contrast, technological externalities and some psychological externalities are generally viewed as warranting some sort of policy response, when circumstances permit these responses to be effective.

Externalities can also be *positive* or *negative*. Positive externalities exist where someone receives no payment for benefits they have produced for others. In the smartphone industry, firms frequently co-opt the innovations of rivals to produce goods that are similar, but not so similar as to infringe patents, and so do not require the payment of license fees. Research and development thus creates positive externalities for copycat firms (Pindyck and Rubinfeld 2013: 665). Positive externalities also arise from market transactions (Narveson 2003: 208). Keith buys himself a watch which he uses to remind Louise that her important meeting is starting soon, saving her from arriving late. Louise's benefit did not influence Keith's decision to buy the watch, and she pays nothing for it, yet it is undeniably a utility-enhancing side-effect. Both those

⁴⁷ Although it may seem that psychological externalities appear to draw from classical utilitarianism, they can also be understood in terms of preference satisfaction: in the example above, people have preferences to see Christmas decorations and the satisfaction of these preferences creates feelings of pleasure.

⁴⁸ Responses to pecuniary externalities may, however, be justified for social reasons.

directly involved in a market transaction and those who circumstantially benefit from it are made better-off, yet purchase decisions are generally based on a good's private benefits and costs rather than its wider social effects. Consequently, markets tend to undersupply goods with positive externalities.

Negative externalities, on the other hand, occur when someone imposes a cost upon another for which they do not provide compensation (Koutsoyiannis 1980: 541). Suppose a cattle farm releases pollution into a river, which affects a downstream corn farm that uses the water as a production input. The corn farm is required to treat the water in a costly process, which the cattle farm does not pay for. The social costs of the cattle farm's beef are thus greater than its private costs, as the price received by the producer does not include the costs of pollution (Begg et al. 2000: 266). As purchase decisions are typically based on a good's private costs and benefits and not its wider social impacts, markets tend to oversupply goods that produce negative externalities.^{49, 50}

Externalities can also be categorised according to whether they stem from *production* or *consumption*. Production externalities are caused by a firm expanding its output, which imposes costs or benefits upon others for which the firm neither pays nor receives payment (Salvatore 2003: 606). If the aforementioned cattle farm expands production, it will release more pollution into the river, increasing the corn farm's spending on water treatment. This constitutes a negative production externality. Conversely, consumption externalities stem from the consumption pattern of consumers. These consumption habits can impose costs or benefits upon others, for which consumers neither pay nor receive payment (Koutsoyiannis 1980: 544; Perloff 1999: 656). Keith's watch, for example, provides a positive consumption externality when it helps Louise arrive on time for her important meeting.

Last, externalities can exist between firms, between firms and people, and between people. Positive inter-firm externalities exist between orchard owners and beekeepers. If an apple farmer increases the size of his orchard, he will create more nectar for the bees, to the beekeeper's benefit; on the other hand, if the beekeeper expands her production, there will be more bees to pollinate the apple farmer's orchards (Meade 1952: 56). Firms can also impose externalities upon individuals. Manufacturers' pollution may harm people's health, but polluters tend not to bear these costs. Similarly, people can impose externalities upon firms –

⁴⁹ Some externalities can be simultaneously positive and negative. Some people enjoy the smell of tobacco emanating from a nearby smoker, or the sight of wind turbines against a bucolic backdrop, while others find these things offensive (Bagus 2011: 118; Perloff 1999: 656).

⁵⁰ A graphical treatment of both negative and positive externalities can be found in Appendix II.

transport pollution can harm the productivity of farmers (Nicholson and Snyder 2010: 520). Last, people can impose externalities upon other people. Someone smoking in a public space imposes an externality upon non-smokers, although this externality can be either positive or negative depending on the preferences of those affected (Koutsoyiannis 1980: 544).

In short, externalities are costs or benefits that are imposed upon a third party (Bannock et al. 1984: 164). Where externalities exist, the producer price of a good or service does not reflect its full social value, which results in an inefficient level of supply (Koutsoyiannis 1980: 542). Although all externalities are forms of market failure, they are generally deemed to only warrant an intervention if they are a technological externality, or in some cases a psychological externality.⁵¹

3. Externalities in Animal Agriculture

Animal agriculture produces numerous production-side externalities. The rise of large-scale industrial animal farms has created pecuniary externalities by forcing many small-scale farmers, who often used traditional husbandry techniques that were more environmentally- and animal-friendly, out of business. In the United States, the number of pork producers declined 87.8% between 1980 and 2005, while output rose from 69 million pigs per year to over 103 million (Rollin 2008 15; Singer and Mason 2006: 43).⁵² The decline of small-scale agriculture has decimated many communities and left them to contend with high levels of unemployment.

Of additional concern are the industry's technological externalities, which a 2011 UK government report claimed that, if left unchecked, 'will lead to irreversible environmental damage and long-term threats to the viability of the food system' (Foresight 2011: 33). Intensive farming produces huge quantities of animal waste which, if improperly managed, can release ecologically disastrous levels of acids, nitrogen and phosphorus into the local environment and water sources (Fearing and Matheny 2007: 163; Mason and Finelli 2006: 119; Singer and Mason 2006: 31-32, 65). Animal agriculture produces further technological externalities through its

⁵¹ The question of which psychological externalities warrant policy interventions is a difficult one that lies beyond the scope of this thesis. A satisfying account must be able to discriminate between seemingly irrelevant psychological externalities, such as a dislike of socks and sandals, hateful psychological externalities, such as a dislike of people of a certain race, and more appropriate psychological externalities, for instance concern for farm animal welfare.

⁵² UK figures tell a different story, with a decline in the number of producers corresponding with a decline in total herd size (AHDB 2019b, 2019c). As we will see in Chapter 5 (p.114), this result is anomalous, and is primarily a result of legislative interventions that left the domestic industry unable to compete with cheaper imports.

contribution to water shortages, deforestation, loss of biodiversity, soil depletion and land erosion (Durning and Brough 1995: 154-59; Francione and Charlton 2015: 71). The industry also harms the global environment as it is a major producer of greenhouse gases. UN research concluded that livestock supply chains are responsible for 14.5% of the world's greenhouse gas emissions – more than the transport sector's direct emissions – primarily in the form of methane, but also nitrous oxide and carbon dioxide (Gerber et al. 2013: 15; Koneswaran and Nierenberg 2008: 580; Bailey et al. 2014: 4).⁵³ One report found that the emissions of the world's largest three meat producers – Brazilian firm JBS and American firms Cargill and Tyson – were higher than that of France in 2016 (Majot and Kuyek 2017). In the UK, animal agriculture was responsible for about 10% of the country's greenhouse gas emissions in 2017, with the meat and dairy sectors the biggest polluters (DBEIS 2019; Garnett 2009: 492).

Pollution from animal agriculture has been linked to illnesses including nervous-system impairments, cancer and 'blue-baby syndrome', as well as the rise of antibiotic-resistant diseases such as MRSA (Durning and Brough 1995: 155; Safran Foer 2010: 175-76). High stocking densities also provide ideal conditions for diseases to mutate into forms that are both more virulent and transmittable to humans, with a 2005 UN task force concluding that high-intensity farming was a fundamental cause of outbreaks of avian influenza (Singer and Mason 2006: 34).

The consumption of animal products produces additional technological externalities through its contribution to acute foodborne illnesses. Intensively-produced meat products are a source of foodborne illnesses caused by contamination from intestinal pathogens such as campylobacter, E. Coli and salmonella. The Food Standards Agency estimates that, in 2018, campylobacter imposed a cost of £712.6 million upon the UK, salmonella £212 million and E. Coli £3.9 million, with costs primarily borne by individual sufferers and carers (Daniel et al. 2020).⁵⁴ While not all cases of campylobacter, salmonella and E. Coli can be attributed to animal products (dogs and cats can be carriers of campylobacter, person-to-person transmission of salmonella is possible, and exposure to contaminated faeces can cause E. Coli infections, for instance), a significant proportion of the societal costs of these illnesses stem from exposure to contaminated animal products.

⁵³ Goodland and Anhang's (2009) research, which factors in animal agriculture's indirect contributions to greenhouse gas emissions, concludes that the industry accounts for at least 51% of global annual emissions.

⁵⁴ The FSA's model is one of the more comprehensive cost of illness models, as, unlike many alternatives (see Tam and O'Brien 2016 and Santos et al. 2010 for instance), it assigns a value to the pain and suffering endured by individuals, in addition to the direct costs of medical care and indirect costs associated with loss of productivity.

The consumption of some animal products also contributes to chronic diseases. A 2015 study by the World Health Organisation concluded that processed meat is carcinogenic. Red meat was found to be ‘probably carcinogenic’, meaning that positive associations between its consumption and cancer have been observed, but certain confounding factors cannot be ruled out (IARC 2015). Each additional 100g portion of red meat consumed per day increases a person’s risk of colorectal cancer by 17%, while each additional 50g daily portion of processed meat increased the risk by 18%. Increased rates of prostate and pancreatic cancers were also linked to consumption of red meat, while the consumption of processed meat was connected to a higher risk of stomach cancer (Bouvard et al. 2015: 1599). Consumption of red and processed meat has also been linked with higher rates of diabetes, heart disease, strokes and overall mortality (Springmann et al. 2018a: 2; Sinha et al. 2009). These diet-related illnesses impose significant costs upon healthcare systems, governments and societies: Springmann et al. (2018b: 40) estimate that the UK’s annual expenditure on health problems caused by red and processed meat consumption will reach £5 billion by 2020, through government healthcare expenditure, patient spending, opportunity costs of informal care and loss of productivity as a result of illness and death. In addition to these direct costs are the nonfinancial costs of illness relating to pain, suffering and grief, which are harder to quantify but no less significant. Springmann et al.’s model does not consider these nonfinancial costs and so likely underestimates the full social costs of red and processed meat-related illnesses.

Although these externalities are serious issues that surely demand government action, they lie beyond the focus of this thesis, which is concerned with farm animal welfare.⁵⁵ Animal agriculture produces externalities that affect animal welfare in a multitude of ways. A farmer who pays to vaccinate himself against avian influenza simultaneously protects his fowl against the risk of his infecting them, even if this was not his intention when purchasing the vaccine. A tree-planting initiative intended to make the local community more aesthetically pleasing also improves air quality, a benefit that is enjoyed by both humans and farm animals with access to an outdoor space.

These positive externalities must be contrasted with the negative externalities produced by animal agriculture, however. Intensive production harms vast numbers of animals – in 2018, the UK farmed roughly 10,004,000 cattle, 4,969,000 swine, 181,818,000 poultry and 34,832,000 sheep, as well as millions of farmed fish, the numbers of which are not counted (DEFRA 2019a: 6).

⁵⁵ I will, however, return to health and environmental externalities (p.78) to consider how responding to these issues might impact farm animal welfare.

Intensively-farmed animals are frequently kept in densely-populated production systems that offer few environmental stimuli and deny them the opportunity to express important natural behaviours (Pollan 2011: 218; Singer 2015a: 99-102, 121-22). Farm animals commonly respond to these conditions with so-called vice behaviours such as feather pecking and tail biting. Many producers address these vice behaviours through noncurative operations such as beak trimming and tail docking, which are typically performed without anaesthesia.

Although it is clear that farm animals can be both positively and negatively affected by externalities, this thesis is primarily concerned with human responses to farm animal welfare, which are prime examples of psychological externalities (McMullen 2016a: 33). As discussed in Chapter 1 (p.26), many people derive satisfaction from the belief that the animals that provide their food have lived happy lives (FAWC 2011: 6; IGD 2007: 16; Brook Lyndhurst 2010: 40), indicating that farm animal welfare creates positive psychological externalities. Conversely, low welfare can distress people who care about animals (Norwood and Lusk 2011: 294-95, 311; Lusk and Norwood 2011: 474). These two ways of valuing animal welfare may at first sight seem to be two sides of the same coin but in fact demand different responses from consumers. If consumers mainly experience negative psychological externalities from animal agriculture, they are likely to respond by boycotting low-welfare products. Should, on the other hand, animal welfare be experienced primarily as a positive externality, consumers will more likely purchase higher-welfare products.

Importantly, psychological externalities can be experienced by both consumers and non-consumers of animal products. Vegetarians and vegans may be upset by low-welfare conditions, but consumers of higher-welfare goods may also feel the same way. Conversely, a consumer of lower-welfare goods can also derive satisfaction from the existence of higher animal welfare, even if she does nothing to contribute to its provision.

4. Externality Policy Aims

Because externalities indicate that a good's producer price does not reflect the full value of its social costs or benefits, markets in externality-bearing goods will not distribute resources efficiently. Consequently, there is scope for policy to increase social utility by internalising externality effects within a market so that a good's price reflects its full social benefits and costs.

The aim of policy is thus to encourage people to support goods with positive externalities and to discourage the support of goods with negative externalities. In animal product markets, this

may take the form of encouraging purchases of higher-welfare goods and discouraging purchases of lower-welfare ones in order to internalise psychological externalities. Policies such as taxes, subsidies, cap and trade, and market creation could be used to achieve this end.

5. Taxes

Taxes were first proposed as a means of responding to externalities by Arthur C. Pigou, in his 1920 book *The Economics of Welfare*. For this reason, they are often referred to as *Pigouvian* or *Pigovian taxes*. Externality taxes are intended to make prices reflect the full social value of a good; they are thus levied on participants in an externality-producing market.⁵⁶

Because taxes reduce demand for a good or service by increasing its price, they are most effective when used to target goods that create negative externalities. There are thus limits to how they should be used as economic farm animal welfare policy tools. We have seen that low animal welfare can cause people displeasure and so can be viewed as a negative psychological externality. A tax on animal products or low-welfare modes of production may be a suitable policy in these cases. In those cases where people view farm animal welfare as a positive psychological externality, however, a tax will not be an appropriate response as it reduces the supply of animal products from low-welfare systems, but does not guarantee a corresponding increase in supply from higher-welfare systems.

The mere implementation of a tax in order to address negative externalities does not guarantee success, however, as three factors can limit its effectiveness. First, an externality tax must be levied on the externality itself to avoid producing counter-intuitive outcomes. For example, a government may try to tackle pollution by levying an externality tax upon output in a polluting industry. When the tax is levied on output rather than pollution, however, a firm that uses low-pollution technology will pay the same tax per unit of output as more polluting firms. This reduces the incentive to develop low-pollution technology, contrary to policymakers' intentions (Nechyba 2017: 748). It is therefore more appropriate to levy taxes based on firms' contributions to negative externalities, rather than output. Accordingly, a welfare-motivated externality tax ought to be based on the welfare provided to animals, rather than the number of animals farmed or quantity of goods produced.⁵⁷

⁵⁶ A graphical treatment of externality taxes is found in Appendix III.

⁵⁷ This could be linked to a tiered labelling system of the kind discussed in Chapter 6, or simply levied on the number of animals subjected to certain practices. Such a tax could be levied on the producers themselves, or on the final products.

Second, an externality tax must be applied consistently across all relevant industries. If a welfare-motivated externality tax is applied to the beef industry but not the dairy industry, for instance, the market may respond by reallocating cattle to the dairy industry, which will prove counterproductive should dairy cattle receive lower levels of welfare than beef cattle. Consequently, if low-welfare taxes are not levied consistently on all animal industries, they may in fact diminish welfare (Cowen 2006: 42).

Third, the political viability of an externality tax must be considered, in particular the size it must be to have a meaningful effect upon consumer demand. Demand for animal products tends to be inelastic in the short run (Fearing and Matheny 2007: 170; Cowen 2006: 44; Malone and Lusk 2016: 519; Bennett 1997: 286; Tiffin et al. 2011), meaning that a price rise is likely to cause a smaller decrease in demand: if a tax increased the price of eggs by 5%, demand would decrease by less than 5%. The price inelasticity of animal products means that any tax aiming to substantially reduce demand will have to be large: Springmann et al. (2018a; 2018b) predict that a health-motivated externality tax of 79% on processed meat and 14% on red meat would reduce processed meat consumption by 22% and, because of substitution effects, would leave the consumption of red meat unchanged.

A large welfare-motivated externality tax may not be politically viable. A survey of the literature finds no instances of an animal welfare tax being enacted, and just a single instance of an animal welfare tax being called for: the conservative Christian Democratic Union's Barbara Otte-Kinast, the agriculture minister for Lower Saxony, Germany, called for an animal welfare tax on animal products, the proceeds of which would be used to support welfare improvements (Schulz 2020). The lack of support for a welfare-motivated tax implies that this policy is simply not politically viable in the current climate.

A welfare tax, along with other food taxes, also has implications for food justice. While, with certain goods at least, questions of justice may not be quite so salient, this is not the case with meat, which is a key source of dietary protein for many people. Increases in food prices disproportionately disadvantage people on lower incomes (Lang 2019: 815; Noor 2019), who spend a higher percentage of their incomes on subsistence; a welfare tax thus may deprive people on lower incomes of a key source of protein. Springmann et al.'s health-motivated red and processed meat tax, for instance, stands to have the greatest impact upon those on lower incomes, who not only spend a higher proportion of their income on food, but also consume more red and processed meat than those on higher incomes (Maguire and Monsivais 2015: 185).

When introduced in isolation, an animal welfare tax may even contravene the UN's International Covenant on Economic, Social and Cultural Rights, which recognises the universal human right to adequate food. 'Adequate' is defined not only as being sufficient to prevent hunger, but also as being acceptable within a given culture, and thus affords significance to the 'non-nutrient-based values attached to food and food consumption' (United Nations Human Rights Officer of the High Commissioner 1966). Animal products constitute a large part of the UK's culinary heritage – a traditional Christmas dinner, for instance, would surely be incomplete without a roast turkey or chicken, and a mind's-eye representation of a full English breakfast will likely include some, if not all of, eggs, bacon, sausages and black pudding. As such, a welfare-motivated externality tax risks alienating the worst-off in society from certain food styles.

Consequently, policy ought to ensure the fair representation of different food styles to preserve the dignity of everyone in society. A fair representation of food styles requires that different styles are similarly priced, to give people a freer choice of which food style to follow (Korthals 2012: 104-7). A tax on lower-welfare animal products, which are invariably cheaper than higher-welfare ones, denies people on lower incomes a free choice about which food style to follow and ignores the cultural and social significance that animal products may hold for them. This is not to say that internalising these externalities is undesirable, but rather to point out that, in this case at least, concerns about market efficiency must be weighed against questions of justice.

Although the issue of political viability may discourage policymakers from implementing an animal welfare tax in isolation, this does not mean that externality taxes should be entirely rejected as a means of responding to animal welfare issues. Research from Chatham House indicates that climate change-motivated interventions in animal product markets, such as the levying of taxes or removal of subsidies, are more likely to win public support, particularly when trustworthy information about the impact of meat and dairy production on climate change is provided. In addition, when the rationale for these policies was convincingly and clearly communicated, any public backlash was likely to be short-lived: people in fact expected governments to take action (Wellesley et al. 2015: 39-41). In short, a climate-motivated externality tax may be politically viable; one 2017 investor report in fact anticipates that meat taxes will become part of the political landscape over the next decade (FAIRR 2017: 2; Carrington 2017). Such taxes will have an impact upon farm animal welfare, and thus should be explored further.

We see further support for this brief change of focus from a welfare tax to an environmental tax when we consider the informational challenges facing psychological externality policies. For a government to determine the socially-optimal provision of psychological externalities, it must first calculate their size. Identifying people's subjective valuations surely requires an unrealistic level of access to information about personal preferences and attitudes. For one, this method becomes prohibitively expensive and time-consuming when dealing with externalities that impact large numbers of people. In addition, maintaining an optimal supply would be resource intensive, as a government would have to constantly respond to changing market conditions and public attitudes (Lusk and Norwood 2011: 475; Lusk 2011: 572).

Moreover, subjective valuations may not be a reliable source of information. There are some externalities to which people will simply be unable or unwilling to assign an accurate value (Perloff 1999: 688). Even in those cases where people are able to accurately calculate the value of a psychological externality, there is often reason to be sceptical of the truthfulness of their responses: those affected by negative psychological externalities have an incentive to exaggerate their dislike, as, by doing so, they induce governments to introduce higher taxes to further reduce these harms (Schotter 1997: 565). Unless governments can reliably access the information required to calculate the value of psychological externalities, it is unlikely that externality taxes will produce a socially-optimal outcome. This does not entail, however, that the policy should be rejected: externality taxes can still be used to limit the supply of goods that bear negative psychological externalities, although policymakers must accept that they cannot reliably know whether they have achieved a socially-optimal level of supply.

In contrast, the technological externalities created by animal agriculture, namely illness and pollution, are less dependent on subjective valuations and so may be easier to calculate or model. For health externalities, survey data reveals the proportion of a population that engage in a given habit, medical research exposes the probabilistic connection between the habit and the risk of disease and health services can provide information about the costs of treating the disease. From this, we can estimate the national costs of treating the disease.⁵⁸ Something similar can be done for environmental externalities. To calculate the size of an optimal emissions tax, we need to know the size of a given sector's emissions and the social cost of carbon: see Ricke et al. 2018 or Interagency Working Group on Social Cost of Greenhouse Gases, United States Government 2016 for two estimates. With this information, we can calculate a

⁵⁸ This method was used by Springmann et al. (2018a; 2018b) to estimate the costs of red and processed meat-related health issues, as discussed below.

product's carbon footprint and thus determine the optimal emissions tax. Taxes may therefore be better suited to internalising technological externalities, which are less afflicted by informational difficulties than psychological externalities.

These informational difficulties shed light on the policies that will be best suited to responding to psychological externalities. A tax is less effective at responding to psychological externalities because it is a *centralised policy*, meaning it flows from governments. For a centralised policy to achieve a socially-optimal outcome by internalising externality effects, governments must be able to determine the socially-optimal level of supply. While this is challenging when dealing with psychological externalities,⁵⁹ the value of technological externalities can often be calculated with a reasonable degree of accuracy.

Because taxes targeting animal agriculture's technological externalities do not suffer the same informational and political issues that hinder psychological externality taxes, policymakers may wish to consider whether technological externality taxes can provide a vicarious response to animal welfare psychological externalities. Springmann et al. (2018a; 2018b) designed a consumer externality tax to internalise the UK's healthcare costs associated with consumption of red and processed meat, although the authors predict numerous secondary benefits. An increase in price would likely lead to a reduction in the consumption of processed meat, which would lead to fewer processed meat-related deaths – the model predicts 6,100 fewer annual deaths, a 19.42% decrease. Reduced illness and mortality from processed meat translates into a reduction in healthcare expenditure of about 20% (£800 million). The tax is also expected to raise about £2bn towards healthcare costs. Reduced consumption would also feed into reduced production, meaning that the processed meat industry's environmental impact would be diminished by a health tax.

Springmann et al. (2016a; 2016b) use a similar methodology to propose an environment-motivated externality tax on animal agriculture. They conclude that, in high income countries, consumption taxes of about 27% on beef, 13.4% on milk, 10.7% on poultry, 8.3% on pork and 6.6% on eggs could reduce greenhouse gas emissions by about 7.1% for beef, 2.1% for milk, 2.3% for poultry, 1.3% for pork and 1.3% for eggs. The taxes are also expected to raise over £61 billion

⁵⁹ Although a centralised response may not provide an optimal response to psychological externalities, other options, namely *decentralised policies*, where each individual is given the freedom to contribute towards their desired level of the externality, may yield more success. Such policies are generally market-based, and will be discussed later.

per year for 35 high-income countries (including the UK), and could prevent over 115,000 annual deaths through the associated health benefits of reduced red meat consumption.

This does not guarantee that technological externality taxes will reduce the size of negative psychological externalities stemming from farm animal welfare, however. Environmental and health taxes will have the greatest impact upon those sectors that contribute the most to these negative externalities. A health-motivated tax will overwhelmingly impact the red and processed meat sectors due to the links between these products and cancer, stroke and diabetes. Such associations have not been found with white meat, although research in this area is limited: the WHO, who released the 2015 report detailing the associations between red and processed meat and chronic illness, has not conducted a similar study for white meat. Similarly, an environmental tax will have the biggest impact on the biggest polluters, namely the cattle industry. The UN's Food and Agriculture Organization identifies cow's milk and beef as producing higher levels of pollution, both per sector and per kilogram of protein, than pork, chicken meat and chicken eggs (Gerber et al. 2013: 15-17). This is a result of ruminants' enteric fermentation, which is responsible for about 82% of agricultural methane emissions (DEFRA 2011: 10). Nitrous oxide emissions, which are produced from waste management and the production of animal feed (Sonesson et al. 2009: 7), do not appear to have been analysed on a sector-by-sector basis.

We can thus anticipate health- and environment-motivated externality taxes to primarily reduce the production and consumption of red meat. The impact on psychological externalities relating to farm animal welfare will therefore be greatest where there exists significant concern for ruminant welfare relative to other farm animals. In the UK, however, concern is greater for the welfare of chickens than of cattle and so these taxes will likely do little to internalise animal welfare psychological externalities (Clark et al. 2017: 122; Eurobarometer 2005: 105; IGD 2007: 20-21; Brook Lyndhurst 2012: 24).⁶⁰

It is also likely that these taxes would have a limited effect upon animal suffering in the UK, due to the greater number of animals involved in poultry production relative to cattle production (Reese 2018: 60): in 2018, the UK produced 188 million poultry, in contrast to 9.9 million cattle (DEFRA 2019a: 4). Environment and health taxes may even exacerbate negative psychological

⁶⁰ The surveys and focus groups that support this claim were conducted in 2005 and 2007 respectively. Consequently, perceptions may have changed, especially in light of the EU's 2012 ban on battery cages. Up-to-date research on consumer concerns about farm animal welfare will give us a better idea of the impact that an environment- or health-motivated meat tax could have upon psychological externalities.

welfare externalities and increase animal suffering if more expensive red meat leads to increased poultry consumption, as one elasticity analysis indicates is likely (Tiffin et al. 2011: 43).

In addition to potentially exacerbating psychological externalities and animal suffering, an environment-motivated animal product tax could also discourage higher-welfare agriculture. The enclosed nature of caged and barn egg systems better enables manure-related ammonia and dust to be contained than in higher-welfare organic and free-range systems. Additionally, more extensive heating is required in cage-free systems due to their lower stocking densities. This requires a greater use of energy during periods of cold weather (Xin et al. 2011: 271-72; Singer and Mason 2006: 106); depending on how this energy is produced, this can contribute to greenhouse gas emissions. Consequently, cage-free eggs could face higher environmental externality taxes than caged alternatives, meaning that these taxes may inadvertently encourage lower-welfare production. Although animal agriculture's technological externalities are serious problems that warrant attention for their own sakes, taxes introduced to this end cannot be relied upon to either reduce animal suffering or internalise animal welfare psychological externalities.

In short, there are limits to how externality taxes can be used to improve animal welfare. First, animal welfare must be understood as a negative psychological externality, rather than a positive one: a reduction in lower-welfare animal production must be viewed as desirable. Second, the value of psychological externalities is difficult to calculate, rendering it almost impossible to determine the optimal level of a welfare tax. Third, a welfare-motivated tax would likely need to be quite large to meaningfully affect consumer behaviour, which raises concerns about political viability.

Although they may be more politically viable than a welfare tax, health and environment taxes are unlikely to address negative psychological externalities stemming from farm animal welfare. The greatest impact of these taxes would fall upon red meat sectors, while, in the UK at least, people report greater concern for poultry welfare; a health- or environment-motivated tax is more likely to exacerbate animal welfare psychological externalities by encouraging consumers to substitute poultry for red meat. This does not mean that there are no good reasons to implement a health- or environment-motivated animal product tax, only that animal welfare is rarely one of them.

I also raised concerns about the societal impact of an animal welfare tax introduced in isolation: such a policy may exclude people on lower incomes from consuming animal products, and thus raises questions about food justice. This problem arises because an externality tax discourages

the purchase of lower-welfare, cheaper animal products, but does not encourage the consumption of other products. One way of resolving this issue may therefore be to pair an externality tax with a policy, such as an externality subsidy, that encourages the consumption of other types of animal product.

6. Consumer Subsidies

Subsidies increase demand for a good by lowering its consumer price or increasing its producer price.⁶¹ Because subsidies increase demand, they are most effective when used to target goods and production methods that create positive externalities. When applied to animal product markets, subsidies should therefore be used to encourage a greater supply of, or demand for, higher-welfare goods. Animal welfare subsidies can also be used to lower the costs of charitable donations.

a. Product Subsidies

Product subsidies can be granted to either producers or consumers. A producer subsidy ameliorates the costs of higher-welfare production by granting money to producers who meet certain welfare standards or products.⁶² Several countries already use subsidies in this way. The RAUS⁶³ and BTS⁶⁴ programmes in Switzerland grant subsidies to producers who provide their animals with outdoor access and high-quality housing (Voegler 2017: 32; Mann 2005a: 369; Mann 2005b: 142), while the German state of Lower Saxony subsidises producers who refrain from noncurative operations such as tail docking and beak trimming (Voegler 2017: 32). Animal welfare subsidies may also feature in the UK's post-Brexit agricultural policy. Currently, farming subsidies are largely determined by the EU Common Agricultural Payment (CAP) system of direct payments based on the amount of land farmed, but Brexit provides the opportunity for the UK to forge its own path. The proposed *Agriculture Bill* [HC] (Bill 266: 2017-19) seeks to make subsidies more dependent on producers' contributions to a range of public goods, including 'protecting or improving the health or welfare of livestock'. Despite calling animal welfare a public good, the bill treats farm animal welfare as an externality: encouraging and

⁶¹ A graphical treatment of subsidies can be found in Appendix III.

⁶² Consumer subsidies subsidise end products.

⁶³ Regelmäßiger Auslauf im Freien (roughly 'regular outdoor exercise') – a subsidy program for farms that provide their livestock with regular outdoor exercise, i.e. free-range production.

⁶⁴ Besonders Tierfreundliche Stallhaltungssysteme (roughly 'particularly animal-friendly stable systems') – a subsidy program for farms that provide their livestock with animal-friendly housing.

discouraging production and consumption is a hallmark of externality policies, while public good policies ensure the direct provision of goods.

Subsidies may be more politically viable than taxes. Two factors cause people to react more negatively to taxes than to subsidies: an apparent intrinsic aversion to taxes (Olivola and Sussman 2015: 570, 573) and *loss aversion*, the attachment of greater significance to losses than to gains (Kahneman 2012: 282-86). As subsidies make a wider range of goods accessible to people on lower incomes, issues of food justice also seem less relevant.

Nevertheless, the political viability of subsidies holds true only to an extent: if a subsidy has to be particularly large – and thus expensive – to be effective, the policy will not be politically viable, especially when funded by general tax revenue. John McInerney argues that consumers will be largely unresponsive to price changes in higher-welfare goods; consequently, subsidies must be large to be effective.⁶⁵ To understand why, we need to look at the groups McInerney thinks will (and will not) be affected by a price decrease. He states that ethically-motivated consumers already avoid low-welfare goods and so cannot respond to price decreases by substituting higher-welfare animal products for lower-welfare ones. McInerney argues that consumers of lower-welfare goods will also be unresponsive to small price decreases because they do not attach great significance to animal welfare; if they cared about farm animal welfare, they would not buy lower-welfare goods. It is only when the price premium between higher- and lower-welfare goods is eliminated, McInerney claims, that price decreases will have a meaningful effect upon demand (McInerney 2004: 14).

There are, however, several ways to contest McInerney's argument. First, he focuses on substitution effects and affords little attention to income effects. In this example, substitution effects refer to the relative desirability of higher- and lower-welfare goods, and how this changes as the prices of these goods change. When higher-welfare goods become cheaper, people who

⁶⁵ At first glance, we might think that a subsidy must be large to be effective simply because demand for animal products is generally inelastic (Tiffin et al. 2011: 16-17). Elasticity, however, is not as salient for subsidies as for taxes, due to the two policies' different aims. A tax seeks to discourage people from consuming lower-welfare goods, and does not seek to make higher-welfare goods more appealing. When assessing the effects of a tax, we consider *own-price elasticity*, how demand for a product changes as its price changes. A welfare subsidy, on the other hand, encourages shoppers to substitute higher-welfare goods for lower-welfare ones. To assess the impact of this policy, it is more appropriate to consider *cross-price elasticity*, how demand for a good responds to price changes in another good. Studies of pork consumption in Scotland have shown that when the price of animal-friendly pork increases, shoppers reduce their purchases and increase purchases of organic or conventionally-produced pork, indicating that higher- and lower-welfare pork are treated as substitutes (Akaichi and Revoredo-Giha 2016). Although perceptions of substitutability are not always symmetrical (milk may be viewed as a substitute for a milkshake, for instance, but milkshakes may not be viewed as a substitute for milk), I assume symmetrical substitution between higher- and lower-welfare goods.

value welfare attributes might find them more attractive relative to lower-welfare goods. Income effects will also occur, however. When higher-welfare goods are subsidised, existing consumers are made better-off – their income now allows them to buy more. With more money to spend, these consumers will buy more of the products they like, including higher-welfare goods. Welfare subsidies could therefore encourage ethically-motivated consumers to increase their purchases of higher-welfare goods. Income effects are likely to be small, however, due to the own-price inelasticity of animal products.

Second, as McNerney himself recognises (2004: 39), some consumers may be willing to pay more for higher-welfare goods, but are unwilling to pay current prices; these people may refuse to participate in animal product markets. A subsidy might entice some of these ‘exiled compassionate omnivores’ back into animal product markets, even if it does not eliminate the price premium between higher- and lower-welfare goods. McNerney mistakenly assumes that there are no substitutes for higher-welfare goods. This is not the case: meat substitutes such as Quorn, as well as products such as tofu, pulses and more, can all be viewed, to a greater or lesser degree, as substitutes for animal products. A welfare subsidy could therefore encourage consumers of these goods to consume higher-welfare animal products. To calculate this effect, we need to know how many people abstain from participating in animal product markets for price reasons, and the cross-price elasticity between higher-welfare goods and non-animal product substitutes; such information was not accessible at the time of writing. Chapter 1’s discussion of the meat paradox (p.30) suggests that this effect will be small, however, because many consumers dissociate living animals from their produce or simply do not think about ethical issues in purchase environments. Instead of opting out of the market due to price concerns, consumers are perhaps more likely to overcome or ignore their qualms about eating lower-welfare produce.

Third, McNerney’s argument assumes a model of the consumer in animal product markets that could be described as unrealistic. He states that consumers of lower-welfare products will not be affected by a decrease in the price of higher-welfare goods as it is improbable that they care about animal welfare. Even if a consumer of low-welfare goods does not directly care about animal welfare, however, McNerney’s conclusion does not necessarily follow. We saw in Chapter 1 (p.23) that many consumers hold self-interested preferences for higher animal welfare, due to perceived associations with nutrition, safety, quality and taste. These associations, whether justified or not, entail that a decrease in the price of higher-welfare goods can appeal to consumers of lower-welfare products, irrespective of direct concern for animal welfare. The size of this effect is dependent upon both the number of consumers who associate animal

welfare with other quality attributes and the cross-price elasticity of higher- and low-welfare goods, i.e. how far consumers perceive the two types of good to be substitutes.

Fourth, McInerney's account of ethically-motivated behaviours in this domain is also suspect. He appears to assume that consumers have well-developed and consistent ethical views that they are aware of and use to make ethically-informed purchases: people who care about animal welfare will not buy low-welfare goods, and people who do not care about animal welfare will. The reality seems less straightforward. As we saw in Chapter 1 (p.30), and will return to in greater detail in Chapter 7, consumers are often detached from their ethical views in animal product markets. Despite stating ethical concern for farm animal welfare, many consumers avoid thinking about the ultimate fate of a farm animal at points of purchase (Harper and Henson 2001: 11). This detachment allows consumers who care about animal welfare to still buy low-welfare goods.

When consumers do think about farm animal welfare, the focus is primarily on positive elements of animal agriculture that contribute to the idea that animals have lived happy lives (IGD 2007: 16; Brook Lyndhurst 2010: 40); this suggests that higher-welfare purchases are frequently associated with positive psychological externalities.⁶⁶ Consequently, some higher-welfare purchases may be motivated by impure altruism, the personal pleasures associated with contributing to higher-welfare goods. When a subsidy reduces the price of higher-welfare goods, the price of achieving this feeling of warm glow also decreases, making higher-welfare purchases more attractive even when they are not price-equivalent with lower-welfare ones. This argument is effectively an extension of the previous one, treating the intrinsic satisfaction of a higher-welfare purchase as a product attribute akin to nutrition, safety and taste. It thus extends the reach of the previous argument to cover consumers with ethical preferences for farm animal welfare.

For these reasons, subsidies could in fact be effective in encouraging consumers to substitute higher-welfare goods for lower-welfare ones, even when they fail to make higher-welfare goods price-competitive. They do not have to be large in order to have a meaningful effect and may therefore be a viable method of increasing the supply of higher-welfare goods.

Even if subsidies did need to be large to be effective, they can be combined with other policies to mitigate their impact on government spending. As noted above, an externality tax might be

⁶⁶ As argued in Chapter 1 (p.30), detachment from welfare issues renders it entirely conceivable that a consumer can fail to consider animal welfare when purchasing lower-welfare goods, but still derive satisfaction from their contribution to animal welfare when buying higher-welfare ones.

combined with an externality subsidy. A tax on low-welfare production would raise revenue that could contribute to the costs of a welfare subsidy. This can address concerns about food justice by reducing disruption to consumption patterns and better-enabling people of all levels of income to freely choose their food styles.

As with externality taxes, however, several factors can limit the effectiveness of an externality subsidy. First, a badly designed subsidy may not achieve its objectives. Although McInerney (2004: 53) argues that subsidies granted to producers will be as effective as end-product (consumer) subsidies, David Harvey and Carmen Hubbard (2013: 113) suggest that, as the structure of a market can encourage producers to hold onto the subsidy and fail to pass any of it on to consumers, producer subsidies could be less effective, and might in fact hinder the further development of higher-welfare products. Consequently, policymakers might want to focus more on consumer subsidies and move away from the current system of direct payments to producers.

Second, a subsidy that is applied inconsistently may not achieve its objectives. If, say, an EU government wanted to move egg producers away from battery cage production in the run-up to the 2012 legislative ban, they might have introduced a subsidy for enriched cage production. Such a subsidy would be of overall detriment to hen welfare if it was large enough to encourage barn, free-range and organic producers to switch to lower-welfare enriched cage systems. An effective subsidy must therefore be applied to all relevant forms of production.

Third, participation is voluntary: subsidies cannot force producers to provide their animals with higher welfare if they do not want to. While, in many respects, the voluntariness of this policy can be perceived as an advantage, it may also hinder its effectiveness. To this end, subsidy payments must be large enough to provide a genuine incentive for producers to consider higher-welfare production methods. In some sectors, this may not be a serious concern; agricultural subsidies around the world keep food prices at an artificially low level, and often leave farmers dependent upon them in order to remain competitive (Pollan 2011: 52-55). In the UK, subsidies make up a significant proportion of farm income, and, excepting horticulture, specialist pig farms and poultry farms, provided the largest stream of farm revenue in 2016/17 (DEFRA 2017a). In the livestock and dairy industries, where subsidies made up 88% and 51% of business income respectively, changing the terms of these subsidies – as proposed in the *Agriculture Bill* [HC] (Bill 266: 2017-19) – would effectively force producers to comply: those who do not are unlikely to remain competitive and will be forced out of the market. This may not, however, be so effective in the pig and poultry industries, where producers are less dependent on subsidies:

they comprised 19% and 14% of business income respectively in 2016/17. In addition, Agriculture and Horticulture Development Board research (2017: 22-23) anticipates that pig farmers' incomes will increase after Brexit, meaning that subsidies are unlikely to become a more important revenue stream in the near future.⁶⁷

Fourth, as with other centralised policies, informational issues can limit a subsidy's success in achieving socially-optimal levels of supply when dealing with psychological externalities. Externality subsidies can still be used, however, to encourage a given level of supply, but policymakers must accept that they cannot know whether this supply is socially-optimal.

b. Charity Subsidies

In addition to using subsidies to encourage higher-welfare production, subsidies can be used to target charitable donations. Governments in many countries (including the UK) offer tax relief on charitable donations (Nechyba 2017: 1050-51). Tax relief reduces the costs of donations, which encourages more people to donate more money (Roberts 1987: 420). Tax relief could therefore encourage donations to charities that support farm animal welfare. Charities support farm animal welfare in a myriad of ways, including lobbying for new legislation (World Animal Protection), investigating and prosecuting those who break animal welfare laws (RSPCA), raising awareness of welfare issues (CIWF), operating high-welfare labelling schemes (Soil Association), producing research and policy assessments (UFAW), providing training in high-welfare methods for those working with animals (HSA), encouraging public education on welfare issues (BHWT) and rehabilitating animals (Thornberry Animal Sanctuary).⁶⁸

Because charities perform so many different functions relating to farm animal welfare, tax relief enables consumers to contribute to the issues that matter most to them. Every year, the British Hen Welfare Trust (BHWT) finds homes for about 60,000 hens, which, at two years of age, are at the end of their commercial laying lives and would otherwise be sent to slaughter (BHWT 2019). The BHWT's work does not, however, directly alleviate the animal suffering and associated psychological externalities created by egg production, but rather provides life to some of the sector's 'waste' products. Consequently, a person's attitudes towards farm animal welfare will determine whether they support the BHWT. If Michael is more responsive to animal

⁶⁷ The impact on poultry farmers' income is not discussed in the report, so it is possible that this sector will become more dependent on subsidies post-Brexit.

⁶⁸ Many of these charities improve farm animal welfare in multiple ways: the RSPCA, for instance, performs all of the above functions.

suffering than animal happiness, for instance, he is unlikely to support the BHWT's efforts because it does little to relieve suffering in production systems. The BHWT's work will primarily be regarded as effective by those who are concerned about whether farm animals are given good lives. Yet even for these people, the BHWT's work may leave concerns unaddressed as it does not improve animal welfare within production. There are, however, a plethora of other animal welfare charities that work on production-side welfare issues and can therefore address animal agriculture's negative psychological externalities.

Subsidising charitable donations thus represents a highly equitable approach, as it enables people to directly address the elements of animal agriculture from which their psychological externalities spring. The policy also facilitates the expression of a range of ethical attitudes: abolitionists with views similar to Gary Francione can express their ethics by contributing to charities that promote vegan lifestyles, while those concerned about animal welfare can donate to charities that pursue this cause both on the farm and in the corridors of power.⁶⁹

Consumer subsidies are limited, however, by the fact that they cannot entirely eliminate free rider incentives. The existence of good farm animal welfare in society can be of satisfaction even to people who do little to support higher welfare. Consumers can thus reap the psychological benefits of higher welfare without contributing to the associated costs. Even though tax relief for charitable donations reduces the costs of charitable actions, it cannot eliminate the consumer incentive to enjoy the benefits of better welfare without contributing – this is known as free riding. As tax relief does not remove the opportunity to free-ride, it can only ever be a partial response to psychological externalities and is best-used in conjunction with other policies.

Further problems are raised by informational issues. Different elements of animal welfare may require different levels of subsidy to reach an optimal level of provision. This adds to the complexity of providing an optimal supply of animal welfare, and, more importantly, may not be a politically viable option: governments that adopt a policy of variable tax breaks may be accused of playing favourites with different causes. Fortunately, this political dilemma is avoided by the same informational difficulties that prevent other centralised policies from achieving a socially-optimal level of supply. As with welfare-motivated externality taxes and product subsidies, a more realistic use of this policy is to encourage an increase or decrease in

⁶⁹ Indeed, charitable donations appear to be an effective way for people who do not participate in animal product markets to express concern for farm animal welfare: one study in Belgium found that vegetarians donate more to animal charities than meat eaters and flexitarians, with no significant difference between each group for donations to human charities (De Backer and Hudders 2015).

supply, rather than provide an optimal level. For this, equally subsidising charitable donations irrespective of their cause may be an appropriate policy.

7. Cap and Trade

Cap and trade, commonly used to address pollution, offers another externality policy response. When used to tackle pollution, a government must first determine the maximum level of pollution that it deems acceptable (the cap). It then distributes or sells permits to firms which entitle them to emit a certain amount of pollution, totalling the government's desired level (Parkin 2000: 437). Permits may be sold to competitors if a producer has more than it needs. Heavy polluters find themselves having to buy permits, which increases their production costs, while less-polluting firms can sell their permits to subsidise theirs. This gives producers an incentive to adopt low-pollution technology, as doing so reduces the need to purchase pollution permits, and the selling of permits opens up a new revenue stream.

Cap and trade is intended to reduce negative externalities and so will be most effective when used to limit the production of lower-welfare animal products, rather than increasing the supply of positive externalities. This is not to say that a cap and trade policy cannot increase the supply of higher-welfare goods, only that this is achieved by making lower-welfare goods less appealing.

An animal welfare cap and trade could work by either allocating or selling permits which allow producers to use a given low-welfare practice, such as a higher stocking density, upon a certain number of animals, with a centralised body determining the level at which the practice ought to be capped. Producers who do not use their full allocation could sell to those who face a shortage. Additionally, animal welfare charities could purchase permits and withdraw them from circulation, limiting the extent of the production practice. Cap and trade discourages the use of a lower-welfare practice by making it more expensive relative to other methods of production: it can subsidise higher-welfare producers by granting them the opportunity to sell excess permits, and also serves as another cost for those producers who must buy permits.

There are, however, several concerns regarding the use of cap and trade policies as a means of addressing animal welfare psychological externalities. First, cap and trade faces the same informational problems that plague externality taxes and subsidies. Given that cap and trade creates a market in permits, one could be forgiven for assuming that it is a decentralised response. This assumption is incorrect, however, as, although market mechanisms will efficiently distribute a fixed level of supply, a centralised body must first set this level of supply.

Governments are unlikely to possess full information when calculating the value of psychological externalities, which entails that using a cap and trade in this context will not lead to an optimal outcome (Nechyba 2017: 750-51). This is not to say that there is no place for cap and trade in responding to psychological externalities: as with welfare-motivated externality taxes and subsidies, the policy is still an effective tool for reaching a desired level of supply, but cannot guarantee a socially-optimal supply.

As with externality taxes, cap and trade is likely better-suited to dealing with the technological externalities produced by animal product industries, namely pollution. Many countries, such as New Zealand, already run pollution cap and trade programs which are able to incorporate animal industries (Bailey et al. 2014: 8). New Zealand's incorporation of livestock emissions in its cap and trade scheme can only be regarded as partial, however, as the country requires the reporting of agricultural emissions but not trading, meaning that there is significant scope for policy development in this area.

There are numerous technical concerns about how the design of emissions trading systems can affect their efficacy in reducing greenhouse gas emissions, and whether it is appropriate to use this policy in agricultural industries. A discussion of these issues is beyond the scope of this thesis, but can be found in Ritter 2017, Thomassin 2003, Bullock 2012 and Ancev 2011. My interest lies with how the policy can influence animal welfare psychological externalities, and in this regard it appears that cap and trade faces the same concerns as environment- and health-motivated externality taxes. As with a pollution tax, a pollution cap and trade should have the greatest impact on the biggest polluters, i.e. red meat sectors. Consequently, the policy will likely encourage consumers to substitute white meat for red meat, and so will only address animal welfare psychological externalities in societies where there is significant concern for livestock welfare and relatively little concern for poultry welfare. In the UK, evidence suggests that consumers are most concerned about welfare in poultry sectors (Clark et al. 2017: 122; Eurobarometer 2005: 105; IGD 2007: 20-21; Brook Lyndhurst 2012: 24), meaning that a cap and trade is in fact likely to exacerbate negative psychological externalities if it causes consumers to substitute white meat for red meat. As with environment- and health-motivated taxes, this is not to say that there are no good reasons to use an emissions cap and trade in animal agriculture, only that this tool may not be an effective way of addressing animal welfare externalities.

While there are clear difficulties associated with using cap and trade to internalise psychological externalities, the policy may be useful in phasing out extremely low-welfare practices. Naturally, it would be inappropriate to use a cap and trade policy in a manner that maintained any use of

extremely low-welfare practice in the long term. The long-term solution to any such practice must be prohibition, but cap and trade could be used as a tool to facilitate a transition away from extremely low-welfare practice. This may be appropriate in cases where introducing an immediate ban would impose large costs upon producers and so could face significant political obstacles. Because governments control the level at which a cap is set, they can gradually reduce the supply of permits over time, which would cause them to become more expensive. Increasing permit prices will raise costs for those producers who remain reliant upon extremely low-welfare practices, and thus provides an increasingly strong incentive to adopt other production methods.

Such a use of cap and trade might have been able to increase compliance with the EU's 2012 ban on battery cages, which could not be described as an immediate success. Despite there being more than twelve years between the legislation being passed and the ban coming into effect, an estimated 14% of EU hens were still in battery cages at the start of 2012 (Blandford and Harvey 2014: 37-38; BBC Newsround 2012; BBC 2012). If a cap and trade had been introduced upon the passage of the EU's Council Directive 1999/74/EC of July 1999, the costs associated with battery production could have been slowly raised over the following twelve years, providing an additional incentive for producers to adopt other methods of production and be ready for the ban's implementation in 2012.

In short, cap and trade represents another limited way of addressing animal welfare externalities. Informational issues limit the policy's use in addressing psychological externalities, and, when used to address animal agriculture's technological externalities, might even exacerbate concern for welfare. Despite these concerns, cap and trade could be an effective way of transitioning away from extremely low-welfare practices before bans come into effect.

8. Market Creation

A fourth externality policy involves creating new markets by better-defining property rights. This idea was first proposed by Ronald H. Coase in his 1960 paper 'The Problem of Social Cost'. Coase recognised that, where property rights for a production input remain undefined, firms have little incentive to treat the input as a cost in the production process; this tends to lead to its overuse. Not having to pay for emitting air pollution, for instance, effectively turns clean air into a free input in the production process. Firms will emit more air pollution than they would if clean air was owned by someone, and thus needed to be paid for. If property rights are

defined,⁷⁰ and if those affected by the externality can engage in costless negotiations, an efficient solution can be arrived at. To illustrate, suppose that there are two firms operating on a river: a cattle farm and a water treatment facility. The cattle farm releases pollution into the river, and an increase in its production therefore increases the costs of treating the water, which the farm does not cover; the farm thus imposes a negative externality upon the water treatment facility. Assume that the cattle farm considers reducing production by 5000kg (about five cows) and calculates that doing so would cost £16,000 in lost revenue (AHDB 2019a). The cattle farm's actions would also reduce the water treatment plant's costs by £20,000, which makes a reduction in beef production a Kaldor-Hicks efficient action; the losses of the farm are outweighed by the gains of the water treatment plant (Schotter 1997: 578). There are currently no clearly-defined property rights in this example, however: without knowing whether the cattle farm has the right to pollute the river or whether the water treatment plant has a right to pollution-free water, the two parties will not negotiate (Perloff 1999: 673). The cattle farm will continue producing at a profit-maximising level, which imposes a substantial negative externality upon the water treatment facility and leads to a suboptimal outcome. If a government defines property rights, however, the firms have an incentive to negotiate and arrive at an optimal solution. Notably, it does not matter who gets the property rights; as long as they are defined, and costless negotiation is possible, a Kaldor-Hicks efficiency improvement can be achieved.

Assume that the cattle farm is given the right to pollute the river. This creates a market in the emission of pollution into the river. We have seen that if the farm reduces production by 5,000kg, which leads to a £16,000 reduction in its revenue, the water treatment plant stands to save £20,000. The water treatment facility can therefore maximise its profits by offering the cattle farm between £16,000 and £20,000 to reduce production. Doing so leaves both firms better off, or at least no worse off than they would have been had a solution not been agreed upon (Mansfield 1985: 501). Alternatively, assume that the water treatment plant is given the right to use the water as it chooses. In this case, there is still scope for the cattle farm to emit pollution into the river. The water treatment plant will accept payment for each unit of pollution emitted up to the point that the cost of treating the water exceeds the price of another unit of pollution, i.e. where the marginal revenue gained by allowing another unit of pollution to be emitted exceeds the marginal cost of treating another unit of water. Similarly, the farm will treat the emission of pollution into the river as another production input, and will produce beef up to

⁷⁰ The distribution of property rights can be a result of firms' negotiations or government intervention.

the point where the cost of producing another unit equals the revenue it brings in. Regardless of who owns the property rights, a Kaldor-Hicks improvement will be achieved if the parties can negotiate without cost (Schotter 1997: 578).⁷¹

Defining property rights provides a solution to some externality problems because it recognises that these externalities are connected to a ‘missing market’ elsewhere (Nechyba 2017: 756). The creation of new markets through the definition of property rights forces firms and individuals to consider externalities as costly inputs, and can thus internalise them (Salvatore 2003: 609).

Market creation can also internalise positive externalities. Earlier in this chapter, we saw that positive externalities exist between beekeepers and orchard owners. The bees pollinate the orchard, while the orchard provides nectar for the bees (Meade 1952: 56). The orchard owner could – at least in theory – seek to charge the beekeeper for using it as an input in her production of honey. This new income stream may encourage him to expand his orchard to the point where marginal social benefits equal marginal costs. Similarly, the beekeeper could charge for the pollination services her bees provide. This new income stream could encourage her to expand her beehives, again to the point that the marginal social benefits are equal to the marginal costs.

When applied to issues of farm animal welfare, we see that creating markets for animal welfare overcomes the informational difficulties that centralised policies face, as doing so enables people to contribute to a personally optimal level of animal welfare. For this reason, market creation is a decentralised action and could be an effective means of addressing the psychological externalities, both positive and negative, created by animal agriculture.

To create a new market in this way, someone, or some group, must either be granted, or assume ownership of animal welfare. There appear to be three main ways of doing this. First, the producers of animal products could assume ownership of animal welfare and charge consumers to deliver higher welfare (Lusk 2011). J.L. Lusk suggests that firms are often discouraged from providing higher levels of animal welfare because they cannot directly extract payment from consumers for doing so, and therefore interprets animal welfare as a positive externality. Lusk proposes the creation of ‘animal well-being units’ (AWBUs), which are distributed to producers in proportion to the number of animals raised and the quality of their living conditions. Producers would be able to sell these AWBUs in a market that existed independent of markets

⁷¹ While the party that possesses the property rights appears to have an advantage in negotiations and is thus likely to secure a better deal for itself, this is a problem with distribution rather than efficiency (Nicholson and Snyder 2010: 525).

for animal products (Lusk 2011: 565).⁷² Lusk assumes that people and groups that are concerned about animal welfare but unwilling to buy animal products would then buy these credits, and market mechanisms would promote the efficient pricing and supply of AWBUs, and thus animal welfare (Lusk 2011: 564). The AWBU may also promote innovations in animal welfare by encouraging producers to look for ways to provide better animal welfare at lower costs (Lusk 2011: 571).

Lusk does not explicitly explore how an AWBU market may influence markets for animal products, but it is plausible to suggest that, if it functions as he intends, the decoupling of welfare from product would contribute to a short-run reduction in the consumer price of higher-welfare products. A large part of what makes higher-welfare goods more expensive than conventional ones is the higher costs of providing more welfare, and so the creation of a separate revenue stream for higher-welfare could subsidise animal products and allow producers to charge less for them without harming profit margins. This could encourage consumers to substitute higher-welfare goods for lower-welfare ones, if not for welfare-related reasons then for these goods' perceived superior quality and safety attributes. Access to a higher-welfare revenue stream, coupled with an increase in demand for higher-welfare goods, could also encourage producers to move away from low-welfare practices and towards higher-welfare ones.

There are, however, two issues that undermine the AWBU market's effectiveness as a response to animal welfare psychological externalities. First, as Lusk himself recognises, it fails to fully internalise externality effects. To explore this point, we need to consider who is expected to participate in the AWBU market. For obvious reasons, people who do not care about farm animal welfare will not participate. Consumers of animal products are also unlikely to participate; although these people may have preferences for animal welfare, they are likely to primarily express these preferences in animal product markets (Harper and Henson 2001: 11). Lusk instead suggests (2011: 566) that the viability of the AWBU market relies upon those who care about animal welfare but refrain from the consumption of animal products; the intention of the AWBU market is to provide this group with the opportunity to express their preferences in a market setting, and thus internalise their psychological externalities.

Within the group of people who care about animal welfare but abstain from consuming animal products, however, there will be plenty who will also abstain from the AWBU market. For

⁷² Although the market for AWBUs may be able to function without government involvement, there may still be a role for policymakers to play in the definition of AWBUs, as well as in supporting efforts to ensure transparency and accountability through monitoring and enforcement.

abolitionists like Gary Francione and Anna Charlton (2015), who oppose all forms of animal ownership, or rights-based theorists such as Tom Regan (2004), who believe that farm animals have a right not to be used merely for human ends, participation in the AWBU market is to implicitly endorse the continued treatment of animals as property and human resources, which they are likely to find an unpalatable proposition; as Lusk admits (2011: 566), we expect them to abstain. Lusk's proposal will therefore not fully internalise animal welfare externalities because many people who care about animal welfare will continue to lack an appropriate way of expressing their concern.

Although the primary purpose of the new market is to internalise externalities, rather than to improve animal welfare, it is still worth considering how the AWBU market might impact the lives of farm animals. It seems likely that, because we cannot expect many people to participate in Lusk's AWBU market, it will have little meaningful impact upon farm animal welfare. Expected participants are likely to identify as vegetarians or vegans, a group which, according to a recent report by the supermarket Waitrose, makes up 12.5% of the UK's population (Waitrose and Partners 2018: 6-7). Within this group, participants must have some sort of concern for animal welfare – 55% state that animal welfare was one of the motivating factors in their turning vegetarian or vegan, a finding which corresponds with other surveys (e.g. Mintel 2017e).⁷³ At this point, just under 7% of the population are possible participants in the AWBU market, which seems promising. However, 49% of those identifying as vegetarian or vegan in the Waitrose report admitted to eating meat occasionally, at weekends, or on special occasions. These people are unlikely to participate in the AWBU market, as they can express preferences for greater animal welfare in animal product markets, where their purchases will carry greater private benefits. Ruling this group out leaves a potential participant pool of about 3.4% of the population. Some of this group will likely oppose animal ownership and the use of animals as merely human means, however, or will view participation in the AWBU market as a way of subsidising the consumption habits of others. Numbers for these groups are not available, meaning it is not clear what proportion of the population are likely participants in the AWBU market, although it is likely to be small.⁷⁴ Participation is also likely to be negatively impacted by the opportunity to free ride on others' contributions to the AWBU market, due to the excludability and non-rivalry of psychological externalities.

⁷³ Respondents could provide multiple responses to this question, and so were not prevented from expressing concern for farm animal welfare as a secondary motivator.

⁷⁴ Lusk suggests (2011: 573) that animal welfare charities may also be buyers of AWBUs. However, this is unlikely to increase the number of participants in the AWBU market: charities committed to buying AWBUs are most likely to receive donations for this purpose from people who also support the AWBU market.

When low levels of market participation are combined with high numbers of farm animals, it seems unlikely that the AWBU market can meaningfully improve animal lives. In 2016 alone, approximately 2.8 million cattle, 14.6 million sheep, 11 million pigs and 1 billion poultry⁷⁵ were slaughtered in the UK (DEFRA 2016a-u, 2017b-c). When dealing with such large numbers, the contributions of the subsection of vegetarians and vegans willing to participate in the AWBU market will surely have little effect upon farm animals' lives.

The failure of Lusk's AWBU market to internalise psychological externalities or lead to meaningful improvements in farm animal welfare does not mean that the possibilities of market creation have been exhausted. Lusk's AWBU market relies upon the producers of animal products assuming ownership of animal welfare, but ownership could be given to or assumed by other parties. One alternative is for a government to grant ownership of animal welfare to the general public. Such a distribution forces producers to treat low animal welfare as a production input, the use of which requires compensation to be paid to the public. This effectively amounts to an externality tax on the use of low welfare as a production input, increasing the costs of low-welfare production and reducing the supply of low-welfare goods.

Consequently, the informational difficulties that hinder the efficiency of externality taxes are again relevant. It is also unlikely that public ownership of animal welfare can internalise all animal welfare externalities. This policy is unlikely to be palatable to many abolitionists and supporters of animal rights as it implicitly condones the use of animals for human ends; even if the market contributes to reductions in animal harm, it may create new psychological externalities if people are dissatisfied by it. It is also probable that the pricing of animal welfare will not be agreeable to all members of society, and so their psychological externalities will not be fully internalised. Consequently, this distribution of property rights cannot be relied upon to fully-internalise psychological externalities.

A final way that animal welfare property rights might be distributed is to grant them to the farm animals themselves. By doing this, policymakers would be effectively adopting a weak version of the animal rights position by denying producers and consumers the opportunity to treat farm animals merely as means to human ends. The implications are potentially radical: animal interests would need to be considered by producers *for the sakes of the animals themselves*, rather than for economic reasons. This means that certain common practices, such as the egg industry's culling of male chicks, typically by gassing them (Saul 2015), or the dairy industry's

⁷⁵ The poultry figures cover birds bred for their meat as well as 51.4 million 'boiling fowl', a category that includes spent laying hens and spent breeders.

removal of male calves from their mothers shortly after birth and sending them to slaughter (Levitt 2018a), would likely be prohibited. In short, giving farm animals ownership their own welfare would prohibit producers from inflicting harms upon them solely for commercial gain, meaning that animal agriculture could only exist where it was also clearly in the animal's interests.

Such a distribution of rights would not create a market in the way that Lusk's AWBU or public ownership of farm animal welfare does; we are not buying or selling animal welfare. In fact, consumers do not feature at all in this market. Instead, this distribution essentially creates a two-good barter economy, where producers supply animal welfare and animals supply their produce. Producers seek profit and farm animals seek a life worth living. Granting animals ownership of their welfare thus does not have to entail the end of animal farming. If an animal's existence can be said to be of good quality, then it may be in the animals' interests to live and there may be scope for farm animals to bargain with producers in order to secure a certain standard of welfare.⁷⁶

Ultimately, this bargaining would likely lead farm animals to sacrifice elements of welfare in order to be given, as a bare minimum, a life worth living. This minimum goal corresponds with the recommendations of the Farm Animal Welfare Committee (FAWC), which defines a life worth living as one where 'the balance of an animal's experience [is] positive over its lifetime', any suffering is 'necessary, proportionate and minimal' and where animals' needs and some of their wants are provided for (FAWC 2009: 14).⁷⁷ This is not the same as a good life, which requires a much higher standard of living (FAWC 2009: 16) – a life worth living is likely compatible with various low-welfare practices. A life worth living is just a baseline, however: negotiations between stakeholders may lead to higher levels of welfare, up to a maximum point where further welfare improvements are not economically viable. The effectiveness of this

⁷⁶ This bargaining is clearly hypothetical; if such a market existed then the interests of farm animals would need to be represented by humans, perhaps by animal welfare experts with no ties to animal agriculture, who would ensure that elements of welfare are only sacrificed when it is in the animals' best interests. Such a possibility is not so farfetched: the Nonhuman Rights Project already pushes for the recognition of legal personhood for animals, which, when their rights are violated, would allow nonhumans to seek legal redress with the help of human representation (Reese 2018: 11-12).

⁷⁷ There are numerous practical difficulties associated with the concept of a life worth living. First, it is not apparent that quality of life can be treated as what John Webster calls an 'algebraic sum of positive and negative experiences'. Even if it can be, our ability to assess trade-offs between constituents of welfare in other species must surely be limited. Moreover, our judgements regarding whether an animal possesses a life worth living are made without the contribution of the animal itself; our assessments instead rely upon how we, or their human representatives, think the animal thinks and feels in a given environment (Webster 2016: 35; Yeates 2017: 25). Perhaps the most appropriate response to these issues is to err on the side of caution and aim to provide animals with lives significantly above the quality we believe necessary for a life worth living. Any mistakes about quality of life would subsequently be unlikely to deny animals a life worth living.

distribution of rights as a way of improving animal welfare will therefore be largely dependent on the outcome of negotiations: if current standards do not provide farm animals with lives worth living, animal ownership of welfare will improve standards and reduce the size of negative psychological externalities, but, if animals are already provided with lives worth living, the impact on psychological externalities (and animal welfare) will be uncertain.

Although this designation of property rights may preserve the existence of animal product markets, it may nevertheless reduce the dissatisfaction felt by supporters of animal rights. Some supporters, including those inspired by Tom Regan, believe that animals have a right not to be used merely as means to human ends. These people may be appeased when animal interests are put at the heart of animal agriculture, when animals are only raised when they can be given lives worth living; the granting of welfare rights to farm animals ensures that they are no longer treated solely as a means to human ends. Similarly, abolitionists such as Gary Francione and Anna Charlton maintain that animals have a right not to be treated exclusively as human resources. The key term here is *exclusively*, a term the two authors use repeatedly (Francione and Charlton 2015: 9, 23, 87, 97, 108, 129). When animal needs and interests are placed at the heart of animal agriculture to the detriment of economic concerns, farm animals are no longer being treated exclusively as human resources; in fact, farm animals could be said to possess a degree of self-ownership. Some abolitionists might therefore find that this designation of property rights goes some way towards addressing their concerns.

On the other hand, it would be naïve to assume that this policy overcomes all psychological externalities. Francione and Charlton (2015) would be unlikely to view this as an acceptable solution; ultimately, farm animals would still be treated as property, would likely still endure significant institutional suffering, and would continue to live lives of dependency, unable to survive by themselves in the human world or the nonhuman world. Others may hold that the killing of an animal for consumption can never be in its interests. Additionally, consequentialists who share Richard Ryder's 'painience' position (2000: 213) – which focuses on pains caused to animals but not pleasures – may be unhappy with outcomes that lean towards lives worth living, as harms to animals and low-welfare practices will continue to cause them dissatisfaction.

Consequently, granting ownership of welfare to farm animals will fail to internalise all psychological externalities. Nevertheless, the policy aligns with FAWC welfare objectives and would guarantee a baseline for animal welfare by ensuring that no farm animal is given a life not worth living.

There are, however, three further issues which may limit the effectiveness of any market creation policy. First, outcomes will only be socially-optimal if the costs of negotiation are sufficiently low. While the costs of negotiating may be low when only a small number of parties are involved, bringing large numbers of stakeholders to the negotiating table may create prohibitively large costs (Mansfield 1985: 503). This issue is particularly salient to environmental issues, where the cost of negotiation tends to be high due to the number of stakeholders involved (Nicholson 2010: 523). There will likely be many stakeholders in issues of animal welfare: producers, consumers, vegetarians and vegans, and perhaps the farm animals themselves. Second, even if negotiations are costless, an optimal outcome may not be achieved should any party lack complete information about the costs and benefits involved. As will be discussed in Chapter 6, informational asymmetries abound in animal product markets, making it harder to achieve an optimal outcome. Third, firms might engage in strategic bargaining, which militates against an agreement being reached (Perloff 1999: 675).

In short, none of the three proposed forms of market creation fully-internalise psychological externalities, although they could still bring about radical changes in animal product markets. Ownership of animal welfare could be given to or assumed by producers, who would then sell credits connected to this welfare to concerned individuals who do not consume animal products. I challenged the effectiveness of this solution on two fronts: first, supporters of animal rights would likely find such a solution unpalatable and would abstain from the AWBU market; consequently, some psychological externality effects will not be internalised by Lusk's market. Second, given their limited number, it is not apparent that the participants in Lusk's market would bring about meaningful improvements to farm animal welfare.

Alternatively, ownership of animal welfare could be granted to the public, in what essentially amounts to another way of framing an externality tax on low welfare. Perhaps most radically, ownership of animal welfare could be granted to the animals themselves. Doing this would not necessitate the end of animal product industries, provided it could be shown that it is in an animal's interests to sacrifice certain elements of welfare in order to be brought into existence. At a minimum, farm animals would be provided with lives worth living, although, depending on negotiations, they could achieve higher levels of welfare.

9. Summary

In this chapter, I have proposed that animal agriculture produces both positive and negative psychological externalities stemming from concern for farm animal welfare. I have considered

four policy responses to externalities. When governments can calculate the socially-optimal level of supply of an externality-producing good, taxes, subsidies and cap and trade policies can push markets towards these levels. When dealing with psychological externalities such as the feelings of satisfaction and dissatisfaction produced by farm animal welfare, however, it is nearly impossible for governments to calculate the socially-optimal level of supply. In such instances, centralised policies are best used to promote a given level of supply, although policymakers must recognise that they do not know whether this will be a socially-optimal one.

I also explored the possibility of using taxes, product subsidies and cap and trade to indirectly improve farm animal welfare by addressing the health and environmental technological externalities created by animal agriculture. These policies could in fact penalise higher-welfare production, and would exacerbate negative psychological externalities when they facilitate a transition from red meat consumption to white meat consumption in societies such as the UK, where welfare concerns are greatest for poultry.

Creating markets in animal welfare could be a novel way of internalising some psychological externalities. Regardless of whether ownership of animal welfare is given to producers, the public or the animals themselves, this policy is also unlikely to fully internalise psychological externalities, as people opposed to animal ownership may not be receptive to any policy that promotes the commodification of animals. Irrespective of its effectiveness as an externality policy, market creation, in particular when ownership of animal welfare is given to the animals themselves, has the potential to radically reform our relationships with farm animals and animal agriculture.

Chapter 5

Public Goods

In this Chapter

- An introduction to the concept of public goods.
- A defence of the claim that farm animal welfare does not meet the technical criteria for a public good but should, in certain circumstances, nevertheless be treated as one for policy reasons.
- A discussion of public good policies and how they can be used to provide and protect farm animal welfare.

1. Introduction

The American sociologist K.A. Davis (1949: 157-62) once noted that conflict ‘is an ever-present in human relations’ and, for many people, this conflict starts at home. Chores can be a major source of domestic conflict: who will sweep the floors, empty the bins, do the washing up? These chores are frequently not finished, leading to tension as the home becomes dirty, messy and cluttered. Paradoxically, each member of the household professes a keen desire for cleanliness. How can it be that everyone wants to live in a clean and tidy house, yet the reality is rather different?

The word economics stems from the Ancient Greek word for household management, οἰκονόμος, and the discipline can shed some light on this domestic paradox. A clean and tidy house possesses elements of a *public good*, something that all housemates will benefit from irrespective of who does the cleaning. When you can reap the benefits of a clean house without putting in any effort, the motivation for hoovering the stairs quickly ebbs away – after all, someone else will surely get around to doing it eventually. But when everyone holds this attitude, grime and dirt start to build up and the house becomes squalid.

The concept of public goods does more than explain why household chores are so frequently left unattended to. It can explain why markets for goods with certain characteristics are often inefficient. In this chapter, I introduce the technical concept of a public good and consider how far farm animal welfare can be understood as one. I contend that farm animal welfare does not satisfy the technical conditions of a public good and is better understood as a psychological

externality. Nevertheless, I propose that, from a policy perspective, treating farm animal welfare as a public good may be useful in certain circumstances. I then assess how legislation and public good taxes can be used to provide and protect farm animal welfare.

2. Public Goods

a. Excludability and Rivalry

The two unique characteristics of a public good are best understood in relation to its counterpart of a *private good*. Many commodities encountered in our day-to-day lives are private goods: a chocolate bar is a prime example. There are two characteristics that make a chocolate bar a private good. First, it is *excludable*: it can be denied to those who do not pay for it (Snyder and Nicholson 2012: 628). Second, a chocolate bar is *rival in consumption*: when eaten, nobody else can consume that particular chocolate bar (Salvatore 2003: 611).

Not all goods are excludable. Tourists can access foreign public radio without contributing to it. Because those who have not paid for it cannot be prevented from enjoying its benefits, it is *nonexcludable* (Mansfield 1985: 491). Other goods are not rivals in consumption. When you watch a television program, you do not diminish the enjoyment experienced by other people watching that same program. Such goods are *non-rivalrous*: each subsequent person who makes use of the good imposes no additional cost upon the producer (Nicholson and Snyder 2010: 534). Most goods are neither fully rivalrous nor fully non-rivalrous, but exist on a dynamic scale between the two. For instance, local goods are often non-rivalrous only to an extent – they might have a fixed capacity, so heavy use comes at a cost to other users (Nechyba 2017: 1040-41). The arrival of a couple of extra people at a public swimming pool may do little to diminish someone's enjoyment when swimming alone, but eventually the pool becomes so packed that each additional swimmer detracts from the enjoyment of others.

There is disagreement in the literature about what constitutes a public good. Some maintain that public goods are both nonexcludable and non-rivalrous (Pindyck and Rubinfeld 2013: 690; Schotter 1997: 588; Begg et al. 2000: 281; Nicholson 2010: 531; Reiss 2013: 235).⁷⁸ Others maintain that non-rivalry is sufficient, although recognise that public goods will often be nonexcludable too (Salvatore 2003: 611; Nechyba 2017: 1040; Mansfield 1985: 490; Perloff 1999: 682). Given that

⁷⁸ The continuous nature of rivalry raises questions regarding how much non-rivalry is required for a good to count as a public good.

many classic examples of public goods, such as national defence, clean air and streetlighting, are both nonexcludable and non-rivalrous, both criteria will be assumed necessary here.

Because public goods are nonexcludable and non-rival in consumption, their costs do not reflect their full social value; as such, they can be understood as a special kind of externality. Consequently, markets for public goods, where they exist at all, are likely to undersupply public goods. The market's inability to provide a socially-optimal supply of public goods is largely a consequence of the free rider effect (Kaul 2012: 729).

b. The Free-Rider Problem

The excludability and rivalry of private goods means that the amount you are able to consume is determined by what you are willing to pay: if you are not willing to pay for a chocolate bar, you will not be able to consume it; if you are willing to pay for just one chocolate bar, you will be able to consume only one. Because, on the other hand, public goods are nonexcludable and non-rivalrous, the quantity that a person is able to consume is instead determined by total supply (Parkin 2000: 394; Mansfield 1985: 492; Begg et al. 2000: 281). When people can enjoy public goods without paying for them and their enjoyment is unlikely to lead to these goods becoming exhausted, they have little incentive to contribute to the good's supply. The problem of people using public goods without paying for them, known as *free-riding*, is thus especially salient in large societies, where consumers may believe a public good will be provided irrespective of their contributions (Salvatore 2003: 614). One person withdrawing support for a public good will barely affect its supply, allowing a free-rider to enjoy its benefits without contributing to its provision. When many people behave in this way, however, the supply of a public good can be severely reduced.

This does not mean that markets cannot exist in public goods: there are several reasons why people might choose to pay for them, the most obvious being that they have preferences for public goods. This motivation is captured by models of *pure altruism*, which assume that people have preferences only for private consumption and the total supply of a public good (Andreoni 1990: 464; 1988a: 57). Models of pure altruism are pessimistic about private contributions to public goods, however: they anticipate that, even when overall demand for public goods is high, few people will contribute because of the free-rider effect.

The pure altruism model is evidently incomplete, as many people voluntarily contribute to public goods. This is frequently seen in the literature, typically in *public good games*. In games

of this type, players are given an endowment, and must choose how much to put in a public good pot. At the end of the game, each player receives a payoff that is a function of whatever they withheld from the pot in addition to a percentage (often 40%) of what all players have placed in the pot. Assume there are four players, each with an endowment of £10. If everyone placed their entire endowment in the shared pot, each player would receive £16. Should, however, one player decide to free-ride and contribute nothing, that person will receive £22, while the others receive £12. If two players free-ride, they each receive £18, while the other two receive £8. If three free-ride, they receive £14 and the one public good contributor receives £4. If everyone free-rides, they leave with their initial endowment of £10. Clearly then, the best outcome for the group as a whole is for everyone to contribute to the shared pot. However, each player can make themselves better-off by free-riding, leading to a Nash equilibrium (an equilibrium where nobody can make themselves better off by changing their strategy) that is not Pareto-optimal (Andreoni 1988b: 293; Keser and van Winden 2000: 25-26).

Despite free-riding being the ‘optimal’ strategy in public good games, the behaviour of participants in experiments indicates many people have reasons for contributing to public goods. The hypothesis of strong free-riding, where just about everyone avoids contributing to public goods, can be rejected, as most people in public good games contribute to the public good pot. There is, however still a clear weak free-rider effect: participants in single-shot public good experiments tend to contribute between 40 and 60% of their endowments to the public good, a far cry from full free-riding, but also significantly below the optimal level of contribution (Marwell and Ames 1981; Thaler 2016: 145; Ledyard 1995: 113; Dawes and Thaler 1988: 189).

Outside of the laboratory, there are also many examples of people contributing to public goods. People donate to political parties, pick up litter and pay for fireworks displays they could watch from external vantage points for free (Thaler 2016: 144; Bergstrom et al. 1986: 25). In 2017, people in the UK self-reported as giving £10.3 billion to charity, with about 60% of people claiming to be donors. Among donors, the median monthly donation was £20 (CAF 2018). Given the extent of public engagement with public goods, it is thus clear that models of pure altruism do not capture all the reasons why people choose to contribute.

The unrealistic predictions of pure altruism models led James Andreoni (1989) to develop the theory of *impure altruism*, which, as we saw in Chapter 1 (p.32), recognises that acting altruistically – for instance by contributing to public goods – can produce a feeling called variously a warm glow, a satisfaction of conscience, an intrinsic satisfaction and a private benefit. This feeling can constitute a motivating force for pro-public good behaviours (Dawes

and Thaler 1988: 192; Clark et al. 2003: 239; Deci 1971: 105). Feelings of intrinsic satisfaction are excludable in nature – free-riders cannot enjoy the satisfaction that can accompany pro-public good behaviours – and so can offer an explanation for why people contribute to public goods.

Other writers have claimed that social norms may provide another motivation for voluntarily contributing to public goods (Holländer 1990; Cartwright and Patel 2010; Rege 2004). The desire to follow social norms, in turn motivated by the desire to receive social approval and avoid social disapproval, attaches further private benefits to public good contributions and costs to free-riding. As will be discussed in Chapter 7 (p.178), social norms could be a useful policy instrument for promoting pro-farm animal welfare behaviours.

c. *Optimal Supply*

Despite there being a range of reasons why people might want to contribute to public goods, the free-rider effect will still cause markets to undersupply these goods. Consequently, the provision of public goods is largely regarded as a government responsibility. For governments to provide public goods effectively, they need, at least in the theoretical domain, to be able to calculate public demand for these goods. The aggregate demand for a *private good* is calculated by summing the quantities demanded at various price points. For instance, when the price (represented in Figure 4 as marginal willingness to pay) of a bag of apples is £3, Nellie demands three and Oliver two: aggregate demand is therefore five bags. For *public goods*, we must instead consider people's willingness to pay for a given quantity. If Nellie is willing to pay £3 for a third public park and Oliver £1, total willingness to pay for the third park is £4. The contrast between the two processes is illustrated in Figures 4 and 5.

The different methods stem from the non-rivalry of public goods: one person's consumption of a public good does not preclude another's. If two people each demand five units of a public good, only five units need to be supplied, whereas ten units of a private good would be required. Summing individual demand would overestimate the optimal supply of public goods; willingness to pay for different levels of supply is instead the appropriate measure.

In sum, public goods are nonexcludable and non-rivalrous in nature, meaning that people can enjoy a good's benefits without paying for it, and one person's use of public goods does not limit another's. Because they can free-ride on others' contributions, consumers have little incentive to pay for public goods; despite this, we see a degree of private support for public goods, which implies that there are private benefits attached to pro-public good behaviours. Because

Figure 4 – Aggregating Demand for Private Goods
(Adapted from Kolstad 2011: 97)






Figure 5 – Aggregating Demand for Public Goods
 (Adapted from Begg et al. 2000: 268)



voluntary support for public goods is likely to be limited, markets in public goods are liable to failure and the burden of supply largely rests on the government, which can calculate the optimal level of provision by summing willingness to pay for a given level of supply.

3. Farm Animal Welfare as a Public Good

Many economists, organisations and policymakers (see Norwood and Lusk 2011, McMullen 2016a, FAWC 2011, UK Parliament 2018, and RSPCA 2018, for example) identify farm animal welfare as a public good. J.L. Lusk and F.B. Norwood (2011: 474), for instance, state that farm animal welfare is ‘non-rival (the happiness I receive from knowing a hen benefits from my consumption of cage-free eggs does not prohibit you from also being happy at the hen’s better life) and non-excludable (I cannot keep you from enjoying the fact that I eat cage-free eggs)’.⁷⁹ They therefore contend that farm animal welfare is likely to be underprovided in the market because of the free-rider problem. While some people may be satisfied in the knowledge that farm animals enjoy a particular level of welfare, they can enjoy a similar degree of satisfaction even when they do not contribute to its provision (Norwood and Lusk 2011: 315). For instance, you may derive pleasure from the existence of the growing organic industry because you believe it provides animals with higher welfare than conventional animal agriculture, yet continue to buy conventionally-produced goods because they are cheaper. Due to consumer dissociation from the suffering and death produced in animal agriculture, as discussed in Chapter 1 (p.32), people can have an asymmetric relationship with farm animal welfare, where they derive pleasure from higher welfare but overlook the low welfare aspects of animal agriculture that would otherwise cause them discomfort.

Although often desirable, some public goods can be undesirable: such things are known as *public bads*. Elements of public bads pertain to farm animal welfare (Harvey and Hubbard 2013: 108): they are both non-rival (the distress you might feel at the level of animal suffering in a society can also be felt by others) and non-excludable (people cannot be prevented from feeling distressed by animal suffering) (Norwood and Lusk 2011: 311).

One plausible candidate for a public bad in this context is animal cruelty, which, in this thesis, will be defined in terms of public attitudes. Because there likely exists significant public concern for farm animal welfare, there may exist a level of welfare or set of practices that are deemed so

⁷⁹ As noted in the previous chapter (p.72), we can understand this in terms of preference satisfaction: people have preferences for farm animals to consume certain goods, and derive happiness from the satisfaction of these preferences.

detrimental to farm animals that almost everyone opposes them and would be willing to pay to eradicate them. I will refer to these practices and these levels of welfare as ‘cruelty’ or ‘cruel practice’.

Because farm animal welfare bears similarity to public goods, markets are unlikely to supply an optimal level of it, despite there being significant public concern for farm animal welfare and willingness to pay to improve it. Consequently, some degree of provision and protection of animal welfare is likely to be in society’s interests, and there may be scope for government intervention in animal product markets to achieve this.

4. Farm Animal Welfare as a Public Good?

Although many academics and policymakers treat farm animal welfare as a public good, some thinkers have suggested that it does not in fact fulfil the necessary criteria of nonexcludability and non-rivalry. David Harvey and Carmen Hubbard (2013: 109), for instance, recognise that welfare is intrinsically tied to the production and sale of animal products. Without markets for animal products, there would be no need for farm animals, and thus there would be no farm animal welfare. The provision of farm animal welfare is therefore dependent upon the demand for animal products, which are private goods. Rather than being a public good, animal welfare could instead be characterised as a mere attribute of private goods.

Stefan Mann is also sceptical of whether farm animal welfare is best categorised as a public good, although for different reasons. Mann understands farm animal welfare as relating to an animal’s quality of life, which in turn depends upon the living conditions and husbandry practices its owner provides.⁸⁰ This relationship between farm animals and their owners is a private interaction good; it cannot be a public good because it is excludable; Mann recognises that the way one animal is treated has little bearing upon the treatment of another when he writes that ‘a farmer may pamper one cow and torture another’ (2005a: 370-71).

If animal welfare relates to the relationship between owner and animal, welfare improvements are rival and excludable, and thus private, goods. Although many welfare improvements may not be fully-rivalrous, they are certainly closer to the rival end of the spectrum. For instance, although one hen’s use of a dust bath may not preclude another from using it, there are only so many hens that can use it simultaneously. In addition, individual animals can be excluded from

⁸⁰ In this thesis, I have noted that farm animal welfare is likely to be affected by factors including production inputs (p.14). Mann’s account appears to cohere with this definition, and perhaps develops it further in its recognition that many of the constituents of farm animal welfare are provided by an animal’s owner.

making use of welfare improvements (Mann 2005b: 135-36). In short, Mann argues that improvements in farm animal welfare are both rival and excludable, and so are in fact private goods.

This raises the question of why so many have claimed that farm animal welfare is a public good. As noted earlier, J.L. Lusk and F.B. Norwood claim (2011: 474) that animal welfare is a public good because it is ‘non-rival (the happiness I receive from knowing a hen benefits from my consumption of cage-free eggs does not prohibit you from also being happy at the hen’s better life) and non-excludable (I cannot keep you from enjoying the fact that I eat cage-free eggs)’. At first glance, the claims they make seem fairly incontestable. When we consider the human benefits these writers ascribe to farm animal welfare, however, we see the flaw in their argument. Lusk and Norwood, along with others in the literature (see McInerney 2004: 4 for instance), do not see farm animal welfare as a public good in terms of how it benefits farm animals; as we have already seen, such a strategy would be doomed to failure because the benefits to farm animals are rivalrous and excludable. Similarly, they do not see farm animal welfare as a public good because of the benefits accrued by animal owners: these too are rivalrous and excludable.

Instead, Lusk and Norwood see farm animal welfare as a public good in terms of the happiness and sadness that humans draw from its *existence*. Contrast this with a classic public good such as a public radio station. Public radio is non-rival in the sense that one person listening to it does not diminish the enjoyment of anyone else listening to it, and nonexcludable because those who do not pay for it cannot be prevented from tuning in. The nonrivalry and nonexcludability conditions relate to the benefits people draw from the *consumption* of the radio station, in this case by listening to its programming, rather than from pleasure taken from its mere existence.

Through analogy, however, we see that the public is not a consumer of farm animal welfare. First, we might better understand our relationship with current levels of farm animal welfare in society, by considering Poppy, who buys a classic car that is in desperate need of refurbishment. The paint is scratched and faded, the lights do not work, and the car is beginning to rust. Poppy diligently repairs and refurbishes the car, to the point that it looks as good as new. She takes it out for a ride on a Sunday morning and the car catches the eye of numerous admiring passers-by.

By Lusk and Norwood’s reasoning, Poppy’s car would be a public good: the pleasure that one person draws from seeing the classic car in all its glory does not preclude another from feeling the same way, and people cannot be prevented from being delighted by the sight of the car. For

Poppy's car to be a public good, we must say that passers-by are consumers of the car in some way. It seems more plausible to contend that, as the sole owner of the car, Poppy is also its sole consumer. When the car's condition is improved, the benefits accrue to its consumer. The enjoyment experienced by passers-by is unrelated to the consumption of the car, but rather is a *psychological externality* stemming from Poppy's purchase of the car and the time and money she has put into refurbishing it.⁸¹

Poppy is also a chicken farmer. She used to raise her chickens in a barn, but has recently shifted to free-range production. People walking past Poppy's farm take pleasure in seeing chickens roaming the fields and living what are perceived to be good lives. This pleasure fits Lusk and Norwood's definition of a public good: one person's pleasure does not preclude another's, and people cannot be prevented from taking pleasure from the sight of Poppy's poultry roaming the fields. In fact, the mere knowledge that Poppy's chickens are living good lives may be enough to spark pleasure. Just as with her classic car, however, Poppy is the sole owner of her chickens, and is thus the sole consumer. When Poppy shifts production from barn to free-range, many would say that the chickens' welfare improves. Any benefits of this improvement accrue solely to Poppy, who might now own healthier and possibly happier chickens, and the chickens themselves. The enjoyment experienced by passers-by is once more unrelated to the consumption of the chickens, but is a psychological externality created by Poppy's investment in free-range production.

To labour the point, because farm animal welfare concerns the relationship between an owner and an animal, welfare improvements can only be consumed by the owner and the animal; consequently, the benefits of improved farm animal welfare are 'consumed' by the animal's owner, and perhaps the animal itself. Where other people derive enjoyment from the presence of higher welfare, this is not as consumers, but as mere passers-by. Their benefits are thus best understood as psychological externalities.

But what of our relationship with farm animal welfare that our purchases have contributed to? If passers-by stop off at Poppy's farm to buy her produce and financially contribute to her provision of higher welfare, are they not in some sense consumers of this welfare too? Another analogy sheds light on this matter. In Christmas 2019, Oxfam launched a series of 'life-changing

⁸¹ Some have recognised that the consumption of some goods appears to be visual in nature (Urry 2000: 148-9; Zuev 2015: 79). It seems implausible to suggest that passers-by are visually consuming Poppy's car, however: as J.E. Schroeder notes (2004: 230), 'visual consumption begins with images', such as adverts, photographs and paintings. Perhaps passers-by are consuming the *image* of Poppy's car, but it remains implausible to claim that they are consuming the car itself.

gift cards'. For £25, you could buy a Christmas card for a friend or relative, while also providing a pig to a women-led co-operative farm in northern Rwanda. The pigs provide manure and piglets, which can be sold to fund essentials such as food and education, and empower these women to become more self-sufficient (Oxfam 2020).

When Qasir buys an Oxfam poverty-busting pig gift card, however, we want to say that he becomes a consumer of the gift card itself, not the pig. The pig is owned and consumed by the owners of the co-operative farm in northern Rwanda. Qasir may derive a feeling of warm glow from his purchase, and so may experience benefits beyond those immediately provided by the gift card, but he is clearly not a consumer of the pig.

It is useful at this point to introduce a distinction used in the marketing literature. Economics tends to bundle consumers, buyers and shoppers together as consumers, but, in this instance, recognising the distinction between the three is enlightening. Consumers are a product's end users. If you eat a steak, you are the consumer of that steak. When you have your hair cut, you are the consumer of that service. When you visit an art exhibition, you are a consumer of the exhibition. Buyers, on the other hand, are those who front up the money for a purchase. Shoppers are those people directly involved in the purchase decision, who choose which goods to buy, irrespective of whether they are buying these goods with their own money or whether they are the end users of them (Shankar et al. 2011: S29).

When buying a poverty-busting pig gift card, Qasir is the shopper, buyer and consumer of the card. If he gifts it to a friend, the friend is the consumer of the gift card, while Qasir remains the shopper and the buyer. Neither Qasir nor his friend are consumers of the pig itself, however. Qasir buys the pig, an Oxfam agent presumably shops for the pig, and it is consumed by the co-operative farmers of northern Rwanda.

So it is the case for purchases of higher-welfare products. If Qasir pays a premium for free-range chicken, he plays the role of consumer, buyer and shopper for the chicken meat. For access to free range, Qasir is merely the buyer: Poppy shops for the land upon which the chickens roam, and either Poppy or Poppy and her poultry are consumers of this land and any associated welfare benefits. Qasir does not roam the fields that his purchase has contributed to, so he is clearly not a consumer of the access to free range that his money has bought: the access to free range can thus be understood as non-commodity output of agriculture that is not consumed by Qasir (Ryland 2014: 854). Any pleasure that Qasir feels as the result of his purchase is either related to his consumption of the chicken, or psychological externalities relating to his (charitable) contribution to the birds' welfare.

We are not done yet, though. Like many people, Qasir associates farm animal welfare with quality attributes including safety, healthfulness and taste (Brook Lyndhurst 2012: 21; Matthews 1996: 41-42; Jago et al. 2000: 163-64; Harper and Henson 2001: 21-22; Verbeke et al. 2010: 285; Heng et al. 2013: 425; RSPCA 2006: 23). When he buys free-range chicken, does he not gain in some way from the higher-welfare aspects of his purchase? If so, could we say that he directly benefits from the provision of higher welfare, and could be said to be a consumer of this welfare? The answer is no. We have seen that animal welfare relates to the relationship between the animal owner and the animal. When the animal has been slaughtered, processed and packaged, this relationship is severed and there is no welfare to speak of in the final product. The perceived benefits (taste etc.) associated with an animal's welfare when it was alive are product attributes, and are directly enjoyed by the consumer, but there is no welfare to speak of, to consume, or to directly benefit from. Higher welfare benefits the consumer indirectly, by supporting those immediate quality attributes such as safety, healthfulness and taste that consumers directly benefit from.

In sum, farm animal welfare is not a public good because the public does not consume farm animal welfare. Farm animal welfare concerns the relationship between farm animals and their owners, and thus is rival and excludable in nature. Any pleasure (or pain) that people feel as a result of the prevailing level of farm animal welfare in society is better understood as a psychological externality stemming from the level of welfare that an owner provides their animal. Even when people pay for higher-welfare products, they are not consumers of this higher welfare: the premiums they pay for higher-welfare goods are more plausibly viewed as charitable donations, to the benefit of farm animals and their owners. And even when people derive private benefits from higher-welfare goods, we do not want to say that they are consumers of higher welfare. If welfare concerns the relationship between an owner and their animal, this relationship is dissolved when the animal is slaughtered, and thus there is no welfare directly contained in the animal product. End users might benefit from attributes provided by the welfare that an animal enjoyed in life, but they are not consumers of the welfare itself. If the public is not a consumer of farm animal welfare, we cannot understand it as a public good.

Despite farm animal welfare not being best-understood as a public good in a technical sense, it may still be prudent to continuing treating it as if it were one, at least from a policy perspective. Many of those in the policy debate, including former DEFRA secretary Michael Gove, identify farm animal welfare as a public good (McInerney 2004: 3; FAWC 2008: 2; FAWC 2009: 44; RSPCA 2018: 2; Tasker 2018). If this thesis is to resonate with the policy debate, it must surely

engage with the issue on the same terms. Additionally, the need for government intervention in public good issues is widely accepted and so treating animal welfare as a public good might make it easier to justify government action in an area that has significant public support. Last, due to the similarities between public goods and psychological externalities, some public good policies could also be effective tools for responding to psychological externalities.

Even if it is deemed inappropriate to use public good policies to directly provide farm animal welfare, these policies could still be used to indirectly support animal welfare. First, there may be public goods which contribute to farm animal welfare. One example that warrants further investigation is education. At this point, there have been few studies concerning the link between education and purchases of higher-welfare goods, so firm conclusions cannot be drawn. If a causal connection between education and support for animal welfare can be established, however, education policy may be an indirect means of satisfying public preferences for farm animal welfare. There may exist other public goods whose support will benefit farm animal welfare.⁸² Further research into the demographics of those buying higher-welfare animal products may indicate what these public goods are.

Second, supporting farm animal welfare may benefit the supply of other public goods. Animal cruelty is incompatible with a society's values, and these values may be a public good in their own right. The benefits of living in a society that generally endorses ethical behaviours are numerous: physical, emotional and financial support for those who need it, protection of rights, action against cruelty, fair treatment for all, etc. These benefits can be seen as both non-excludable and non-rivalrous, thus constituting a public good. Applying these ethical codes consistently throughout society may also require the protection of farm animal welfare.

Third, there may be public goods directly related to farm animal welfare that can be supported with public good policies. Harvey and Hubbard (2013: 113) argue that, rather than animal welfare itself, it is the regulation of animal welfare that constitutes a public good. One person's 'consumption' of animal welfare regulations does not limit another's, and people cannot be excluded from consuming animal welfare regulation, which applies to a whole society and not a select group. Consequently, there is scope for legislation to play a role in any public good response to farm animal welfare.

⁸² There may also exist public goods whose support would be detrimental to farm animal welfare. As discussed in the previous chapter (p.80), support for environmental public goods such as clean air could ultimately harm farm animal welfare.

In sum, although many writers consider farm animal welfare to be a public good, the human pleasures and pains associated with it are derived from the existence of welfare rather than its consumption and so are more accurately categorised as psychological externalities. Nevertheless, the concept of public goods remains useful when discussing policy responses in this area: policymakers continue to treat farm animal welfare as a public good; public good policies may be appropriate tools for responding to psychological externalities; there may be public goods that benefit farm animal welfare, and supporting animal welfare can help develop other public goods; and aspects of farm animal welfare are public goods. For these reasons, I will spend the remainder of this chapter examining how public good policies can influence the provision of farm animal welfare: doing so coheres with the current policy conversation and can produce proposals that will have a meaningful impact on farm animal welfare. Additionally, if it is agreed that animal welfare is not a public good, this chapter's analysis can be adapted to target public goods that indirectly benefit animal welfare. I begin with legislative interventions.

5. Legislation

An unfunded mandate is a type of legislation that requires producers to provide public goods at their own expense. Given that unfunded mandates require everyone in a society to obey the new legislation, they are appropriate public good policies because they guarantee a good's provision.

Unfunded mandates can be used to both require the use of certain practices and prohibit others, and so are appropriate for dealing with both public goods and public bads. The policy is limited, however, by the extent to which producers can extract payment from their customers. For this reason, unfunded mandates may be most effective when applied to public goods that are bundled with private goods; the policy has been used to target air pollution, for instance, by requiring producers to install pollution-limiting devices on vehicles. The policy may thus be appropriate for addressing issues in farm animal welfare, as animal welfare, which we are treating as a public good, is bundled with animal products, which are private goods.

One way that unfunded mandates can address public good problems caused by farm animal welfare is by imposing minimum welfare standards that prohibit cruel practice (Doonan et al. 2009: 233). There are several examples of unfunded mandates being used to protect animal welfare, particularly relating to the use of caged egg systems. Switzerland's 1981 Animal Welfare Act effectively banned cage systems (Häne 2000: 22), while the European Union's Council Directive 1999/74/EC of July 1999 introduced a ban on the use of conventional cages in egg production that came into effect in 2012 (Druce and Lymbery 2006: 130). Both Austria and

Germany have incoming unfunded mandates banning the use of enriched cages in egg production, to be enacted in 2020 and 2025-28 respectively (Stevenson 2018: 18), and, as we saw earlier, California eliminated the use of battery cages in egg production with the passage of Proposition 2 by popular vote in 2008. Beyond egg production, sow stalls were banned in the UK through an unfunded mandate in 1999 (Levitt 2018b), and other unfunded mandates were established in *The Welfare of Farmed Animals (England) Regulations 2007*.

Any legislation that limits consumer choice or leads to consumer price increases is likely to face at least some opposition, but consumer research suggests that anti-cruelty legislation may attract broad support. Focus groups indicate that the existence of animal welfare legislation helps many consumers to feel comfortable about not considering welfare issues when buying animal products; they instead adopt an attitude of ‘if it’s on the shelf it must be fine’ (Brook Lyndhurst 2010: 33; Harper and Henson 2001: 21). It thus appears that many consumers believe that responsibility for eradicating cruel practice at least partially rests with the government, and so some degree of legislation may be appropriate. Legislation can even be a popular response: 62% of British respondents to the 2005 Eurobarometer survey (2005: 64) claimed that animal welfare was not treated with enough importance in food and agricultural policy, indicating that there exists appetite for further government action.

Although political viability may not present a major threat to unfunded mandates, several factors could undermine the policy’s effectiveness. First, legislation must be carefully worded, otherwise the intended welfare improvements may not follow. This was seen in California, where the wording of Proposition 2, which was intended to eliminate the use of cages in egg production, was suspect: it stated that farm animals should not be confined in a manner that prevents them from lying down, standing up, fully extending their limbs and turning around freely. This prompted at least one producer to install larger cages (Lusk 2011: 563), while others simply removed the connecting walls between cages, giving hens more space but maintaining existing stocking densities (Reese 2018: 17- 18).

Second, there must be robust penalties for breaches of unfunded mandates, otherwise firms will have little incentive to conform. This factor may be behind the initially limited effectiveness of the EU battery cage ban in 2012, as the size of penalties and the enforcement of the directive was left to member states (Council Directive 1999/74/EC art. 13). Fifteen states (including the United Kingdom) were not rigorous in enforcing the law when it first came into effect: consequently, an estimated 14% of EU hens were held in non-compliant cages at the start of

2012, despite producers having had over twelve years to prepare for the ban's implementation (Blandford and Harvey 2014: 37-38; BBC Newsround 2012; BBC 2012).

Third, unfunded mandates will, in many cases, be ineffective if firms are unable to extract payment, either directly or indirectly, for their efforts in increasing the provision of public goods. Unfunded mandates can be expected to increase producers' costs – if they reduced costs then market forces would likely have already pushed producers towards the desired production methods and there would be no need for legislation. If producers cannot collect remuneration for these higher costs, it may become untenable for them to remain in the market.⁸³ In many cases, producers can simply raise prices to accommodate the greater costs of higher-welfare production. This may be difficult, however, if governments do not apply unfunded mandates to imports.

When unfunded mandates are not applied to imports, they force domestic producers to adopt more costly production practices but do not remove lower-welfare imports from the market. These cheaper, lower-welfare imports can undercut domestic producers. Consequently, this use of unfunded mandates can be expected to do little to improve animal welfare, but will instead drive socially-unacceptable production practices abroad while undermining the competitiveness of the domestic industry (Sumner et al 2011: 248; Mann 2005b: 143; Doonan et al. 2009: 233). This was an effect of the UK's 1999 ban on the use of gestation crates in pig farming, which could not be applied to EU imports due to EU competition law and the nature of the customs union. Despite surveys indicating that consumers would support the new, higher domestic welfare standards, the market was flooded with cheaper meat from European countries with lower standards: imports of Danish pig meat increased 50% between 1997 and 2007 and German imports increased 400% in the same period. The impact on the domestic industry was catastrophic: the number of British pig farmers nearly halved, from 10,000 in 1999 to 6,000 by 2009, and the country's self-sufficiency in pig meat fell from over 80% at the end of the decade to under 50% in 2018 (AHDB 2018). A similar story played out in the Swedish pig industry, where a ban on gestation crates was introduced in 1994 to complement bans on farrowing cages and tail docking. The country joined the EU in 1995 and experienced a surge in

⁸³ This is of particular concern for firms competing in export markets. If international markets have little appetite for higher welfare standards, the price increases associated with an unfunded mandate might harm a firm's international competitiveness (Blandford and Fulponi 1999: 421).

cheap imports from Denmark and Germany. Whereas Sweden had once been almost entirely self-sufficient in pig meat, by 2014 it was only 68% self-sufficient (Levitt 2018b; O'Dwyer 2014).⁸⁴

Policymakers have at least three ways of responding to the threat of lower-welfare imports. First, they could extend an unfunded mandate to apply to imports. This path was taken by the Californian state government, which passed bill AB 1437 to require imported animal products to meet the standards demanded by Proposition 2 (Mullally and Lusk 2018: 652). One issue with this policy is that exporting producers will need to be monitored and certified by trustworthy organisations to ensure that standards are being met. This could impose significant costs upon these producers, making trade less profitable and thus less desirable. An exception may exist where there is unmet demand for higher-welfare goods in the exporter's domestic market, in which case welfare certification may provide additional benefits by allowing producers to specialise within their own domestic market.

Another issue with this response is that it may not be compatible with World Trade Organisation (WTO) regulations. The WTO does not permit technical barriers to trade, i.e. standards and regulations that are deemed discriminatory against other countries or that give domestic producers an unfair advantage (WTO 2014: 5). The US-Mexico Tuna-dolphin I dispute of 1991 suggests that an import ban on lower-welfare production could be regarded as a form of trade discrimination: the WTO ruled that import bans are not permitted on methods of production, but must instead be applied to broad categories of goods. This means that an import ban on eggs may be permissible, but a ban on caged eggs would not be (Bowles et al. 2017: 1; Stevenson 2015; Blandford and Fulponi 1999: 419). On the other hand, the organisation's General Agreement on Tariffs and Trade (GATT) delineates circumstances under which violations of WTO measures are exempted. One such circumstance allows measures to be taken to 'protect public morals' (WTO 1986: 37). Although this argument has not yet been presented to the WTO, it could be claimed that the protection of farm animal welfare is a matter of public morals, and so a ban on lower-welfare imports would be justified. Peter Stevenson attempts to make this argument, pointing out that the 2013-14 EU ban on seal imports was upheld by the WTO, in a

⁸⁴ Interestingly, the limited empirical evidence available in this area does not unanimously point towards welfare improvements being deleterious to domestic production. Switzerland's Animal Welfare Act phased out the use of caged systems of egg production between 1981 and 1992, but did not change its import policy (Häne et al. 2000: 21-22). One might have expected cheap caged eggs to flow into the country and crowd out domestic producers, who were now required to use more costly production methods. However, the reverse happened, and the domestic market share increased from 62% in 1992 to 74% in 2006. Mann (2005b: 143) attributes this to the psychological effects of the ban, which indicated to consumers that eggs produced using caged systems were ethically unacceptable. These psychological effects were surely assisted by a mandatory labelling scheme that required imported caged eggs to carry the label 'Produced in battery cages, which are not permitted in Switzerland' (FAWC 2006: 10).

case that established animal welfare as an issue relevant to public morals (Stevenson 2018: 11). Whether this case sets a precedent that would justify a ban on animal products produced using lower-welfare methods remains to be seen; for now, though, policymakers should not rely upon an import ban as a means of protecting domestic production and farm animal welfare from lower-welfare imports.

A second response is to impose tariffs upon lower-welfare imports in order to safeguard the competitiveness of the domestic industry (Stevenson 2018: 4). This increases the costs of lower-welfare production to the point that domestic producers cannot be priced out of the market by cheap lower-welfare imports. This will not be possible, however, if imports are protected under the terms of a free-trade agreement, highlighting the need for governments to defend domestic standards in trade negotiations.

Additionally, this use of tariffs risks falling afoul of WTO regulations. The WTO operates a 'most favoured nation' principle to prevent trade discrimination. This requires tariffs on goods to be applied equally to all countries, in the absence of a free trade agreement or customs union (WTO 2019). We have already seen that the WTO does not permit discrimination of goods based on production method, meaning that a low welfare tariff would likely be interpreted as discriminatory.

A third way of protecting domestic producers from cheaper, lower-welfare imports may be to introduce clearer product labelling to inform consumers which goods have been produced using lower-welfare practices. This policy may have helped preserve the domestic Swiss egg industry after battery cage production in the country was prohibited in 1992 (FAWC 2006: 10). As I discuss further in Chapter 6, consumers are given relatively little information about animal welfare at points of purchase, making it difficult to discriminate between products according to welfare attributes. Consumers would be able to avoid lower-welfare imports if they could identify these goods at points of purchase. This in turn could encourage international producers to improve welfare standards in order to remain competitive in the British market.

Once more, however, WTO regulations may scupper this policy. Article 2.2 of the WTO's Agreement on Technical Barriers to Trade permits the introduction of technical regulations such as labelling schemes in pursuit of a 'legitimate objective'. Although animal health is considered to be a legitimate objective, the WTO has not ruled that animal welfare is one (WTO 1995: 118; FAWC 2006: 7; McInerney 2004: 62). Until the WTO definitively rules on this issue, it is not clear that there is scope to enact animal welfare policies that discriminate between products on the basis of production method.

At this point, it appears that the UK's relatively stringent animal welfare legislation threatens to undermine the competitiveness of its domestic animal product industries, especially in the context of Brexit. If it is not possible to protect domestic producers from lower-welfare imports, British producers would likely need to lower their own welfare standards to remain competitive, which would require the relaxing of the country's animal welfare legislation (CIWF 2016: 3; UK Parliament. House of Lords European Union Committee 2017: 15). The stringency of WTO regulations does not, however, mean that the UK is powerless to protect its domestic industries from lower-welfare imports. One solution would be to enact high tariffs on all animal products to render the import of lower-welfare goods nonviable, before negotiating trade deals which grant higher-welfare animal products favourable access to the UK market.⁸⁵

This strategy could even strengthen UK animal product industries in the wake of Brexit. In certain areas, UK welfare standards are more stringent than the EU's. For instance, sow stalls are banned in the UK (and Sweden), but can be used in other EU member states from the weaning of the previous litter until the end of the first four weeks of pregnancy (CIWF 2019). EU competition law prevents the protection of domestic industries against other member states, however, which prohibits the UK from banning EU imports that fail to meet its domestic welfare standards (Stevenson 2018: 17). Consequently, cheaper EU imports undercut the British pig farming industry and lessen the incentive to enact further farm animal welfare legislation in other areas. Well-handled post-Brexit trade deals could enable the UK to become a world leader in farm animal welfare while maintaining a thriving domestic industry. Without import restrictions on lower-welfare goods, however, domestic producers are likely to call for existing welfare standards to be weakened in order to remain competitive.

To be effective, animal welfare regulations must therefore be carefully worded, robustly supported by legal sanctions for noncompliant producers, and (ideally) applied to imported products. When these concerns are met, legislation promises to be a powerful tool that can prevent cruel practices and buttress a society's ethical codes.

Even when implemented effectively, however, legislation targeting low-welfare production can still produce negative effects. Legislation that limits market freedom can harm people's utility by reducing the range of products they can buy. Amartya Sen (1988: 290-92) argues that the ability to make choices has intrinsic value and can be understood as central to personal

⁸⁵ The EU appears to be adopting a similar strategy, with a recent free trade deal with Brazil, Argentina, Paraguay and Uruguay allowing eggs to be imported duty-free only if they meet the EU's animal welfare standards (Busby 2019).

development; consequently, reductions in market freedom may be viewed as *prima facie* undesirable. J.C. Harsanyi (1977a: 652-54) makes a similar point from a preference utilitarian perspective, recognising that having free personal choice can be utility-enhancing.

Additionally, as higher-welfare production is generally more expensive than lower-welfare production, unfunded mandates can be expected to disrupt consumption patterns. This will disproportionately affect those on lower incomes. For one, people in higher income brackets are more likely to already buy higher-welfare goods, meaning that reducing the availability of lower-welfare options will be less disruptive to their consumption habits, but more disruptive to those on lower incomes, who are more likely to buy lower-welfare animal products (IGD 2007: 31; Brook Lyndhurst 2010: 42). Additionally, wealthier consumers will be better-able to absorb higher prices, while lower-income consumers may struggle to maintain their existing consumption patterns.

Consequently, policymakers ought to be disciplined in their use of unfunded mandates, and restrict this policy to targeting cruel practice. Given that farm animal welfare is a matter of concern to almost everyone in the UK, and given that animal cruelty faces almost universal opposition, the benefits of using unfunded mandates to tackle cruel practice will likely outweigh the costs of reduced choice and higher prices. It is when unfunded mandates are used to ban practices at higher levels of welfare that the benefits may be outweighed by limitations to market freedom, which disproportionately affect those on lower incomes.

6. Taxes

One of the criticisms I made of unfunded mandates was that they can lead to higher prices when they require producers to make costly capital investments. This issue can be mitigated by public good taxes. Public good taxes are one form of funded provision that eradicates the free-rider problem by making participation in the supply of a public good involuntary: people no longer choose whether to contribute to a good's provision, but are compelled to by law (Salvatore 2003: 614).

It is useful to highlight the differences between public good and externality taxes. The two policies have different goals and are levied on different groups: externality taxes attempt to make prices reflect the full social value of a good; they are thus levied on participants in the externality-producing market. Public good taxes, on the other hand, seek to ensure the provision of a good that benefits almost everyone in a society, and so are levied on the whole

society. Where animal welfare is a matter of concern for almost everyone in a society, i.e. at levels of cruel practice, policymakers may be justified in levying taxes on the entire population. At higher levels of welfare, however, there may be significant groups of people who are uninterested in the further provision of animal welfare, meaning that a mandate for direct government action, which would impose costs upon the entire population, is weaker. Other forms of intervention may still be desirable though, particularly if they make markets more reflective of society's preferences for animal welfare.

There appear to be two plausible ways of using a farm animal welfare public good tax. One use is to support farm animal welfare research with the aims of developing new technologies and informing government recommendations about how animals should be kept. Alternatively, tax money could compensate producers for making welfare improvements, in what we may choose to call a *funded mandate* (Blandford and Fulponi 1999: 417-18).

Both uses of a public good tax face informational issues in determining the optimal level of supply. As discussed earlier, the optimal supply of public goods is determined by summing people's willingness to pay for a given level of provision, which thus requires knowledge of each person's demand curve. The information issues that affect a government's attempts to calculate the optimal supply of psychological externalities – discussed in the previous chapter (p.81) – also pertain here. In addition, there are incentives to provide misleading information about demand for public goods. If, as proposed in the Lindahl model, a public good tax charges people according to their personal demand, they have an incentive to underreport. This allows them to pay a lower level of tax whilst free-riding off the contributions of others (Nicholson and Snyder 2010: 540). If, for instance, a public good tax was introduced to provide laying hens with access to free range, some people might choose to state a lower willingness to pay in the belief that others will reply honestly, allowing them to enjoy living in a society where laying hens have access to free range without having to pay as much for it.

Alternatively, if tax payments are independent of an individual's demand, people have an incentive to exaggerate personal benefits, which would encourage governments to provide more of a public good while placing the greater part of the burden upon the rest of the population. As Begg et al. (2000: 282) put it: '[w]e are all for safer streets if we do not have to contribute to the cost'. Returning to our public good tax to provide laying hens with access to free range, animal lovers may choose to exaggerate the value of their personal benefit: doing so causes the government to overestimate the public good's value and supply more of it, while the cost burden of this additional supply would be spread over the entire population. These informational

difficulties mean that public good taxes are unlikely to be suitable policy tools for ensuring a socially-optimal level of supply for public goods.

As argued earlier, however, the aim of public good policy with respect to farm animal welfare should not be to achieve a socially-optimal level of supply, but rather to eradicate cruel practice. At the level of cruel practice, policymakers can be confident that the costs of prohibiting animal cruelty will be outweighed by the benefits, namely a reduction in negative psychological externalities and an increase in positive ones. While externality policies, which seek to provide a socially-optimal level of animal welfare, are more likely to be hindered by informational issues, the aims of public good policy in farm animal welfare are less ambitious and so largely avoid these informational issues.

Despite circumventing these informational difficulties, there are two areas where animal welfare public good taxes may be vulnerable to criticism. First, opposition to this policy can be expected from many in society, as new costs framed as taxes are often unpopular with the public (Carattini et al. 2017: 1). Unfunded mandates are likely to be more acceptable to the public than funded ones as they are not presented in the form of a direct tax, though they too are likely to lead to higher consumer expenditure, albeit through price rises. Additionally, public good taxes may be opposed on ethical grounds, particularly by animal abolitionists such as Gary Francione and Anna Charlton (2015), who oppose all forms of animal ownership: these taxes effectively force animal abolitionists to make financial contributions to an institution they regard as morally repugnant.⁸⁶

Second, producers who already use higher-welfare systems may oppose a public good tax. The tax provides no benefit to them, but creates new competitors in markets for higher-welfare goods, which increases competition and may force some out of business. Given that one aim of an animal welfare public good tax is to improve farm animal welfare, using it in a way that could harm existing higher-welfare producers is inappropriate. Even producers who use conventional or low-welfare production methods, and thus stand to be the primary beneficiaries of a funded mandate, might have cause for concern. Although a funded mandate would cover the capital costs of new housing, many higher-welfare production systems have higher running costs⁸⁷

⁸⁶ Any policy that incentivises certain forms of animal product production, be it a tax, subsidy or new legislation, will also be subject to this criticism. A discussion of the viability of the abolitionists' argument lies beyond the scope of this thesis, which is concerned with making society more reflective of public concern for farm animal welfare. Implicit in this is the assumption that consumers are free to make their own minds up about the acceptability of consuming animal products, which runs counter to the animal abolitionists' claims.

⁸⁷ Exceptions to this rule do exist, for instance in pig farming, where some forms of group housing are cheaper than lower-welfare sow stalls (Stevenson 2011: 5-6).

than lower-welfare alternatives or cause overall productivity decreases (Stevenson 2011: 5-7). While the expected cost increase of moving from low-welfare to higher-welfare production may amount to just a few pence per unit, this could still push some producers out of business if it leads to higher prices and reduced consumer demand (McInerney 1998). In such cases, consumers would be effectively hit twice by the funded mandate: once by the cost of the tax itself, and again by the subsequent increase in product prices.

7. Summary

I have considered a range of ways that farm animal welfare is affected by, and can in turn affect, public goods. Although the public's benefit from farm animal welfare is better understood as a psychological externality and not as a public good in the technical sense, it may be worth continuing to treat it as one from a policy perspective, at least at the level of animal cruelty. At its lowest levels, i.e. cruel practice, direct government intervention to secure the provision and protection of farm animal welfare can be justified as being in almost everyone's interests. Government action might take the form of legislating in defence of animal welfare; most people in society would prefer to have laws that enshrine ethical proscriptions against cruel practice. These laws could be combined with public good taxes in order to either develop a stronger understanding of farm animal welfare through scientific research or lessen the financial burden upon producers who are required to change their production methods.

Care must be taken, however, to ensure that public good policies are effective in achieving their aims. Poorly-worded laws, or ones that are not applied to imports, may have little impact on overall animal suffering, and are thus unlikely to ease public discontent. Tax-funded welfare improvements must be considered carefully in order to avoid harming existing higher-welfare producers.

Despite the UK having some of the most robust anti-cruelty measures in the world, there appears to be scope for further development. In some areas, new laws may be appropriate. This was illustrated by the public outcry that accompanied the release of footage taken by animal campaign group Animal Equality, who investigated three farms owned by Moy Park, a British chicken producer that supplies 30% of the UK market, and sells to retailers including Tesco, McDonald's, Co-op, Ocado and Sainsbury's. The footage showed chickens suffering from skin irritation, breathing difficulties and leg injuries that left them unable to stand, lame birds collapsing under their own weight and carcasses mouldering among the flock. Importantly, the public outcry did not relate to legal violations on the part of Moy Park, who appears to have

broken no laws (Animal Equality 2019; Grant 2019; BBC 2019). This indicates that public appetite exists for further government action, perhaps relating to the use of fast-growing species that are susceptible to health problems. Additionally, public awareness of production practices and interpretations of cruel practice can evolve, which may require new policy interventions, while the changing political landscape, namely the challenges posed by Brexit, means that care must be taken to ensure that existing anti-cruelty measures are not undermined.

Chapter 6

Information Asymmetries

In this Chapter

- An introduction to information asymmetries.
- A discussion of how labelling can overcome information asymmetries.
- An assessment of welfare labelling in the UK.
- A defence of the claim that logo welfare labels could be crowding out higher-welfare goods through adverse selection.
- A proposal for a more effective welfare label.

1. Introduction

In May 2019, the value of star fund manager Neil Woodford's Woodford Equity Income Fund plummeted £600m after a series of investments failed to perform, prompting customers to withdraw their money. Woodford was forced to suspend further withdrawals, effectively trapping customers' money in an undesirable fund. The investment group Hargreaves Lansdown had confronted Woodford with concerns about the state of the fund in November 2017, yet continued to recommend it to new investors (Makortoff 2019a, 2019b; Farrell 2019).

In short, investors knew less about the state of the fund than Hargreaves Lansdown and Woodford. This asymmetric information led people to make investments they would not have otherwise made, which both trapped and devalued their money. Asymmetric information thus caused many investors to act in ways that were not reflective of their preferences.

Information asymmetries create suboptimal outcomes in a range of markets, including those for animal products. I begin this chapter by detailing how information asymmetries in animal product markets produce suboptimal outcomes, before considering how information provision in the form of welfare labelling can overcome this market failure. I analyse the RSPCA Assured label and build upon this analysis to propose a more effective welfare label.

2. Farm Animal Welfare as a Product Attribute

Information asymmetries exist when producers and consumers possess different levels of information about a good or a market. Markets afflicted with information asymmetries will inefficiently allocate resources because consumers (and producers) are liable to unintentionally behave in ways that do not reflect their interests (Begg et al. 2000: 48). Asymmetries tend to be of particular disadvantage to consumers: producers are likely to at least possess complete information about their own goods (Norwood and Lusk 2011: 328). Information asymmetries can be attributed to the existence of different types of product attributes.

There are three main classes of product attribute. *Search attributes* are characteristics that can be identified in a good before purchase, such as shape, size and colour. If you have a preference for apples that are red or shirts that are medium-sized, a cursory examination of a product will likely tell you whether it will satisfy this preference. Given that information about search attributes is available at points of purchase, most consumers will possess full information about these qualities and will make choices that satisfy their preferences for search attributes.⁸⁸

A second class of product attribute is the *experience attribute*, which delineates qualities, such as taste and comfort, that can only be identified after a product's use or consumption. Determining your favourite brand of, say, canned tuna will require the purchase of several brands to compare taste (Nelson 1970: 311-12), but once you have found your favourite, you are unlikely to intentionally buy a brand that you believe to be inferior, price considerations etc. notwithstanding. In cases such as these, experience attributes are unlikely to contribute to market inefficiencies in the long run. There are circumstances, however, where experience attributes may cause inefficient outcomes. If a good is expensive, for instance, it may not be feasible for a prospective consumer to sample all available options, although consumer and industry reviews may enable them to make a more informed choice about a product's experience attributes.

The last class of product attribute is the *credence attribute*, which is a quality, such as nutritional profile and environmental impact, that is not easily identifiable either before purchase or after consumption (Darby and Karni 1973: 68-69; Wirth et al. 2011: 49). Credence attributes are likely to be significant sources of market inefficiency because consumers are dependent upon external assistance if they are to possess information about these attributes.

⁸⁸ One exception is price: consumers may not possess information about how a good is priced elsewhere, and so may inadvertently pay more for a good than they need to. The rise of online shopping and price comparison websites has, however, mitigated this to an extent.

Credence attributes are the most relevant type of product characteristic for this thesis, as they abound in food markets. Qualities such as nutritional profile, food safety, environmental impact, origin, production methods and the ingredients used to process goods could all be regarded as credence attributes (Caswell and Mojdzuska 1996: 1250-51; Moser et al. 2011: 122-28; Blythman 2015: 79). Farm animal welfare is also a credence attribute (Fearing and Matheny 2007: 165; McInerney 2004: 12; FAWC 2006: 5): examining an animal product before purchase will rarely yield an insight into how the animal lived, and, in many cases, people cannot taste the difference between lower- and higher-welfare goods (Haspel 2010; López-Alt 2010; Dransfield et al. 2005: 68; Jonsäll et al. 2002: 77). Consequently, consumers may struggle to identify a product's animal welfare or other credence attributes at points of purchase and may make purchases that violate their preferences for these qualities.

Informational issues can also negatively affect producers, because, when consumers lack information about a product's credence attributes, their behaviour does not reliably signal to producers the extent of demand for these attributes, causing gaps in the market to be left unfilled (Fearing and Matheny 2007: 164). Although market research can help producers ascertain what consumers want, it may not be as reliable as evidence derived from consumer behaviour, especially in matters of farm animal welfare. As discussed in Chapter 5 (p.120), research conducted in 1999 suggested that British consumers would support the domestic pork industry if prices increased as a result of a ban on sow stalls. The research was unduly optimistic: after the ban was introduced, consumers turned to cheaper German and Danish imports, contributing to a decline in the number of British pig farmers, from 10,000 in 1999 to 6,000 in 2009. The country's self-sufficiency in pork similarly declined, from over 80% at the end of the decade to under 50% in 2018 (Levitt 2018b; AHDB 2018).⁸⁹ Care must therefore be taken when using market research to reveal consumers' preferences for credence attributes.

In short, inefficiencies in animal product markets can be caused by consumers' inability to identify a product's animal welfare attributes, and consequent inability to signal preferences for these attributes to producers. Governments and producers can respond to this form of market failure by providing external cues such as product labels to assist consumer decision-making (FAWC 2006: 5; IGD 2007: 36).

⁸⁹ We also saw in Chapter 5 that this decline may have been the result of inadequate product labelling, which left consumers unable to discern which products had been made according to standards that were prohibited in the UK. We might therefore conclude that consumer research may be an effective way to learn about consumer attitudes, but not necessarily barriers to purchase.

3. Welfare Labelling

Labels provide information about product attributes that may not otherwise be readily-accessible. At their most effective, labels reduce uncertainty and enable consumers to reliably form expectations about a good's credence characteristics, while also assuring them that these expectations are justified (Napolitano et al. 2010: 537; Fearing and Matheny 2007: 165; Tonsor and Wolf 2011: 431; Blandford and Harvey 2014: 36). They thus transform experience and credence attributes into search attributes (Caswell and Mojduszka 1996: 1251). From a behavioural perspective, labels can be a particularly effective way of influencing behaviour because they operate at the point of purchase, which reduces *search costs* – the time, effort and money expended examining different options – for consumers with preferences for a product's credence attributes (Tonsor and Wolf 2011: 431).

Labelling programmes have desirable attributes for consumers. First, labels preserve a high degree of market freedom. While they help consumers with preferences for credence attributes to identify the products that satisfy these preferences, they do little to disrupt the consumption patterns of those uninterested in the labelled attribute (Harper and Henson 1999: 14; Caswell and Mojduszka 1996: 1250). Labels thus have *option value* – they offer assistance only to those who wish to use them. Second, labels have *existence value*: their mere presence can allay fears about product quality and safety (Caswell and Padberg 1992: 465). Reliable and trustworthy labelling programmes require a network of regulation and inspection, and thus can reassure consumers that products are being monitored and assessed. Labels can thus facilitate trust in food systems even to consumers who do not use them in their decision-making.

Labels can also benefit producers. By enabling consumers to reliably express preferences about credence attributes, labels help producers to derive information about consumer preferences. This allows producers to identify potential gaps in the market where demand is unmet. In addition, labelling programmes may foist fewer costs upon producers than other policies (Tonsor and Wolf 2011: 430). Just as labelling schemes have option value for consumers, some also have option value for producers: labels managed by the RSPCA, Red Tractor and Soil Association are all voluntary, enabling producers to opt in if they can afford to make the investment necessary to meet the requisite standards, but not forcing them to do so.

Because product labelling offers benefits to both consumers and producers, it deserves consideration as a means of tackling asymmetric information in animal product markets. There are numerous factors that influence the effectiveness of labels, however. Not just anyone can

reduce inefficiencies by creating a welfare label: consumers will not rely upon a label they deem to be untrustworthy, and so any effective label must be supported by credible institutions (Blandford and Fulponi 1999: 414-15).

Numerous studies have been conducted about the relative trustworthiness of labelling institutions. The government is generally not seen as a suitable institution for developing and managing a labelling scheme (Harper and Henson 1999: 21; Harper and Henson 2001: 26). This perception may stem from the fact that labelling programmes must monitor subscribing producers; this can be resource-intensive and thus may be perceived to be beyond the powers of the government. Instead, labels might be better-run by the private sector (Codron et al. 2005: 278).

This does not, however, mean that producers or food industries are viewed as trustworthy industries in this context. These organisations are frequently regarded as the least credible sources of information about animal welfare (Brook Lyndhurst 2010: 37). In the United States, the United Egg Producers (UEP), the industry's trade association representing many of the country's egg producers, released its 'Animal Care Certified' label in 2002. Despite the name implying a higher level of welfare for egg-laying hens, its standards did little to improve upon existing practices, allowing producers to continue using low-welfare husbandry practices such as beak trimming and forced moulting. The label also required birds to be provided with 67 square inches of space, while a member of the UEP's advisory committee described a provision of 80 inches as meagre. After legal action claiming that the label amounted to false advertising, the label was replaced with one stating 'United Egg producers Certified: Produced in compliance with UEP animal husbandry guidelines' (Singer and Mason 2006: 40-41). In the UK, the Red Tractor and Lion Mark labels are also industry-led.

While consumers do not trust governments or the food industry to run welfare labels, survey respondents often state the belief that the most trustworthy labels will be administered and monitored by independent groups such as animal protectionists, consumer rights organisations and organisations specialising in food issues (Harper and Henson 2001: 22, 26; Grunert 2006: 152; Kjaernes and Lavik 2007: 24; Toma et al. 2012: 598). This does not, however, mean that there is no role for governments to play in supporting labelling programmes. Government regulation may be necessary to ensure that labelling organisations can be held accountable for the provision of full and accurate information, even if it may be preferable to leave the monitoring of labels to private bodies (Caswell and Padberg 1992: 461). Governments may also be able to offer support through funding, as many organisations that run labelling schemes, such as the

RSPCA and Soil Association, are charities.⁹⁰ A trustworthy labelling scheme could therefore be run by a charity specialising in issues of animal welfare, but supported by a government (FAWC 2006: 25; Blandford and Fulponi 1999: 415).

The RSPCA Assured label therefore has significant potential to become an effective labelling scheme as it is managed by a type of organisation that is widely regarded as credible in this domain – in one focus group, participants even singled out the RSPCA as an organisation that would be well-suited to running an animal welfare label (Brook Lyndhurst 2010: 64) – and is thus worth examining in greater detail.

4. The RSPCA Assured Scheme

The RSPCA Assured label, originally launched as Freedom Food in 1994, is a voluntary programme that claims to be the only labelling scheme in the UK specifically dedicated to farm animal welfare (Food Ethics Council and Pickett 2014: 3; RSPCA Assured 2018; FAWC 2006: 12). The label appears on products in the form of a logo (see Figure 6), which indicates that the product in question has met RSPCA Assured welfare guidelines. The RSPCA has standards for the raising of pigs, laying hens, chickens, turkeys, ducks, beef cattle, veal cattle, sheep, dairy cattle, trout and salmon. In 2017, the organisation claimed that over 279 million animals – 73% of which were salmon – were covered by the scheme (RSPCA Assured 2017a).

There are at least two ways that the RSPCA Assured label enables consumers to express preferences for animal welfare. First, the label is relatively extensive, as it can be found on a range of products from multiple species. This enables consumers to more broadly express welfare preferences. Second, the scheme covers multiple facets of animal welfare, including non-production elements such as transport and slaughter, meaning that consumers with preferences in these areas can use the label to inform their purchase decisions.

With consumers better-able to express welfare preferences, production-side inefficiencies may also be reduced: the label enables producers to form a better idea of demand for higher welfare in a range of product categories. In theory, the RSPCA Assured scheme should help producers be more responsive to welfare preferences and empower consumers to express preferences for farm animal welfare.

⁹⁰ Government-charity co-operation in animal welfare issues is not unheard of: the Labour Party's 50-point Animal Welfare Manifesto (2019) committed the party to working with groups such as the People's Dispensary for Sick Animals (PDSA) to expand access to affordable veterinary care and the Dogs Trust to tackle puppy smuggling.

Figure 6 – The RSPCA Assured Label



In practice, however, there are two reasons why these benefits may not be fully realised. First, the RSPCA Assured label (and other logo labels) might be limiting the production of the highest-welfare goods by reducing incentives for producers to go beyond a scheme's minimum requirements. This occurs through a process called *adverse selection*, where higher-quality goods are forced out of the market despite there existing substantial demand for them (Pindyck and Rubinfeld 2013: 634; Schotter 1997: 540; Blandford and Fulponi 1999: 414). A producer of extremely high-welfare pork can differentiate her product from lower-welfare alternatives by subscribing to the RSPCA Assured scheme. She cannot, however, distinguish her product from competitors that merely meet the scheme's minimum requirements: both products will be labelled as RSPCA Assured. For convenience, I will refer to extremely high welfare products as RSPCA Assured+. It is likely that RSPCA Assured+ products will cost more to produce and will therefore be more expensive than RSPCA Assured products. The higher-welfare attributes of RSPCA Assured+ products, however, will not be evident to consumers: the nature of the RSPCA Assured scheme means that RSPCA Assured+ pork can only be advertised as RSPCA Assured. Consumers are also unlikely to adjust their purchase behaviour in response to their experience of the product: blind taste tests have found consumers are unable to distinguish between a range of animal products produced at different levels of welfare (Haspel 2010; López-Alt 2010; Dransfield et al. 2005: 68; Jonsäll et al. 2002: 77).

Consequently, a consumer with a preference for extremely high animal welfare faces a choice between the more expensive RSPCA Assured+ pork or cheaper RSPCA Assured pork, both of which they encounter as 'RSPCA Assured'. They therefore cannot tell whether an RSPCA Assured+ product satisfies their preference for extremely high welfare or is merely an expensive product that meets the minimum RSPCA Assured requirements. Rather than risk spending a larger sum of money on a product that may possess the same welfare attributes as its competitors, the consumer is likely to choose the cheaper RSPCA Assured pork. When this behaviour is repeated, RSPCA Assured+ producers are crowded out of the market, even if there exists demand for their products (Akerlof 1970: 488-89; Salvatore 2003: 641; Caswell and Mojuszka 1996: 1250; McMullen 2016a: 53-54). In essence, consumer risk aversion reduces willingness-to-pay for attributes that consumers cannot identify. This creates a vicious cycle where producers of the highest quality goods are driven out of the market, which further increases the risk of buying a lower quality good or 'lemon'. This leads consumers to place an even lower value on the good, which prices more producers of higher quality goods out of the market. Eventually, the only producers left are those who meet the minimum requirements for the RSPCA Assured scheme, despite there existing demand for products that exceed these

standards. To address the problem of adverse selection, a label must therefore be capable of recognising different levels of animal welfare.

Second, for the RSPCA Assured label to realise its potential to enable consumers to make informed and welfare-motivated purchases, it is surely necessary for consumers to possess at least some awareness of its underlying standards (Duffy and Fearn 2009: 672-73; McEachern and Seaman 2005: 582). One precondition for understanding a label's welfare standards is recognition of the label itself, and the limited research in this area paints a pessimistic picture. In one survey, just 22% of respondents recognised the RSPCA Assured label (McEachern and Warnaby 2004). With many consumers failing to identify the RSPCA label, it is little wonder that just 9% of respondents in another survey claimed familiarity with its underpinning standards (McEachern et al. 2007). Due to low familiarity and even lower awareness of the label's standards, it is unlikely that many consumers are using the RSPCA Assured label to make informed and welfare-motivated purchases.

To develop a more effective welfare labelling programme, it is necessary to understand why so few consumers have familiarised themselves with the standards of the RSPCA Assured label. First, it is possible that many consumers do not wish to be informed about the detail of modern animal agriculture, and instead prefer to trust the labelling organisation to deliver 'high welfare' goods, where high welfare conforms to whatever interpretation consumers hold of the concept. In support of this idea, a Brook Lyndhurst survey (2012: 23) found that 86% of respondents who were aware of the Freedom Food (now RSPCA Assured) label believed that it represented higher animal welfare. When contrasted with the 9% of respondents who claimed awareness of the label's standards in McEachern et al.'s study (2007), it is likely that some consumers associate the RSPCA Assured label with high welfare without direct knowledge of its standards. When these consumers use the label to make welfare-motivated purchases, they are thus trusting the RSPCA to deliver their preferred level of animal welfare: they are using the RSPCA label to express *uninformed welfare preferences*.⁹¹

Using welfare labels to express uninformed welfare preferences may be a useful strategy: animal welfare is a complex issue and most consumers do not have the time to become knowledgeable in this field (Norwood and Lusk 2011: 337) – the RSPCA label's lack of on-product information

⁹¹ Trusting labels to deliver high welfare may not always lead to ideal outcomes, as there will be cases where a label does not deliver upon a consumer's interpretation of high welfare: the RSPCA Assured label, for instance, permits beak trimming on free-range chickens. Presumably, consumers who place their trust in labels prefer to take the risk that their preferred welfare outcomes will not be delivered rather than take the time to engage with a label's standards.

compounds this issue. Additionally, many consumers do not want to engage rationally with animal welfare, as they find this distressing. Instead, as seen in Chapter 1 (p.26), engagement with farm animal welfare often takes place on an emotional level, with consumers seeking cues that an animal has lived a happy life (Brook Lyndhurst 2010: 40; Harper and Henson 2001: 19). By allowing consumers to make quick purchase decisions without having to engage with the uncomfortable details of animal agriculture, logo labels thus facilitate emotional engagement (the logo suggests an animal has lived a happy life) and discourage rational engagement (no information about animal agriculture is provided on-product). An effective labelling scheme must therefore cater to the needs of those wishing to express uninformed welfare preferences: many consumers want to use labels as a beacon for quick decision-making, without having to engage with its underlying standards.

Trusting labelling organisations to deliver good animal welfare transfers responsibility for securing animal welfare from consumers to the labelling organisation, which saves consumers time and energy in their decision-making. Some consumers, however, will want to know more about a label's underlying standards. This does not necessarily mean that they want to possess a full understanding of a label's requirements, but rather that they want to know whether it satisfies their particular welfare concerns: these consumers want to use labels as a way of expressing *informed welfare preferences*. The most effective labels will also facilitate this form of engagement, and must therefore strike a balance: labels must provide a cue for consumers who prefer to trust the labelling organisation to deliver high welfare, while also making it easy for interested consumers to access and engage with their standards.

There are at least three reasons why the costs of engaging with the RSPCA Assured label's underlying standards are prohibitively high for many consumers: first, the label appears on a product as a small badge or stamp, which indicates that the product has been produced in accordance with its welfare standards, but conveys no information about these standards at points of purchase (Brook Lyndhurst 2010: 37). Anyone who wishes to understand the labels' welfare criteria must therefore expend considerable effort to do so, which raises search costs. Those interested in learning more must first navigate their way through the scheme's website. If they want to identify the areas in which the scheme's standards exceed legal minimum requirements, they must also consult the relevant legislation and regulations. The curious are thus forced to consult numerous – often highly technical – sources, further increasing search costs.

Second, the RSPCA Assured standards are not consistent, which facilitates consumer confusion. This is perhaps best-illustrated by the different requirements for free-range eggs and chicken meat. Consumers could not be blamed for assuming that, given that both chicken meat and eggs are produced by birds of the species *Gallus*, chicken meat described as free-range would be produced to the same welfare standards as free-range eggs. This is not so: RSPCA standards diverge considerably in key areas, notably beak trimming: free-range laying hens may only be debeaked before they are one day old, while meat chickens may be debeaked at up to ten days of age (RSPCA 2017a: 42; RSPCA 2017b). This might represent a significant difference in welfare: it appears that beak trimming at an earlier age produces less pain, both in the short and long term, than when the procedure is conducted later. De Mol et al.'s FOWEL account (2006) ranks beak trimming as the fifth most significant element of laying hen welfare (behind access to food, space per hen, availability of perches and water, and access to nests), and finds that welfare is most adversely affected when the procedure is carried out on chicks of eight days of age or older. In contrast, the model finds that welfare is at its highest when the procedure is conducted on chicks younger than eight days. Consumers who mistakenly conflate free-range egg and meat standards in RSPCA Assured goods may therefore end up buying goods produced with practices they do not support.^{92,93}

Third, UK animal product markets feature multiple labels with differing animal welfare criteria (e.g. British Lion, Red Tractor and Soil Association). These different standards can further contribute to consumer confusion (Passantino et al. 2008: 398; Verbeke and Viaene 1999: 61).

When combined, the need to actively seek out information, consult multiple sources and avoid confusion with rival labels may mean that the costs of becoming informed are simply too high for many consumers.

Several ways of reducing these search costs have been proposed. Gemma Harper and Spencer Henson (2001: 29) suggest that combining simple label logos with education campaigns about a label's standards will help consumers make informed and welfare-motivated purchases. For this to work, however, consumers must both be exposed to and pay attention to the information campaign, and must also remember the logo and recall its standards at points of purchase. Given that modern grocery shopping is often rushed and viewed as a chore (Brook Lyndhurst 2010: 5),

⁹² It is likely that very few people will be affected by this particular issue – opposition to beak trimming seems to largely be targeted at the entire practice (see Heng et al. 2013 and Nicol 2018), rather than its use on chicks of eight days or older. Nevertheless, the point of this argument is to show that using the same descriptor to denote different standards in different products is likely to cause confusion.

⁹³ Additional differences can be found in relation to maximum flock size, the use of colonies, the extent of perch availability and environmental enrichment (RSPCA 2017a: 42; RSPCA 2017b).

these assumptions are perhaps optimistic: consumers may easily overlook a small welfare logo, and, as I will argue in the next chapter, may be deterred by their decision-making environment from actively considering ethical issues at points of purchase. Consequently, a logo label and education campaign combination may not be successful in enabling consumers to express informed welfare preferences; a more effective alternative will likely make welfare information more readily available and easier to access, ideally at points of purchase.

One way of providing more welfare information at points of purchase has been trialled in some Danish supermarkets, which have added a welfare barcode to packaged meat. When scanned at an in-store kiosk, this barcode provides detailed welfare information, relating to factors such as the animal's genetics, feed, medication and slaughter date, as well as showing pictures of the farm upon which the animal was raised (Pollan 2011: 244). Although this is an interesting example of how technology can address market failures, the welfare barcode may not sufficiently reduce search costs to the extent that consumers regularly engage with animal welfare standards. Due to the rushed nature of the modern grocery shop, expecting consumers to take the time to scan and read the information on a welfare barcode may simply be unrealistic.

In short, an effective label must enable the expression of both uninformed and informed welfare preferences. Consumers who use labels to express uninformed welfare preferences do not want to engage with a label's underlying standards, but instead trust labelling organisations to deliver good animal welfare. These consumers engage with animal welfare on an emotional level and use labels as cues that an animal has lived a happy life. To this end, labels should be simple in appearance and not contain detailed information about animal agriculture. Logo labels may be well-suited to this purpose.

Welfare labels can also be used as a means of expressing informed welfare preferences. Consumers who use welfare labels to express informed welfare preferences might also engage with animal welfare on an emotional level, but have strong feelings about the use of certain practices. Due to a lack of on-product information, which can contribute to consumer confusion about underlying standards, logo labels are not well-suited to facilitating the expression of informed welfare preferences. Attempts to resolve these issues through education programmes and technological solutions fail to appreciate the nature of purchase environments, where consumers are frequently rushed and primed to focus on the economic aspects of purchase decisions. Asking consumers to take the time to scan a barcode, or expecting them to recall welfare information in an environment that discourages them from thinking about such things

is thus unlikely to resolve the problem. A more plausible solution is to reduce search costs by directly providing welfare information on the product (Hoogland et al. 2007).

5. Textual Labels

Textual product labels provide information about a range of product attributes. In the UK, the Food Standards Agency requires all food labels to state the product name, list of ingredients, allergen information, date of minimum durability, country of origin, and more (FSA 2018). With regards to animal welfare, EU law requires table eggs to provide textual information about the production system, i.e. enriched cage, barn, free range or organic (Commission Regulation 589/2008; FAWC 2006: 9). Labels on table eggs are not required to provide information about other features of welfare, and other animal products and products containing animal material, including eggs, do not require any form of welfare labelling. Producers of other animal products can use marketing terms including 'free range', 'organic' and 'cornfed', provided they meet certain production conditions (Council Regulation 834/2007; Commission Regulation 543/2008).

Although limited, the textual labels currently required by law may offer a more effective means of enabling consumers to express informed welfare preferences than their logo counterparts. The information provided by textual labels makes it easier for consumers with preferences for broad production systems and organic produce to express these preferences in the market, and allows producers to reliably identify the strength of demand for different modes of production. Indeed, the long-term trend of increasing sales of free-range eggs in the UK, increasing from 32.4% of production in 2008 to 53% by Q2 2019 (DEFRA 2019b; CIWF 2013; DEFRA 2013) has surely been at least partially facilitated by the introduction of mandatory labelling of production systems (FAWC 2006: 9).

On the other hand, it is clear that current textual labels are inadequate in three ways. First, the information provided by these labels is incomplete, as it only covers broad production methods. Consumers thus lack the requisite information to express informed welfare preferences about issues including the use of specific production practices, transport and slaughter (FAWC 2006: 20). Without this information, consumers are likely to make choices that violate their welfare preferences, resulting in markets providing suboptimal outcomes.

Second, the limited nature of current textual labelling may distort public understanding of animal welfare. Although textual labels enable consumers to reliably express preferences for the

use of particular systems in egg production, this may lead people to overemphasise the importance of broad production systems as a constituent of animal welfare; textual labels may encourage people to place a greater emphasis on the importance of access to free range and may even lead people to equate access to free range with high welfare. Although we cannot be certain that labelling has influenced public perceptions of hen welfare, it appears that many people do equate free-range and high-welfare: in group interviews led by Gemma Harper and Spencer Henson (1999: 14; 2001: 25), participants identified free-range eggs as their most common 'high-welfare' purchase, indicating an association between free-range and high welfare. A further study by Pettersson et al. (2016) lends support to this idea, finding that, when consumers of free-range eggs were asked to rate the welfare of free-range hens on a scale of one to ten, the mean score was 8.7, with 85.9% of respondents providing a score of eight or above. Many people also believe that free range is the most important constituent of hen welfare: 56.1% of participants in the Pettersson et al. study considered access to an outdoor space to be the most important general factor affecting hen welfare.⁹⁴

Scientific opinion, however, disagrees: access to free range is generally considered to be a minor constituent of laying hen welfare. Pettersson et al. repeated their survey with a group of animal welfare specialists at the University of Bristol, whose answers differed markedly from the public's. The mean welfare score was 5.9, and only 14.7% rated access to an outdoor space as the most important constituent of welfare. Perhaps most tellingly, not a single specialist scored free range systems 8 or above. Their responses accord with De Mol et al.'s FOWEL model of laying hen welfare (2006), which ranks access to free range as the 19th most significant element of hen welfare out of 25. While access to an outdoor space does make a positive contribution to welfare, its impact may be far lower than many other factors. Free-range production is also compatible with a range of practices that are detrimental to hen welfare, such as large flock sizes, suboptimal space per hen and beak trimming, meaning that access to free range may not be a guarantee of high welfare.

Irrespective of whether the current state of labelling is responsible for the exaggerated importance given to free-range production in hen welfare, attention appears to be diverted away from other important areas of welfare: many consumers focus on broad production systems at the expense of specific production practices and issues in transport and slaughter (Brook Lyndhurst 2010: 39). Lack of awareness of welfare issues in these areas may contribute to market

⁹⁴ The authors (2016: 2011) recognise the limitations of their survey, which was sent out to subscribers to the e-mail list of a free-range egg company. Consequently, respondents were surely more likely to hold a favourable opinion of free-range production systems.

inefficiencies if consumers buy free-range products in the mistaken belief that they are supporting higher welfare standards; because their understanding of animal welfare is limited, their actions may be less effective than intended. These consumers may also be vulnerable to manipulation by producers, who, by emphasising the free-range aspects of production, could mislead them into thinking that their purchases are of greater benefit to farm animal welfare than they actually are.

Given that purchases informed by existing textual labels are not reliable indicators of preferences for many elements of animal welfare, producers may struggle to discern demand for contributors to welfare beyond production systems. A consumer who is opposed to, say, beak trimming, or one who prefers hens to be provided with sufficient space for effective dust bathing, cannot currently identify products that satisfy these preferences, and so will be unable to signal their demand for these goods. This in turn hides the true demand for the use or disuse of individual production practices from producers and prevents markets from efficiently satisfying consumers' preferences.

Third, as with logo labels, the process of adverse selection can reduce producers' incentives to exceed minimum welfare requirements demanded by textual labels. Producers who exceed the requirements for, say, free-range production may struggle to advertise this fact to consumers, who may be unwilling to pay more for products when the benefits of doing so are uncertain.

In short, current textual labels do little to enable consumers to express informed welfare preferences. These labels only enable consumers to express welfare preferences for broad production methods and are thus incomplete. They may also distort consumers' understanding of animal welfare and limit producer incentives to exceed minimum standards.

One possible way of responding to the limitations of existing textual labels is to expand them to cover other elements of welfare across a range of species. More on-product information could provide consumers with a more balanced understanding of welfare, encourage producers to implement welfare improvements across a range of attributes and, most importantly, enable consumers to express welfare preferences for these attributes.

There are at least three reasons why an expansion of textual welfare labelling cannot resolve the informational issues discussed in this chapter, however. First, there may not be enough space on a product's packaging to provide detailed welfare information: De Mol et al.'s FOWEL model (2006) identifies 25 constituents of laying hen welfare, Bracke et al.'s SOWEL model of pregnant sow welfare (2002) recognises 37, and Ursinus et al.'s COWEL model of dairy cattle welfare

(2009) contains 42. Product labels must already carry information about numerous other product attributes, including brand name, product description, ingredients, nutritional profile, disposal information, weight, storage and use instructions, returns policies, best before / use by date, allergy information and country of origin. Consequently, there simply may not be enough space on a product to provide such a level of detail about farm animal welfare.

Second, expanded textual labelling is likely to be incompatible with the expression of uninformed welfare preferences. Many consumers do not want to receive detailed information about animal agriculture, and, when exposed to such information, will simply ignore or otherwise fail to engage with it (Brook Lyndhurst 2012: 7). Information presented this way simply does not resonate with the emotional nature of many consumers' engagement with farm animal welfare. These consumers instead want welfare labels to provide a simple visual cue that an animal has lived a happy life, rather than detailed textual information (Brook Lyndhurst 2010: 58).

Third, a mass of detailed textual information could also lead to *information overload*, where, overwhelmed by information, consumers disengage with an issue and ignore all available information. This issue is especially pertinent in supermarket choice environments, which already provide consumers with a deluge of information (Brook Lyndhurst 2010: 38). In addition, consumers are forced to choose between a barrage of similar products and tend to view grocery shopping as a chore to be completed in the face of time pressures. The sheer quantity of information on display, in conjunction with the limited time committed to grocery shopping, makes it likely that many consumers make only limited use of information on product labels (Caswell and Padberg 1992: 462). In short, modern shopping environments can largely be characterised as ones of information overload (Grunert 2006: 154; Tonsor and Wolf 2011: 434). If handled badly, welfare labelling will merely exacerbate this information overload and will be consequently overlooked by consumers (Brook Lyndhurst 2010: 38). Welfare labels must therefore be selective about the information they provide.

An effective welfare label must therefore balance the interests of consumers with informed welfare preferences and those with uninformed welfare preferences. To enable the expression of uninformed welfare preferences, labels must provide a simple visual cue and cannot present too much information. To enable the expression of informed welfare preferences, labels must enable consumers to recognise multiple levels of welfare provision and provide easy access to detailed information about how a good was produced.

6. An Effective Welfare Label?

In light of the above analysis, an effective welfare label must distinguish between different levels of welfare and must make it easy for consumers to access detailed information about the label's standards.

a. Tiered Labels

A tiered label is one which utilises a colour-based (bronze, silver and gold, or red, amber and green) or numerical tiered structure to denote different levels of welfare. A tiered label can cater to consumers with uninformed welfare preferences, who can use the label as a cue that farm animals have lived good lives, while avoiding engagement with its underlying standards. The label also offers value to producers, who can move up the welfare scale when they exceed certain standards. This thus reduces the impact of adverse selection, although we can still expect producers to cluster around the minimum requirements for each tier.

This is not an ideal outcome, as it does not entirely remove the problems posed by adverse selection. It is likely to be the best outcome, however. The only system that would be immune to adverse selection is a continuous scale label, perhaps based on species-specific models of animal welfare that lend themselves towards a 0-10 scoring system, such as De Mol et al.'s FOWEL model (2006), Bracke et al.'s SOWEL model (2002) or Ursinus et al.'s COWEL model (2009). Consumers with informed welfare preferences would likely encounter difficulties with these labels, however, as it would not be obvious what a score represents, making it hard to recognise whether a product satisfies specific welfare preferences. A score of say, 5, could be reached using multiple production systems (e.g. a higher-welfare enriched cage system or a lower-welfare free-range system), and so a continuous scale scoring system provides little information about the trade-offs made to achieve a given level of welfare. Consequently, despite being unable to entirely banish adverse selection, a tiered system likely creates a label that consumers can easily interpret.

Tiered labels are not widely used in UK animal product markets, although a similar model is used to inform consumers about products' nutritional content. Inspiration for a tiered label in an animal welfare context can be drawn from the United States, where the high-end food retailer Whole Foods worked with the subsidiary non-profit Global Animal Partnership (GAP) to implement a five-step welfare labelling programme. All fresh beef, pork, chicken, lamb and turkey (except kosher turkey) sold at Whole Foods is graded using the GAP label. Relevant products must reach Step 1 to be eligible for sale: this requires that no cages or crates are used

in production, and also that some stipulations about population density and diet are met. Step 2 requires animals to be raised in an enriched environment that encourages natural behaviours, Step 3 requires outdoor access and Step 4 can be reached if animals are raised on pasture year-round. Step 5 requires an animal-centred approach to husbandry: no physical alterations (mutilations) are permitted. Step 5+ recognises producers whose animals spend their entire lives on the same farm, which means they are also slaughtered on-farm (Whole Foods Market 2019; GAP 2019a; GAP 2019b).

The labels themselves, as seen in Figure 7, indicate the tier in which the product falls, and also supply a brief description of what this means. The use of colour distinguishes between the different levels, but not in a way that might deter consumers from buying lower-tier products: consumers have stated an unwillingness to buy products with a red or bronze label, even when these colours signify a level of welfare above legal minimums (Brook Lyndhurst 2010: 60). In short, consumers with uninformed welfare preferences can seek products marked out with a colour or a number, or use the textual descriptions to identify products made using processes that accord with their interpretations of ‘happy animals’ without having to engage in the minutiae of animal husbandry.

Figure 7 – The GAP Animal Welfare Certified Label (GAP 2019a)



Producers can also use the 5-Step label to identify consumer demand and willingness to pay for certain aspects of welfare. Although the 5-Step label clearly states whether a producer has provided their animals with amenities such as access to pasture or environmental enrichment, there are limits: the more distasteful elements of animal agriculture, such as the use of mutilations, are not explicitly addressed. If they were, they might deter some consumers from paying attention to the label, and could even drive consumers away from a product (Brook Lyndhurst 2010: 59).

Consequently, it seems that no labelling solution can fully-rectify the informational failures in animal product markets. Consumers might have informed welfare preferences for the disuse of

certain practices, and so would benefit from labels that make it clear which products were not produced using these methods. On the other hand, other consumers are likely find labels that mention these practices to be distasteful and so would be unlikely to pay attention to other welfare issues addressed by the label, including those that they support.

Interested consumers must therefore do their own research to determine a label's specifics: for instance, the euphemistic Step 4 label ('Animal Centred') prohibits beak trimming in laying hens (GAP 2019c). Given that providing on-product information about the use of mutilations and similar practices would damage a label's effectiveness among consumers with uninformed welfare preferences, an effective label must compensate by making this information easy to access and understand.

b. Access to Information about Labelling Standards

To engage consumers with informed welfare preferences, an effective label must make its full standards readily available and easily understood. This allows consumers to identify which tier best satisfies their preferences, and then subsequently use the label as a beacon. A consumer with a preference for eggs produced by hens with intact beaks saves time in the supermarket once she learns that products with a Step 4 label come from animals that have not been subjected to mutilations, as she no longer has to examine multiple products until she finds one that satisfies her preferences.

It is imperative that efforts are made to reduce the search costs of finding this information in supermarkets, however: because shopping is viewed as a chore, and, as I argue in the next chapter (p.159), consumers tend to operate automatically in supermarkets, they will generally avoid effortful actions in supermarket, including seeking out welfare information. Eye-catching shelf displays and banners that provide information about a label's standards may draw consumer attention if designed appropriately, however. Advertising and information campaigns could be more effective around tills, where consumers are likely to spend time waiting: consumers face lower search costs when queuing, as they have reduced opportunities to spend their time elsewhere, meaning they may engage more with leaflets.⁹⁵

Perhaps the most important component of a label that responds to the needs of consumers with informed welfare preferences, however, is its online presence. A label's website needs to be easy

⁹⁵ The aim of providing information at tills is not to encourage immediate purchases, but rather to inform consumers so that they may be more receptive to welfare labels in the future.

to navigate. The requirements for each tier must be easy to understand, and, ideally, the label's position relative to legal minimum standards will be made clear. The website itself must be well-advertised, to enable consumers with informed welfare preferences to access it outside the supermarket, when they are not performing the consumer role and may therefore be willing to spend more time learning about a label's standards.

Once again, Whole Foods' Animal Welfare Certified label provides plenty of inspiration in this area. The label itself prominently advertises its website, which is easy to navigate and provides both a brief overview of what each step entails for different species (see Figure 8) and a detailed guide of every standard that producers must meet. It does not, however, show where these standards diverge from legal minimum standards. The label is also advertised in-store in appropriate locations (Whole Foods 2019).

c. Limitations

Despite being superior to the RSPCA Assured label in several ways, the label proposed in this chapter – based upon the Animal Welfare Certified label – encounters at least three limitations. First, there is limited scope to apply it to all products sold in a market. As discussed in Chapter 5 (p.122), a welfare label imposed upon imports may be perceived as a technical barrier to trade and so could violate WTO regulations (WTO 1995: 118; FAWC 2006: 7). Such a label could be imposed upon domestic producers, and retailers could insist that imported products subscribe to the label, but it is unclear whether a government could mandate that imports display this label.

Second, labelling schemes carry costs relating to the development and maintenance of standards, in particular in relation to monitoring and enforcement of standards. These costs are borne by producers, who generally pay a fee based on the quantity of product sold, and passed on to consumers. These costs need to be low, as they will be paid by all consumers who buy labelled products, not just those who have preferences for the label. As discussed in Chapter 2 (p.61), however, there does appear to be substantial willingness to pay for farm animal welfare, and, given that welfare labels can possess existence value through the trust they inspire in food systems, there may be willingness to pay for welfare labels among consumers who do not directly use them in their decision-making. The RSPCA Assured label charges £36.00 as a joining fee, plus varying fees depending on the product. Table eggs are charged at £142.80 for the first 6,000 hens and £68.40 for each subsequent 6,000 hens. A further 5p per 30 dozen eggs sold is

Figure 8 – Explaining the Animal Welfare Certified Label (GAP 2019c)



charged as a license fee (RSPCA Assured 2017b).⁹⁶ Any competing tiered label would surely need to approximate the costs of its rivals.

Third, it is unclear how the label proposed in this chapter, or any welfare label for that matter, should be applied to processed foods. Processed foods such as bacon quiches are composed of multiple animal products, and if some adhere to a label's standards and others do not, or if different ingredients comply with different tiers, it is not apparent how the product should be labelled (FAWC 2006: 14). To label a product according to its highest-welfare ingredient may be misleading, especially if the ingredient is present only in small quantities. For the same reason, labelling according to a product's lowest-welfare ingredient could also mislead consumers. This concern is not insurmountable, but must be addressed before welfare labels can be applied to processed goods.

7. Summary

To summarise, the issue of welfare labelling is multi-faceted and complex, and an effective label must balance a range of competing interests and concerns. It is important that efforts are made to develop effective labels, however, as an ineffective label can exacerbate market failures. Without labels, or with inadequate welfare labelling, consumers may struggle to express preferences for welfare and their understanding of welfare may be warped by a focus on a single issue, while producers may struggle to discern consumer demand for certain practices and have little incentive to exceed a label's minimal requirements.

There does not appear to be a perfect solution, although there are clear steps that can be taken to improve upon the UK's existing welfare labels. A successful welfare label must balance the interests of both consumers with uninformed welfare preferences, who want a broad indicator that an animal has lived a happy life, and consumers with informed welfare preferences, who may care strongly about certain welfare issues and require easy access to information about a label's standards.

I argued that, due to the emotional nature of consumer engagement with animal welfare, a simple logo provides the most effective foundation for a welfare label. Such a label serves as a beacon for consumers with uninformed welfare preferences, and, if information about the label

⁹⁶ In comparison, the Soil Association charges an annual membership fee of £750, plus 0.03 multiplied by sales between £250,001 and £2,350,000, plus 0.0015 multiplied by sales between £2,350,001 and £4,350,000, plus 0.001 multiplied by sales between £4,350,001 and £22,500,000, plus 0.0003 multiplied by sales over £22,500,001 (Soil Association 2019).

is made readily-accessible and easily-digestible, also to consumers with informed welfare preferences.

The question of how to provide this information is a difficult one. Displaying detailed welfare information about potentially distasteful practices may negatively impact the use of labels by consumers with uninformed welfare preferences, due to the emotional nature of their engagement with animal welfare and the risk of information overload. On the other hand, the search costs of finding detailed welfare information are often prohibitively high for consumers with informed welfare preferences. Perhaps the best solution is to reduce search costs for consumers with informed welfare preferences by providing a small amount of information about carefully-selected production practices on the product itself, and supplementing this with clear advertising for an easy-to-understand website.

On the production side, the problem of adverse selection is difficult to eradicate, although the use of a tiered label will at least increase consumer choice. Production is likely to cluster around the minimum standards for each tier, as producers have little incentive to bear the increased production costs associated with welfare improvements when they cannot advertise this to consumers. Providing more tiers increases the range of standards around which producers are likely to cluster, enabling a greater range of welfare preferences to be expressed. This is not a perfect solution – consumers with preferences for levels of welfare between the tiers and beyond the top tier will struggle to find products that fully satisfy their preferences – but it still represents an improvement upon the existing state of affairs.

Chapter 7

A Trip to the Nudge Supermarket

In this Chapter

- An assessment of how far the axioms of rational preferences reflect decision-making.
- An introduction to Daniel Kahneman's dual-system theory of decision-making.
- Support for the claim that choice environments can influence automatic behaviours.
- The proposal that supermarket design influences consumer behaviour by encouraging the expression of self-interested preferences.
- A discussion of how nudges can be used to support farm animal welfare.

1. Introduction

What you intended to buy when you go into IKEA is rarely all you leave with. Somehow, almost everyone reaches the checkout pushing trolleys laden with things they did not know they needed, or even wanted, when they entered – as many as 60% of IKEA purchases were not on people's shopping lists when they entered the store (Penn 2011). This is no coincidence: IKEA stores are designed to increase impulse purchases. Stores have no windows, which disorients shoppers and discourages them from thinking about time, the outside world and everyday life. This confusion and disorientation make it difficult for consumers to think rationally; their behaviour becomes more automatic and more emotional, which increases impulsive tendencies.

IKEA seeks to take advantage of this shift to automatic behaviours. Its stores utilise a one-way system, which makes it difficult to turn back and encourages shoppers to pick up items as they see them out of fear that they will have no other opportunity to do so. The labyrinthine weaves and bends in the route make it impossible to see what is coming next and lends an air of mystery to the IKEA experience. Mystery increases consumer liking, which encourages shoppers to delve deeper into the heart of the store. This, along with the fact that the one-way system forces shoppers to wend their way past every available item, increases the probability that customers will buy more (Jansson-Boyd 2018; Zambasri 2017; Mitchelson 2018; Tucker 2011).

IKEA designs its stores with the intention of nudging shoppers to make impulse purchases. The company's ability to do this flies in the face of mainstream economic theories of rationality, which assume that preferences are relatively fixed and well-defined, and obey the axioms of

transitivity, independence, completeness and continuity. In this chapter, I consider challenges to these axioms, in particular the continuity axiom – if a is preferred to b , then situations appropriately similar to a will be preferred to b . With reference to Daniel Kahneman’s influential dual-system theory of decision-making and priming, I argue that, as IKEA knows, choice environments can influence consumer behaviour and expressions of preferences, leading to violations of the continuity axiom – a might be preferred to b in one type of choice environment, but b will be preferred to a in another.

I contend that the influence of choice environments renders preferences for farm animal welfare unstable. I propose that supermarket choice environments discourage the expression of ethical concern and encourage self-interested attitudes, with implications for the expression of preferences for farm animal welfare. I then consider how policy can integrate nudge principles in supermarket design for pro-animal welfare ends, with particular attention paid to the People’s Supermarket in Holborn, which, in summer 2019, joined forces with the Royal Society for Public Health and Slimming World to nudge customers to make healthier choices.

2. Irrational Preferences

We saw in Chapter 3 (p.67) that one of the characteristics of a perfectly competitive markets is that consumers behave competitively. Competitive behaviour requires not only that consumers seek lower prices, but also that they behave rationally. A rational consumer’s preferences will satisfy four assumptions. Her preferences will be transitive, meaning that if she prefers a to b , and b to c , she will prefer a to c . Her preferences will also be independent: her preference for a over b is not affected by the presence or absence of alternatives. Her preferences sets will be complete: she will be able to rank all possible alternatives and so, for any two options, will either prefer one over the other or be indifferent between the two. Last, the rational agent’s preferences are continuous, meaning that, if she prefers a to b , states of the world appropriately similar to a will be preferred to b .

In the real world, however, people’s preferences and preference sets often do not obey these axioms. First, using multiple criteria to assess options can cause people’s preferences to become intransitive (May 1954: 7; Lewis 2017: 103-5). Intransitive preferences are often considered irrational because they can lead to a ‘money pump’ outcome, where consumers are willing to pay to swap x for y , y for z , but also z for x , *ad infinitum* (Davidson et al. 1955: 145-46). In a 1954 study, Kenneth May identified one instance of intransitive preferences. May presented 62 college students with three hypothetical marriage partners, x , y and z . Each was described in

terms of three qualities: wealth, intelligence and attractiveness. Candidate x was described as highly intelligent, plain looking and reasonably well off; y was intelligent, very good looking and poor; z was fairly intelligent, good looking and rich. None of the options were described as so poor, unintelligent or unattractive as to be immediately eliminated. The students were repeatedly presented with two of the three candidates, and asked to rank them in order of preference, i.e. they were asked if they preferred x to y , or y to z , or z to x , but were never asked to rank all three simultaneously. May found that 17 of the 62 students (27%) held intransitive preferences.⁹⁷

Second, the use of ‘decoy options’ by marketers undermines the strength of the independence axiom. Dan Ariely (2008: 1-22) identifies a decoy option used as a marketing device by *The Economist* magazine. The magazine offers a website subscription for \$59, and a combination print and website subscription for \$125. Ariely asked 100 students which option they preferred: 68 preferred the online subscription and 32 preferred the combination subscription. Ariely then introduced *The Economist’s* third subscription option: a print-only subscription, which cost \$125. We can (rightly) predict that no student would prefer this new subscription as it is inferior to the combination subscription. If preferences are independent, the introduction of the new option should be irrelevant. In fact, whereas 32 students preferred the print and online subscription in a two-option choice, 84 preferred it when the inferior print-only subscription was introduced. In contrast, of the 68 who initially preferred the online subscription, only 16 continued to do so when also offered the combination subscription. The introduction of an inferior third option caused preference reversals by making the combination subscription appear better value relative to the print-only subscription (Tversky and Simonson 1993: 1181).⁹⁸

Third, the assumption of completeness is not always reflective of how people make choices: in situations such as Sophie’s Choice-style moral dilemmas, a person simply cannot choose between a and b . In *Sophie’s Choice* (Styron 2004), the eponymous character is sent to the Auschwitz concentration camp. Upon arrival, she is presented with a grisly dilemma: one of her two children will be sent to the gas chambers, while the other will be allowed to continue living in the camp. Sophie must choose who will be condemned to an immediate death, and who will be spared. If she refuses to decide, both children will be executed. Many people in such situations, like Sophie, simply cannot choose. This inability to choose should not be confused

⁹⁷ Another instance of intransitive preferences was discovered by Amos Tversky (1969), who found that similar gambles were ranked according to payoffs, while more extreme gambles were ranked by probability, leading to intransitive orderings.

⁹⁸ Tversky and Simonson (1993), Huber et al. (1982) and Wedell and Pettibone (1996) provide further examples of violations of the independence axiom.

with indifference: if it were, the addition of a trivial sum of money to one of the options would tip the balance. Sophie's choice would be no easier if she had to choose between (a) sparing Eva and sending Jan to the gas chamber, or (b) sparing Jan, sending Eva to the gas chamber and receiving a token sum of money (Feldman and Serrano 2010: 11). It thus appears that Sophie's preference ranking for the lives of her children is incomplete.

In short, there are certain, predictable situations where people's preferences are unlikely to obey the axioms of transitivity, independence and completeness. So far, however, I have not reviewed whether the continuity axiom – which states that if a is preferred to b , then situations appropriately similar to a will be preferred to b – is ever predictably violated. This review will occupy me for the remainder of this chapter. I will evaluate the influence of choice environments on consumers' preferences, thought processes and decision-making, and argue that choice environments can affect the kinds of values that we express towards goods such as farm animal welfare. I propose that the influence of choice environments can lead consumers to prefer a to b in one context, but prefer b to a in another. The two options have not changed, only the context in which consumers make their decisions, and so this likely represents a violation of the continuity axiom. As such, in contrast with the assumptions of economists and philosophers of economics (see Varian 2010: 118 and Reiss 2013: 35 for instance), I propose that preferences for some goods, including farm animal welfare, can be unstable.

To understand the mechanisms which lead to continuity violations, I must adopt a theory of consumer decision-making. To this end, I will turn to the dual-process theory advanced by Daniel Kahneman, a psychologist whose research into decision-making, largely conducted with Amos Tversky, laid the foundations for behavioural economics and earned Kahneman the 2002 Nobel Memorial Prize in Economic Sciences.

In his 2011 book *Thinking, Fast and Slow*, Kahneman proposes that thought processes can be divided into two categories. System 1 processes are fast, emotional, automatic and effortless. System 2 processes, on the other hand, are slower, rational, conscious and use up cognitive

resources (Kahneman 2012: 20-21).^{99,100} In the following sections, I discuss Kahneman's account of mental processes in greater detail, and show how environmental cues found in supermarkets can influence both System 1 and System 2 processes in ways which hinder the expression of altruistic preferences for farm animal welfare. I then discuss how policy can nudge consumers to express altruistic preferences for farm animal welfare in supermarket environments.

3. Thinking, Fast: Choice Environments and System 1 Processes

Kahneman proposes that System 1 processes are responsible for activities including: detecting that one object is more distant than another; orienting to the source of a sudden sound; completing the phrase 'bread and ...'; making a disgust face when shown a horrible picture; detecting hostility in a voice; answering $2 + 2 = ?$; reading words on large billboards; driving a car on an empty road; identifying a strong move in chess (if you are a chess master); understanding simple sentences; and recognising that a 'meek and tidy soul with a passion for detail' resembles an occupational stereotype (Kahneman 2012: 21; Norman 2013: 206).

In short, System 1 processes require little effort and generally occur automatically, although some, such as those responsible for chewing, can be brought under voluntary control. System 1 processes can be innate, such as those responsible for perceiving the world around us, orienting attention and avoiding losses, or they can become automatic through practice, as with reading, understanding social nuances and making associations between ideas (e.g. knowing the capital of France) (Norman 2013: 45-47).

⁹⁹ Dual-system theories of thinking have been hugely influential in both psychology and numerous other fields. Professor Steve Peters, a sports psychologist who has worked with British Cycling, British Athletics, Liverpool Football Club and the England football team, proposes a dual-systems theory similar to Kahneman's in his *The Chimp Paradox* (2012); Professor Robert Cialdini (2016) uses dual-system thinking in the context of the influence literature to explain marketers' persuasion tactics; Don Norman (2013) discusses dual-system thinking in relation to product design; and dual-system thinking has been used by marketers to explain why people make impulse purchases, and how they can be encouraged to make these purchases (Xiao and Nicholson 2013).

¹⁰⁰ Although popular in both psychology and economics, dual-systems theories have nevertheless been subject to criticism. First, research indicates that System 1 and System 2 are not in fact discrete parts of the brain, but instead share neural pathways to complete different tasks. There may also be far more than two systems that influence our thinking. These are hardly fatal flaws in the model: some writers (e.g. Evans and Stanovich 2013) now refer to Type 1 and Type 2 to distinguish between different types of mental processes without implying that these processes are performed by distinct systems or parts of the brain. Of greater concern are criticisms that dual-process theories only explain how people solve puzzles and complete computational tasks in laboratory conditions, and offer little about how we reason on a day-to-day basis (Grayot 2019). Nevertheless, given the popularity of dual-process theories in behavioural economics, I will continue to use Kahneman's two-system model, but recognise that these theories can be criticised.

System 1 processes are responsible for many of our quotidian actions and behaviours. When crossing a road, we do not have to think about stepping down off one kerb and up onto another; System 1 processes allow us to automatically identify the slight changes in elevation and recognise the need to step down and then up. These processes let us navigate and interact with our environments smoothly and with little effort, which saves our cognitive resources for more strenuous tasks.

Many System 1 processes are thus responsive to environmental cues. For the most part, this is extremely useful, as it allows us to navigate the world around us with little cognitive strain. On occasion, however, the automaticity and responsiveness to external cues that are hallmarks of System 1 processes lead to undesirable outcomes. A prime example is provided by the so-called Norman door, named after author Don Norman, who highlighted the phenomenon in his 1988 book *The Design of Everyday Things*. Almost everyone will be familiar with the Norman door, although they may not recognise it by name. They will also likely have been frustrated by one at some point in their lives. A Norman door features a long, vertically-oriented cylindrical handle, of the kind shown in Figure 9. People who wish to pass through the Norman door are thus generally inclined to grab the handle and pull it – the handle is perceived to signify that the door opens inwards (Norman 2013: 18). Alas, the Norman door opens outwards, and therefore requires a push – some will even be adorned with a push sign, although these signs are generally ignored due to the strength of the victim's automatic response to the 'grab handle'. The Norman door's design unintentionally exploits people's automatic response to a grab handle to produce an unwanted outcome – a futile tug on an uncooperative door (Norman 2013: 1-2). By putting human-centred design at the heart of choice architecture, for instance by building doors that use push plates on the push side, we can create choice environments that nudge people to make better automatic choices (pushing rather than pulling) and produce better outcomes (frustration-free door encounters) (Thaler and Sunstein 2009: 89-90).

4. An Introduction to Nudge

The concept of nudge, also referred to as libertarian paternalism, was popularised in a 2008 book by Nobel-winning economist Richard Thaler and legal scholar Cass Sunstein. Nudge theory recognises that our automatic behaviours are affected by the contexts in which we make decisions (Thaler and Sunstein 2009: 3). Thaler and Sunstein propose that we should adjust choice architecture to encourage people to make better choices and produce better

Figure 9 – The Dreaded Norman Door



Despite many of its buildings having won architectural awards for design, even the University of Durham is not immune to the allure of a Norman Door.

outcomes in ways that do not limit or burden consumer choice (Thaler and Sunstein 2009: 5-6).

The question of what constitutes a better choice or outcome is a contentious one. To avoid charges of paternalism, Thaler and Sunstein argue that policy nudges should not pursue some independent or 'objective' conception of the good, but rather ought to help people better satisfy their preferences (Thaler and Sunstein 2009: 5-6).

Nudges are increasingly becoming as part of the policymaker's toolkit. In the UK, the Behavioural Insights Team (BIT), often referred to as the Nudge Unit, was established in 2010 with the task of using research from the behavioural sciences to develop more effective policies (Halpern 2015: 8-9). One of its best-known interventions was introduced in 2012, where workplace pension sign-ups were changed from an opt-in to an opt-out. By 2015, more than five million additional UK workers had signed up for pensions. A similar scheme in the United States has proved hugely popular, with 9 out of 10 workers affected by the policy supporting it, and 7 out of 10 of those who choose to opt out also supporting it (Halpern 2015: 62-65). Shifting the default option to automatic enrolment did not limit choice or coerce people, but rather took advantage of inertia to better-satisfy stated preferences to be part of a pension plan.

As the pension interventions show, nudges typically seek to influence automatic behaviours. Nudges are therefore undoubtedly relevant to this thesis: in addition to contributing to frustrating encounters with Norman doors, System 1 processes are also a major driver of consumer behaviour in supermarkets. For many people, the weekly supermarket shop is a chore (Brook Lyndhurst 2010: 5) that they want to deal with as quickly as possible; consequently, they seek to limit mental engagement with the task. System 1 processes are faster and more effortless than System 2 processes, and thus allow people to be cognitively disengaged in the supermarket (Kahneman 2012: 46). Heuristics are hallmarks of System 1 thinking and are often used to streamline supermarket decision-making. For instance, if you are looking to buy toothpaste or shampoo, you may automatically simplify your decision-making process by considering only those products that are on offer. When buying bread, you may stick to the brand or particular product that you always choose – doing so requires less time and effort than considering every available option (Binnekamp and Ingenbleek 2006: 173; McEachern and Schroeder 2004: 500; Verplanken and van Knippenberg 1998; Just 2011: 101).

Using heuristics to streamline decision-making in supermarkets may be necessary, especially under time constraints that limit consumers' opportunities to engage in reflective decision-making. People make an average of 200-300 decisions regarding food consumption every day

(Just 2011: 99), but, in 2012/13, the average supermarket trip took just 47 minutes (Statista 2018). Given this, and the sheer variety of goods that most supermarkets now sell – the average supermarket offers over 30,000 items (Schwartz 2005: 12), while a search for ‘tea’ on the Tesco website yields 176 items – using slower System 2 processes is simply not a viable way of conducting an entire grocery shop under time pressures. Instead, automatic System 1 thinking is likely to dominate consumer decision-making in the supermarket, with System 2 processes contributing intermittently.

In the same way that the Norman door’s choice architecture influences people’s automatic behaviours, a supermarket’s choice architecture – the way it is designed, how products are displayed and distributed around the store and how products are marketed in-store – can influence automatic behaviours. Many supermarkets have realised this, and, like IKEA, adjust their choice architecture in ways that increase revenue, often to the detriment of consumers (Evans and Stanovich 2013: 229). Some of these strategies, which nudge consumers to make unhealthy choices, were outlined in a 2019 report published by the Royal Society for Public Health (RSPH) and Slimming World. The report also proposed a range of design features that could nudge consumers to make healthier choices. Many of this report’s recommendations were implemented by the People’s Supermarket, a community interest company based in Holborn, London. Together, the report and the supermarket provide insights into how supermarkets can be designed to promote animal welfare, by guiding automatic behaviours towards higher-welfare products and away from lower-welfare ones.

5. System 1 Nudges

The RSPH and Slimming World report identifies three areas where supermarket design encourages consumers to automatically engage in unhealthy behaviours, of which two are relevant for our purposes.¹⁰¹ First, the use of offers such as ‘buy one get one free’ on unhealthy products encourages consumers to make unhealthy purchases that they may not have otherwise made. This strategy is highly effective: 36% of respondents in the RSPH’s report claimed to have bought unhealthy food simply because it is on offer, while 33% of 18-24 year-olds stated that offers have led them to consume unhealthy food, even though they were not hungry (RSPH and Slimming World 2019: 3, 12).

¹⁰¹ The third technique is to place healthy foods next to unhealthy ones, encouraging consumers to ‘treat’ themselves for making healthy choices.

These outcomes can be attributed to System 1 rather than System 2 because, at the point of purchase, these consumers have not recognised that they have a choice. System 1 processes do not allow for ambiguity or alternative interpretations (Kahneman 2012: 80), meaning that factors such as the feeling or absence of hunger and preferences for healthy living are not considered in decision-making – as we will see below, making trade-offs between competing values is a function of System 2. Instead, heuristic-driven consumers respond to cues such as brightly coloured sale signs or discount banners, and, when their heuristics relate to price or value considerations, it appears that healthfulness and hunger do not enter their decision-making processes. The upshot is that the use of offers can exploit heuristic decision-making by nudging consumers to buy goods without considering whether they actually want them or if they cohere with lifestyle preferences.

Discounts and offers could be used to encourage consumers to act on their preferences for higher-welfare products. Applying discounts and offers to higher-welfare goods will attract the attention of consumers whose purchase decisions are primarily motivated by price heuristics, and thus encourage purchases. Conversely, an absence of discounts and offers on lower-welfare goods will diminish the likelihood that these goods trigger value and price heuristics, and will nudge consumer away from these products.

The second way that supermarkets nudge people to buy unhealthy goods is by placing high-profit items, often branded goods, in areas that are most likely to attract consumers' attention. Consumers' attention is typically drawn to eye-level items (Hidayatno and Komarudin 2009) – one marketing adage runs 'eye level is buy level' – and products placed in highly-visible locations such as end-of-aisle displays and checkouts. Additionally, many supermarkets use attentional strategies to target children, who are more likely to request that their parents buy products located at their eye level. At checkouts, a remarkable 89.5% of foods situated at a child's eye level are unhealthy (RSPH and Slimming World 2019: 7; Horsley et al. 2014).

Many purchases of eye-level goods can be attributed to System 1 thinking. Consumers who are acting on autopilot are liable to satisfice, because satisficing requires less cognitive effort than maximising. Consequently, consumers acting on autopilot will be more likely to select the first product that conforms to their purchase heuristic. When attention is initially drawn to eye-level products, the likelihood of these products being chosen increases.

The People's Supermarket drew attention away from unhealthy foods by placing them on the top and bottom shelves, out of consumers' direct eyelines.¹⁰² Healthier foods were positioned in multiple prime locations, including at consumers' eye levels, in end-of-aisle displays and by the checkout. To further draw attention away from unhealthy foods, healthy products were afforded a greater allocation of shelf space.¹⁰³ When healthy foods are displayed in multiple and more prominent locations, they are easier to find and consumer attention is more readily drawn to them, encouraging purchases. Similarly, when unhealthy foods are displayed in less-prominent locations, they become harder to find and consumers' attention is focused elsewhere.

These insights can also be used to nudge for farm animal welfare. Consumer attention will be drawn away from lower-welfare goods when they are placed above and below eye level, decreasing the likelihood of automatic purchases. Similarly, placing higher-welfare goods at eye level draws consumer attention and increases the likelihood of automatic purchases.

















Many retailers combine discounting and visibility strategies, for instance by placing sale items at eye level or on end-of-aisle displays. These strategies are also combined by online grocery retailers: Figure 10 shows how Tesco use brightly-coloured offers and discounts to attract consumers' attention. The first image is taken from the Tesco groceries homepage, where sales and deals are advertised prominently using bright colours. The second image is taken from a search for 'pizza', and there are two features worth noting: first, four sale items are listed among the first eight results, and so will receive higher levels of consumer attention; the ordering of search results affects which products a consumer sees first. Second, offers are effectively advertised twice, once in yellow and once in red, to maximise the chance of attracting consumer attention.

To summarise, System 1 processes are automatic, fast and require little cognitive engagement. Consumers often view the supermarket shop as a chore, and, when they face time constraints and potentially thousands of choices, may let System 1 processes govern their decision-making. A consequence of this reliance upon heuristic-driven System 1 thinking is that consumer behaviour is vulnerable to influence by supermarket choice architecture. The use of offers plays upon price and value heuristics by encouraging consumers to both impulse

¹⁰² See Appendix IV for images of this, as well as other design features implemented by the People's Supermarket.

¹⁰³ Some of these changes were later reversed to improve consumer experience. The People's Supermarket is located close to the Great Ormond Street children's hospital, and staff informed me that chocolate and sweets were returned to a more visible location by the checkout to assist people in need of comfort food.

Figure 10 – Using Discounts to Grab Online Shoppers’ Attention

 <p>Half Price</p> <p>Snacks and treats ></p>	 <p>Half Price</p> <p>Health and beauty deals ></p>	 <p>£6</p> <p>Plaza Prosecco ></p>	 <p>Half Price</p> <p>Terry's Chocolate Orange ></p>
<p>See all offers ></p>			
 <p>Any 12 for £3</p> <p>Müller yogurts ></p>		 <p>Save £5</p> <p>Selected cleaning and laundry when you spend £15* ></p>	
 <p>Half Price</p>		 <p>Half Price</p>	
<p>Offer</p>  <p>Tesco Stonebaked Thin Four Cheese Pizza 330G</p> <p>Any 3 for £4.00 Offer valid for delivery from 21/11/2019 until 30/12/2020</p> <p>Write a review > Rest of shelf ></p> <p>£ 1.50 £0.46/100g</p> <p>1 <input type="button" value="Add"/></p>	<p>Offer</p>  <p>Tesco Stonebaked Thin Double Pepperoni Pizza 330G</p> <p>Any 3 for £4.00 Offer valid for delivery from 21/11/2019 until 30/12/2020</p> <p>Write a review > Rest of shelf ></p> <p>£ 1.50 £0.46/100g</p> <p>1 <input type="button" value="Add"/></p>	 <p>Goodfella's Stonebaked Thin Margherita Pizza 345G</p> <p>Write a review > Rest of shelf ></p> <p>£ 2.25 £0.65/100g</p> <p>1 <input type="button" value="Add"/></p>	 <p>Goodfella's Stonebaked Thin Pepperoni Pizza 340G</p> <p>Write a review > Rest of shelf ></p> <p>£ 2.25 £0.66/100g</p> <p>1 <input type="button" value="Add"/></p>
<p>Offer</p>  <p>Pizza Express Margherita Pizza 245G</p> <p>Half Price Was £5.00 Now £2.50 Offer valid for delivery from 15/01/2020 until 04/02/2020</p> <p>Write a review > Rest of shelf ></p> <p>£ 2.50 £1.03/100g</p>	<p>Offer</p>  <p>Tesco Stonebaked Margherita Pizza 252G</p> <p>Write a review > Rest of shelf ></p> <p>£ 3.50 £1.39/100g</p>	<p>Offer</p>  <p>Tesco Stonebaked Ham & Pineapple Pizza 375G</p> <p>Any 3 for £4.00 Offer valid for delivery from 21/11/2019 until 30/12/2020</p> <p>Write a review > Rest of shelf ></p> <p>£ 1.50 £0.40/100g</p>	<p>Offer</p>  <p>Tesco Finest Buffalo Mozzarella & Tomato Pizza 420G</p> <p>Write a review > Rest of shelf ></p> <p>£ 4.50 £1.08/100g</p>

buy and consume goods with little regard for broader lifestyle preferences, while shelf and website design can guide consumer attention towards more profitable products.

The People's Supermarket shows that these design features can be used for other purposes, including pro-animal welfare ends. Discounting higher-welfare products and positioning them in highly-visible locations, such as at eye level or on end-of-aisle displays, could nudge consumers to buy these products over low-welfare alternatives.

6. Thinking, Slow: Choice Environments and System 2 Processes

In contrast to fast and automatic System 1 processes, System 2 processes are slow and effortful. These processes are responsible for a range of behaviours, including: focusing on a particular person's voice in a noisy room; looking for a woman with white hair in a crowd; searching your memory to identify a surprising sound; maintaining a faster walking speed than what comes naturally; monitoring the appropriateness of your behaviour in a social situation; counting the occurrences of the letter *a* in a page of text; parking in a narrow space; comparing two washing machines for overall value; filling out a tax form; and checking the validity of a complex logical argument (Kahneman 2012: 22). Because System 2 processes are effortful, they will be disrupted should our attention be drawn away.

Importantly, System 2 processes can be influenced by System 1 processes. For instance, your System 1 processes will orient you towards an unexpected loud noise, allowing System 2 processes to identify the noise and attend to its source. System 1 can also generate the attitudes, impressions, intuitions, intentions and feelings that constitute the building blocks of System 2 decision-making (Kahneman 2012: 24; Norman 2013: 47). This occurs because System 2 is often lazy: exercising self-control – which includes monitoring and controlling intuitive System 1 responses and reactions – is mentally taxing. Consequently, the path of least resistance, in this case uncritically accepting the inputs generated by System 1, is often preferred (Kahneman 2012: 44-46).

The significance of System 1 influence on System 2 processes is magnified when we introduce the System 1 phenomenon of associative priming. Associative processes are hugely useful tools that allow us to make sense of our environments and anticipate likely events based on current situational knowledge. Hearing the sound of barking behind you, for instance, will prime the concept of a dog, meaning that you will not be surprised when you turn around and see a dog (Kahneman 2012: 51).

These associations spread like ripples: an environmental cue primes one concept, which in turn primes another, which may in turn prime a host of other related concepts, albeit less strongly than the initial prime (Kahneman 2012: 52-53). These concepts may be related to beliefs, feelings, attitudes, values and impressions, as well as more concrete objects (Kahneman 2012: 58; Fazio et al. 1986), and can influence reflective decision-making and behaviours.

The influence of primes upon System 2 thinking was illuminated in a study conducted by Moss-Racusin et al. (2012), where science faculties at over 100 universities were asked to rate an application for the position of laboratory manager. The applications were identical in all respects save one: the name. Half of the applications came from a male student (John), while the other half came from a female student (Jennifer).

Despite the applicants being identical in all relevant respects, John was perceived to be more competent and hireable than Jennifer, and was offered a higher starting salary and more mentoring. Notably, this bias persisted regardless of whether the application was assessed by a man or a woman. To frame these results in terms of dual-systems thinking, assessing the applications required reflective System 2 thinking, but System 1 processes may have influenced these assessments. These System 1 processes may have introduced unconscious biases in System 2 thought: if someone holds an association between STEM subjects and males, for instance, they are likely, when asked to think of the archetypal scientist, to imagine a man. This can influence attitudes and behaviours: when hiring someone to work in a laboratory, applicants who resemble an assessor's stereotype of a scientist are likely to be received more favourably.

Because System 2 is often lazy, the reviewers may not have considered what characteristics their archetype of a laboratory manager might possess, and whether they were relevant to the role. Consequently, biases that originated in the knowledge structures primed by associative System 1 processes may have bled into System 2 processes and influenced reflectively-held perceptions and attitudes, and therefore responses to the applications.

A helpful way of understanding these knowledge structures is to frame them in terms of *schema theory*.

a. Schema Theory

The bundle of impressions, attitudes, feelings, beliefs and associations activated by a prime can be referred to as a *schema*. Schemas contain 'the general rules and information necessary for interpreting situations and for guiding [behaviour]' and 'lend a sense of order, structure, and

coherence to social stimuli that otherwise would be complex, unpredictable, and overwhelming' (Norman 2013: 128; Fiske and Taylor 1991: 150); they are heuristics that allow us to understand the world despite our limited informational processing capacities (Fiske and Taylor 1991: 97-98). Schemas influence what we notice, how quickly we perceive objects – with an emphasis on noticing and recalling schema-consistent information¹⁰⁴ – how we interpret what we notice and what we expect from a situation (Flaherty and Mowen 2010: 256; Leung and Morris 2015: 1031). They therefore enable us to 'discriminate among individuals, interpret information, and evaluate others' (Flaherty and Mowen 2010: 255), and thus influence our social perceptions, evaluations and behaviours.

There are numerous types of schema. *Person schemas* allow us to understand particular individuals through their traits and goals. We form unique schemas for friends, colleagues, family members and others, which are populated with beliefs about their characteristics and motivations and colour how we perceive them and their actions. If Rachel perceives Stiofan to be a malicious person, she may think he has acted deliberately when he spills her drink. If she thinks Stiofan is clumsy or good-natured, she may instead be inclined to think that her drink was spilled accidentally. In short, her 'Stiofan schema' colours her perceptions of Stiofan's actions.

Self schemas relate to self-perceptions of characteristics and traits and help us manage information relating to ourselves. The economic and citizen schemas, which are the focus of much of this chapter, are types of self schema. *Event schemas* help us to organise expectations and navigate probable sequences of events in frequently occurring situations, such as ordering food in a restaurant (Fiske and Taylor 1991: 117-18). *Role schemas* help us to manage information and expectations related to the performance of certain roles. There are two types of role: *achieved roles* are those positions intentionally and effortfully attained, such as career-related roles; *ascribed roles* are acquired at birth or without effort, and relate to characteristics such as age, sex and race (Fiske and Taylor 1991: 119). If we recognise that someone is performing a familiar role, our interactions with them will be guided by beliefs about what this role entails. In schematic terms, role-associated stimuli (for example, a firefighter's uniform and a fire engine) are categorised in a way that activates a specific role schema (firefighter), influencing

¹⁰⁴ This is the case only when schemas are well-developed; when schemas are weak, attention tends to be drawn to inconsistent information (Fiske and Taylor 1991: 128).

interactions with the individual (they may be asked to help rescue a cat from a tree) (Fiske and Taylor 1991: 105).¹⁰⁵

Individual schemas are not always active, but are instead primed by features of the environment. There are many forms of priming, but two are especially relevant to this chapter. Schemas can be activated through semantic priming, the use of tasks that engage directly-related schemas (Leung and Morris 2015: 1031). Exposure to relevant concepts can trigger the activation of the corresponding schema. For instance, Carver et al. (1983) primed test subjects with a video of hostile behaviour and word scramble puzzles which featured hostility-themed words. They found that, compared to a control group, participants were more likely to both behave in a hostile manner and attribute hostile attitudes to a fictional character in an ambiguous passage of text, indicating that the video and word scramble tasks had primed a hostile self schema, with implications for participants' beliefs and behaviours.

Alternatively, schemas can be activated through associative priming, the use of images, sounds, smells or tastes that are not directly connected to a schema but may still be reminiscent of it (Leung and Morris 2015: 1032). Research by Nicolas Guéguen (2012) indicates that exposure to the smell of freshly-baked bread can promote prosocial behaviour. Although it is tenuous to claim a direct association between a bakery smell and a prosocial schema, the smell of fresh bread may have elicited a good mood, which in turn primed a prosocial schema (Fiske and Taylor 1991: 146).

Whether a schema influences behaviour depends upon three factors. First, the schema must be *accessible*, meaning it has the potential to be activated (Higgins 1996: 134). Accessibility is connected to recency and frequency. Recently activated schemas are more accessible than those that have not been activated for some time (Carver et al. 1983: 404-7; Fiske and Taylor 1991: 146). When there is a sufficient temporal delay between schema activation and stimulus exposure, however, people tend to fall back on frequently utilised schemas (Higgins et al. 1985: 66; Higgins 1996: 139).

Second, a schema must be *applicable* to a given context. The greater the overlap between the features of a schema and the stimulus, the more applicable the schema is (Higgins 1996: 135). If an activated schema is not relevant to a situation it will not influence behaviour (Macrae and Johnston 1998: 403). Similarly, schemas will only influence behaviour in those contexts where

¹⁰⁵ There may be a degree of overlap between role and self schemas; people performing certain roles may find that their expectations and beliefs regarding those roles influence perceptions of self, at least while the role is being performed.

they are deemed to be applicable: priming an economic schema, for instance, can reduce empathy, but *only in economic contexts*. Beyond this domain, the schema is not perceived as applicable and so will not influence behaviour (Molinsky et al. 2012: 34).

Third, a schema must be deemed *appropriate* (Higgins 1996: 136). There may be inhibitory factors, both internal (goals or morals) and external (contexts and situational cues) that cause people to view schema-driven behaviours as inappropriate (Macrae and Johnston 1998: 403-4). Many people are averse to acting on stereotypes about characteristics such as gender and race, so may resist using active and applicable schemas.¹⁰⁶ Some may even overcompensate in the opposite direction rather than act upon stereotype beliefs (Leung and Morris 2015: 1032). In short, it is only when schemas are accessible, applicable to the given context and deemed appropriate to use that they are likely to be activated.

When they are activated, however, schemas can explain many features of human psychology. Environmental cues that subconsciously activate an applicable stereotype can affect self-perceptions (Fine 2012: 9). Schemas can also influence behaviours. In one of the classic studies in this field, Bargh et al. (1996: 236-37) primed a group of undergraduates with words relating to an elderly stereotype. Unbeknown to the participants, the researchers then timed how long they took to walk down a corridor at the end of the experiment. The researchers found that subjects primed with the elderly stereotype, of which slowness is associated, walked significantly slower than those in the control group, despite words directly relating to slowness not being included in the priming stage of the experiment. Similar effects have been observed when trait schemas, a form of self schema, are primed. Bargh et al. (1996: 235-36) subjected volunteers to a scrambled sentence test intended to prime either a polite schema, a rude schema or nothing. Subjects primed with the rudeness stimuli were more likely to interrupt in a subsequent conversation than those primed with the neutral stimuli, who were in turn more likely to interrupt than those exposed to the polite stimuli.¹⁰⁷

In sum, exposure to schema-related stimuli activates schemas, which can prime people to make decisions and behave in ways congruent with the beliefs, values, expectations and preferences associated with the activated schema. Our environments can thus prime us to think, feel, behave and make decisions in particular ways. This is especially significant when we consider the influence of economic schemas upon consumer decision-making.

¹⁰⁶ People who do this may be using System 2 processes to evaluate the impressions generated by System 1 processes.

¹⁰⁷ See Bargh and Ferguson (2000) for further examples of priming behaviours through schema activation.

b. Economic and Citizen Schemas

The most relevant schema to this thesis is the economic self schema. Multiple studies have suggested that unobtrusive exposure to economic concepts primes ‘economic’ attitudes and behaviours.

Research by Molinsky et al. (2012) indicates that exposure to economic concepts led to dampened feelings of empathy. One common economic prime is money, exposure to which¹⁰⁸ has been claimed to prime numerous behaviours and attitudes, including: self-interested behaviours at the expense of regard for others (Kouchaki et al. 2013); the prioritising of self-interest and preferences for distance from others, and the performance of ‘socially insensitive’ actions (Vohs et al. 2006); and a business-like mindset, which leads people to both adopt an impersonal approach to social interactions and hold a negative view of the expression of emotions in themselves and others (Jiang et al. 2014). Kay et al. (2004) suggest that other objects associated with business, such as boardroom tables and briefcases, are also economic primes that can encourage competitive behaviours. Similarly, Liberman et al. (2004) conclude that calling a Prisoner’s Dilemma game ‘The Wall Street Game’ elicited significantly more competitive behaviours, i.e. defections, than when it was referred to as ‘The Community Game’, indicating that the reference to Wall Street, an institution explicitly associated with business and economics, primed self-interested behaviours. Moreover, Wang et al. (2011) conclude that having an education in economics, which is likely to make economic schemas more accessible, correlates with more positive attitudes towards greed. Research conducted by Frank et al. (1993) supports this idea, finding that economics students, once again a group which is likely to possess accessible economic schemas, are more likely to adopt self-interested behaviours in Prisoner’s Dilemma games than non-economists.¹⁰⁹

¹⁰⁸ Effective exposure techniques have variously made use of scrambled sentence tasks that evoke the concept of money, images of money (Kouchaki et al. 2013), and even Monopoly money (Vohs et al. 2006).

¹⁰⁹ Many priming studies have been criticised for a lack of academic rigour due to their use of small sample sizes. When combined with relatively marginal results, many of these studies lack explanatory power and there is a statistically significant chance that their results will not be replicated in repeat experiments (Lodder et al. 2019). In fact, Chapter 4 of Kahneman’s *Thinking, Fast and Slow*, which focuses on associative priming, has been criticised for referring to studies which lack sufficient rigour (Schimmack et al. 2017). The replication crisis does not undermine the claims made in this chapter, however: as Kathleen Vohs (2015), who has extensively researched money primes, points out, evidence for the existence of economic primes has been found in 165 studies conducted across 18 countries. Additionally, Lodder et al.’s meta-analysis (2019) concludes that a statistically significant money priming effect has been observed in some word scramble and money handling experiments.

In short, these studies suggest that exposure to economic concepts can prime an economic schema which promotes self-interested values and attitudes while dampening expressions of emotion, empathy and other-regarding values (Wang et al. 2011: 645-46; Molinsky et al. 2012: 28). Additionally, the attitudes and values primed by economic stimuli appear to bleed into System 2 decision-making in games such as the Prisoner's Dilemma, where participants with a primed economic schema prioritised self-interest and a calculative mindset over prosocial behaviours (Vohs et al. 2008: 209-10).

The economic schema, primed by exposure to economic concepts, contrasts with the *citizen schema*, which, when primed, prompts prosocial behaviours that are more strongly influenced by social values and consideration of what is best for the wider community (Conner 2004: 28). A range of studies lend credence to the existence of a prosocial citizen schema and shed light on its nature. Abbate et al. (2013a, 2013b) propose that exposing participants to help-related words in a scrambled sentence task primed prosocial schemas that led people to more-readily offer assistance and also to donate more when requested. Similarly, Macrae and Johnston (1998: 408) contend that scrambled sentence tasks can prompt helpful behaviour, provided the costs of helping were not too great. Research by Van Tongeren et al. (2018) indicates that exposure to superhero images primes a prosocial schema (superheroes are generally regarded as paradigms of altruism and prosociality), which led to higher levels of both self-reported intentions to help others and actual helping. Greitemeyer (2009) noted that priming test subjects with songs featuring prosocial lyrics increased feelings of empathy and helping behaviours. Additionally, the aforementioned Liberman et al. study (2004) concluded that referring to a Prisoner's Dilemma game as 'The Community Game' encouraged more cooperative behaviour than referring to it as 'The Wall Street Game'.¹¹⁰ In short, exposure to prosocial words, lyrics and paragons of prosociality can prime prosocial (citizen) schemas, which lead to more other-regarding and prosocial behaviours.

Due to their respective emphases on self- and other-regarding behaviours, economic and citizen schemas will often conflict and thus elicit different and opposing judgements and behaviours in response to similar objects. Mark Sagoff (2001: 468) illustrates: 'I speed on the highway; yet I

¹¹⁰ Ideally, this study would have had a neutral third category that did not attempt to prime a citizen or economic schema. Without this, there are three possible explanations for the results: (1) exposure to 'The Wall Street Game' primes an economic schema, leading to increased competitive behaviour, but exposure to 'The Community Game' does not prime a citizen schema; (2) exposure to 'The Wall Street Game' does not prime an economic schema, but exposure to 'The Community Game' primes a citizen schema, leading to increased cooperative behaviour; (3) both schemas are activated, leading to increased competitive or prosocial behaviour depending on which prime the subjects were exposed to.

want the police to enforce laws against speeding. I used to buy mixers in returnable bottles – but who can bother to return them? I buy only disposables now, but, to soothe my conscience, I urge my state senator to outlaw one-way containers’. Similarly, we may ‘protest to try to block Wal-Mart from coming to our town and send letters demanding human rights in China and Myanmar, then turn around and buy cheap sweatshop-produced trinkets at these big box retail stores’ (Conner 2004: 28-29). These examples may represent violations of the continuity axiom: we prefer cheap clothes to workers’ rights when in a superstore, but our preferences are reversed in a community context.

c. Schemas and Preferences for Farm Animal Welfare

So far, we have seen that exposure to relevant stimuli can prime schemas, which are knowledge structures containing impressions, attitudes, feelings, beliefs and associations. Primed schemas can feed into System 2 decision-making. We have also seen that exposure to money and other economic concepts can prime self-interested values and ways of thinking, while exposure to prosocial concepts can activate an other-regarding citizen schema.

Modern supermarkets are filled with money-based cues including prices, deals, bargains and commodities, and are therefore likely to prime economic schemas. Consequently, shoppers may be more self-interested and less attentive to other-regarding concerns in supermarkets. In Chapter 1 (p.23), we saw that there are a range of reasons for consumers to hold preferences for farm animal welfare. Some preferences are motivated by self-interest, often connected to a belief that farm animal welfare contributes to a range of other desirable product attributes. Greater attention may be paid to self-interested preferences for farm animal welfare when consumer schemas are active. Other consumers hold altruistic preferences for farm animal welfare, however, and, regardless of whether these preferences are motivated by ethical concern or mere liking for some types of animal, their expression is likely to be dampened in economic choice environments.

This hypothesis is supported by a range of survey and focus group results. When asked to consider why they would buy higher-welfare products, participants in one focus group stated that they were generally only willing to pay more for higher-welfare goods when these goods produced direct benefits, such as healthfulness, for themselves (Harper and Henson 2001: 11). Being asked to view animal welfare as a product attribute led to its commodification and activated an economic schema, which caused participants to give priority to their self-interested values.

These findings are corroborated by the 2007 Eurobarometer survey, which asked respondents to choose three of nine possible reasons for purchasing high-welfare goods. Health concerns dominate, with 15% of UK respondents choosing ‘they come from healthier animals’¹¹¹ as their first choice, and another 15% holding that high-welfare goods are healthier for consumers. A further 11% stated that high welfare entailed higher quality and 10% equated high welfare with better taste. In contrast, the two options that appear to be most indicative of an intrinsic valuation of animal welfare, ‘they come from happier animals’ and ‘they help farmers treat their animals better’, were chosen by 13% and 12% of respondents respectively (Eurobarometer 2007: 78). The survey therefore finds that less than 40% of high-welfare purchases¹¹² are primarily motivated by concern for animal welfare *for the sakes of the animals themselves*: instead, most high-welfare purchases appear to be primarily motivated by an instrumental valuation of animal welfare. When asked to think of animal welfare in relation to animal products, thus priming an economic schema, initial responses appear to be motivated primarily by self-interested preferences for high-quality, safe and healthy food.

In contrast, stimuli that prime citizen schemas (or at least do not prime economic schemas) appear to encourage more altruistic and other-regarding attitudes towards farm animal welfare. The 2016 Eurobarometer survey asked respondents to choose two of five statements that best reflected their understanding of farm animal welfare. Among UK respondents, the statement ‘[animal welfare] contributes to better quality animal products’ received the lowest number of selections, with just 14%. In contrast, ‘[animal welfare] concerns the way farmed animals are treated, providing them with a better quality of life’, and ‘[animal welfare] refers to the duty to respect all animals’ were selected by 42% and 40% of respondents respectively (Eurobarometer 2016: 7). Respondents were therefore more likely to perceive farm animal welfare as being of intrinsic rather than instrumental value when discussing it in a setting that did not prime economic schemas through associations with commodities.

¹¹¹ This option is not particularly helpful for our purposes, as the contribution of higher-welfare goods to animal health could be viewed as desirable for both instrumental and intrinsic reasons. For instrumental reasons, it may be believed that healthier animals produce products that are higher-quality, more nutritious or simply better-tasting, although such thoughts may be better-expressed by selecting these options in the survey. Alternatively, greater animal health may be viewed as desirable for altruistic reasons, for the sake of the animal itself.

¹¹² This figure is composed of the 13% who chose ‘they come from happier animals’, the 12% who chose ‘they help farmers treat their animals better’, and also the 15% who chose ‘they come from healthier animals’, charitably assuming that all did so out of concern for animal welfare rather than a belief that healthier animals make better produce. Although the limitations of the survey mean we cannot ascertain precisely how many, it is likely that some proportion of that 15% valued animal health instrumentally.

In short, priming and schema theory suggest that choice environments can influence the weights given to criteria used in decision-making, with implications for reflective System 2 thinking. System 2 processes are used to compare goods according to multiple criteria; when exposure to economic primes influences the weights given to these criteria, the subsequent assessments will also be affected. Suppose you are looking to buy a new washing machine, and are concerned with two attributes: price and energy efficiency. You have identified two candidates: *a* is cheaper than *b*, but less energy efficient. Reading a leaflet about a social cause that you support is likely to prime a prosocial citizen schema, and prompt you to afford greater significance to prosocial values such as environmental concern. Reading a leaflet about the successes of a famous entrepreneur may instead prime a consumer schema, causing you to give greater weight to price concerns.

Different choice environments appear to evince divergent and context-sensitive preference sets: what I prefer when immersed in economic stimuli may not be what I prefer when surrounded by prosocial stimuli. When the citizen schema is activated, other-regarding values dominate, encouraging people to act on preferences for workers' rights in China and Myanmar over cheap clothes, or preferences for higher animal welfare over cheaper low-welfare produce. These preferences may be reversed, however, when the self-interested economic schema is activated in an economic choice environment such as a supermarket. The context-sensitive nature of some preferences may therefore cause violations of the continuity axiom.

7. System 2 Nudges

Policymakers can use nudges to influence the building blocks that are often used unreflectively in System 2 processes. Drawing from the RSPH and Slimming World report, the People's Supermarket introduced 'nudge points', bright yellow signs displayed in prominent positions that were intended to catch consumers' attention.¹¹³ The People's Supermarket's nudge points drew consumers out of automatic modes of thinking and encouraged them to be more mindful of preferences for healthy living at points of purchase by providing relevant health information and encouraging healthy behaviours. The content of the nudge points is intended to activate schemas associated with healthy living and thus encourage consumers to consider values associated with health in their decision-making. This strategy is supplemented with the provision of in-store nutritional guidance.

¹¹³ See Appendix IV for images.

A similar policy could also be used to encourage consumers to consider a range of ways of valuing farm animal welfare at points of purchase. Information could be provided to promote the consideration of animal welfare at points of purchase, inform consumers about animal welfare issues and empower consumers to make welfare-motivated decisions. An animal welfare nudge point would therefore seek to prime citizen schemas – and thus encourage consumers to consider prosocial values, namely concern for farm animal welfare – in economic choice environments.

Policymakers must, however, be mindful of how consumers engage with farm animal welfare when deciding what kinds of information should be presented on nudge points. Responses to the meat paradox – discussed in Chapter 1 (p.30) – indicate that many consumers will avoid or ignore information that makes them uncomfortable; consequently, information about low-welfare animal agriculture, or those practices used in higher-welfare animal agriculture that elicit discomfort, will not be effective drivers of higher-welfare purchases.

Instead, effective information provision will focus on the idea of happy animals, tapping into the emotional nature of consumer engagement with farm animal welfare. This entails an emphasis on the positive aspects of higher-welfare production, such as access to outdoor spaces and freedom to exercise natural behaviours. This information could be linked to a welfare labelling scheme of the kind outlined in the previous chapter (p.145).

Alternatively, information provision could focus on social norms, a key pillar of the influence literature, which was overlooked in the RSPH and Slimming World report.

a. Social Norms

Social norms are implicit in human interactions, and specify which actions are deemed proper or correct, and which are improper or incorrect (Coleman 1994: 242-9; Young 2015: 360). They are wide-ranging in scope, covering everything from how to answer the phone to proscriptions against littering to appropriate behaviour in a zoo. Norms play an important role in regulating social behaviour and may therefore be a valuable policy instrument.

The use of social norms as an influence tool relies upon the idea that people will do and buy as others do and buy. Social norms affect both System 1 and System 2 processes: television executives have long known that canned laughter causes viewers to laugh longer and more often and rate the material as funnier, while advertisers' claims that a product is best-selling may give us reason to trust the quality of their goods (Cialdini 2007: 115-17). Social norms also offer an

explanation for the bystander effect, where a group of otherwise good people stand idle as an atrocity is committed in front of them, simply because everyone else is doing the same (Cialdini 2007: 128-36). In short, we often look to others to guide our actions, and policymakers can use this tendency to encourage prosocial behaviours by drawing attention to what other people are doing.

The viability of social norm policies as a tool for promoting farm animal welfare, however, is dependent upon the existence of prosocial norms; if people generally do not behave prosocially, there will be no social norms to appeal to. Evolutionary psychology provides a theoretical basis for the idea that co-operative and reciprocal norms are an innate part of human psychology that may have contributed to our species' survival in times of scarcity (Saad 2011: 121-22; Singer 2011: 23-53). Robert Trivers (1971) proposes that natural selection promotes non-kin reciprocal norms when members of a species repeatedly interact with one other and have a long lifespan (to allow debts to be repaid), and live in small, stable social groups. These norms improve a being's (and a species') chances of survival by nurturing friendships, establishing mutually-beneficial relationships with strangers and creating stable norm-following societies by ostracising or otherwise punishing norm-breakers.

Experiments by behavioural economists lend further credence to the existence of co-operative and reciprocal norms. Kahneman et al. (1986) ran a series of experiments where participants were asked to play the *dictator game*. In this game, one player is asked to propose a division of a sum of money between themselves and a partner, which the partner must accept. In Kahneman's experiment, proposers had two choices: they could split \$20 evenly or keep \$18 for themselves and leave their partner with \$2. As the game was played anonymously, meaning that the influence of reputational concerns can be discounted, a purely self-interested player would choose the uneven split.

Players were generally not purely self-interested, however: despite the impossibility of having their offer rejected, 76% of proposers chose to split the money evenly (Kahneman et al. 1986: S291). Kahneman's findings are generally replicated in the literature: one meta-analysis found that 64% of dictators make a positive offer, and 30% offer at least half the endowment (Engel 2011: 589).¹¹⁴

The existence of co-operative and altruistic social norms cannot be definitively divined from these results, however. The learning hypothesis, which states that people might behave

¹¹⁴ Henrich et al. (2001: 75) found that positive offers in dictator games were common in a range of societies.

generously simply because they do not fully understand the game, provides a plausible alternative explanation. The learning hypothesis predicts that the dominant strategy of non-cooperation will be learned with experience, and is lent credence by two pieces of evidence. First, in Kahneman et al.'s dictator game (1986), commerce students were less likely to offer an even split (63%) than psychology students (80%). This indicates that economics training, which often includes a study of game theory, encourages economically 'optimal' behaviours, a finding replicated in the literature (see Marwell and Ames' 1981 paper 'Economists Free Ride, Does Anyone Else?' for instance). Those without economics training are perhaps more likely to instead learn the optimal strategy by repeatedly playing the game.

Results from non-punishment treatments of public good games¹¹⁵ could be interpreted as supporting the learning hypothesis. In numerous studies of multiple-round public good games, average contributions have been found to decay over time, converging on total free-riding by the game's end (Ledyard 1995: 121; Fehr and Schmidt 1999: 844-45; Fehr and Gächter 2000: 985-6).¹¹⁶ The learning hypothesis appears to offer an explanation for these behaviours: as people become more familiar with public good games, they learn economically-optimal strategies and accordingly become less co-operative.

Further public good games conducted by James Andreoni (1988b) cast doubt upon the learning hypothesis' explanatory power, however. Andreoni's study began with a basic ten-round public good game, with the composition of four-person groups remaining fixed throughout. Although participants knew that their partners would not change throughout the game, they did not know who their partners were, thus eliminating reputational motivations for co-operation. Andreoni found that, as with the other studies, levels of co-operation decayed over the ten rounds of the public good game. After the tenth round, however, he informed his subjects that the game would restart, with participants continuing to play in the same groups as before.¹¹⁷

If the learning hypothesis explains decay in public good games, we would anticipate the new game to feature low levels of co-operation from the outset, as participants would by now have arrived at a strategy of non-cooperation. Instead, participants largely reverted to their levels of co-operation in the corresponding rounds of the first game, with a significant proportion even increasing their initial contributions. This finding surely invalidates the learning hypothesis,

¹¹⁵ See Chapter 5 (p.102) for more on public good games.

¹¹⁶ This decay is evident both when subjects are aware of how many rounds there are in the game and when they are not (Isaac et al. 1984; Isaac et al. 1985).

¹¹⁷ Due to budget constraints, only three rounds of the new game were played, although players were told they would be playing ten.

which predicts that experienced players will start new games at low or non-existent levels of co-operation. To further reject the learning hypothesis, Marwell and Ames (1981: 302-3) invited participants from an earlier public good experiment to play again, and found that levels of co-operation were largely unchanged and in line with those of inexperienced players, at about 47% of resources.

With the learning hypothesis rejected, it is appropriate to return to social norm theory. Although co-operative norms can explain why people are so often generous in dictator games, and why initial contributions in public good games tend to be significant, they do not obviously explain why co-operation decays in multiple-round public good games. Public good games conducted by Keser and van Winden (2000) may provide an answer by revealing the nature of co-operative norms. The researchers conducted a 25-round public good experiment (rather than the usual ten) which split participants into two conditions: those in the partner condition played in the same four-person group for the entire game, while people in the stranger condition were randomly allocated to four-person groups every round. After each round, participants were told how much their group had contributed in that round, but not how other groups, or any particular individual, had behaved. The authors found that, across all rounds, partners contributed more to the public good than strangers. In addition, there was no decay in the contribution of partners until the last rounds of the game, in contrast to the strangers condition, where decay was evident throughout.¹¹⁸

Keser and van Winden propose that prosocial norms are governed by two principles, which together compose the theory of *conditional co-operation* (Keser and van Winden 2000; Dawes and Thaler 1988: 191). The first of these principles is *future-oriented behaviour*, which means that behaviour is influenced by expectations of future interactions. Future-oriented behaviour explains why co-operation dropped off towards the end of the game, as participants either believed it was no longer worth continuing to invest in their relationship with other group members, or believed that other players would feel this way.¹¹⁹ This also explains why there were lower levels of co-operation in the strangers condition: participants could not form expectations

¹¹⁸ This result conflicts with Andreoni (1988b), who found decay set in early in his study, rather than towards the end. One explanation for this is that Andreoni's experiment used games of only ten rounds, in contrast to the 25-round games used by Keser and van Winden. In Keser and van Winden's study, decay appeared to set in in the last six rounds, while in Andreoni's study decay was most evident in the last seven rounds. It may simply be that using fewer rounds gave the impression of constant decay in Andreoni's study, when it is in fact a purely end-game phenomenon in partner conditions.

¹¹⁹ In keeping with Isaac et al. (1994: 30), who found that having more rounds in a game correlated with a slower rate of decay, Keser and van Winden hold that people are myopic, looking only a few rounds ahead instead of using a backwards induction strategy. This leads to decay setting in only when the game's end comes into view, rather than a total absence of co-operation.

of other players' future behaviour and did not have a relationship with other group members to invest in, and so had less incentive to co-operate at the beginning.

The second element of conditional co-operation is *reactive behaviour*, where people are influenced by the average group behaviour in the previous round. In both conditions of Keser and van Winden's experiment, participants tended to decrease their contributions if they scored above the group average in the previous round and increase them if below.¹²⁰ In short, the decay witnessed by Andreoni may not be as prevalent as first believed, and can be explained by players putting less value on maintaining their relationship with the rest of the group as the endgame comes into view, and other players responding in kind. These lower levels of co-operation begot even lower levels, contributing to an eradication of co-operative norms by the game's end.

To summarise, experiments using several types of competitive games have produced results which cohere with research from a range of other disciplines to support the existence of social norms that favour co-operative and altruistic behaviours. Alternative explanations for these findings, such as self-interest and the learning hypothesis, can be dismissed. By considering the nature of prosocial norms, we see that co-operation is often not unconditional, but instead dependent upon beliefs and expectations relating to the contributions of others.

i. Social Norms as Policy Instruments

Prosocial norms have the potential to be effective policy tools for improving farm animal welfare. If policymakers can prime citizen schemas, they may be able to promote norm following: research in this area indicates that conformity primes, such as reminders of social norms, can induce people to conform to group norms (Epley and Gilovich 1999). By nudging people to conform to these norms, the rate of norm following increases, which in turn could encourage others to become norm-followers. Social norm nudges thus have the potential to create a virtuous cycle of norm-following by encouraging people to afford greater significance to other-regarding values.

The effectiveness of a social norms policy for encouraging purchases of higher-welfare goods will, however, be affected by whether pro-animal welfare norms exist. Research by Bennett and Blaney (2002) provides evidence for these norms: when told that there existed a social consensus against a given animal welfare issue, participants stated both stronger opposition and a higher

¹²⁰ Because similar levels of reactive behaviour were observed in both conditions, we can rule out the possibility that strategic behaviour, intended to encourage others to contribute more, motivates this phenomenon.

willingness to pay to resolve it. This indicates that there are descriptive norms pertaining to farm animal welfare, which, by informing people about what others feel and do, influence attitudes and behaviours.

The existence of pro-welfare norms sheds light on several focus group studies which examined barriers to purchasing higher-welfare products. These studies found that many people chose not to buy higher-welfare goods because they believed that, in isolation, their purchases had little impact on farm animal welfare (Brook Lyndhurst 2010: 33-34; Schröder and McEachern 2004: 174; Harper and Henson 2001: 10; Harper and Henson 1999: 17-19). These consumers were unwilling to take a stand on animal welfare issues when they believed that they were acting on their own, and so appear to be conditional co-operators. Consequently, if they could be reassured that they are acting as part of a group, these consumers may be more willing to buy higher-welfare goods. Social norm nudges may therefore have a dual effect: they can both prime prosocial schemas that indirectly affect System 2 processes, and also inform consumers about the extent of adherence to social norms, which can directly affect reflective System 2 thinking.

The insight that some consumers are conditional co-operators in matters of farm animal welfare can provide the basis of a policy response. Inspiration for such a policy can be drawn from the UK's BIT, which was tasked with increasing responsiveness to letters requesting the payment of late taxes. Although most people eventually pay, speeding up the process can advance millions of pounds to the government.

Drawing on the work of psychologist Robert Cialdini, the BIT drafted three letters. The control was a simple letter explaining how the tax bill could be paid, which elicited payment in 33.6% of recipients over the following twenty-three days. This letter obeyed the behavioural economic principle that if you want someone to behave in a certain way, you should make it easy (Halpern 2015: 65; Thaler 2016: 313); the letter had been stripped of unnecessary detail and simplified.¹²¹ The first treatment added the (true) sentence 'nine out of ten taxpayers pay on time'. This treatment was inspired by Cialdini's recognition of the power of social norms, that people are more likely to do or buy something if they are made aware that others are doing or buying the same thing (Cialdini 2007: 116). This social norm letter was more successful than the control, eliciting payment from 35.1% of recipients.¹²² The second treatment added the sentence 'most

¹²¹ In this respect the control was not actually a control. Unfortunately, the experimenters did not release any data about responsiveness to the original, more technical letters.

¹²² Although the difference between the two letters is just 1.5%, considering that the policy was essentially costless to implement and brought forwards £9 million of revenue in 23 days, it can be deemed successful (Thaler 2016: 337).

people in *your local area* pay their tax on time’, which drew on Cialdini’s insight that the influence of social norms is enhanced when the other person or people are similar to us: we have greater liking for people who are like us, and are more likely to want to do as they do (Cialdini 2007: 140-42). This treatment was even more effective, with 35.9% of recipients of the local norms letter paying their tax (Halpern 2015: 112-15). In short, when people are conditional co-operators, informing them that other people – especially people like themselves – are norm-followers may encourage them to behave in a similar way.

Policymakers can draw from the BIT experiment to design nudges that encourage conditional co-operators to purchase higher-welfare goods. The BIT’s experiment indicates what kinds of information might be effective. The control in the tax experiment was a letter explaining how the tax bill could be settled, i.e. it informed people how they could perform the desired action. An effective labelling scheme that helps consumers identify higher-welfare goods would play an equivalent role in animal product markets.

The BIT’s first treatment provided information about other people’s behaviour. In an animal welfare context, this may involve providing information about other consumers’ purchases of higher-welfare goods, either by informing consumers about the market share of higher-welfare products and welfare labelling schemes or, in those cases where market share is low, providing information about growth trends in higher-welfare markets. This information could be displayed on nudge points, supermarket displays, or even on the products themselves, with taglines such as ‘one in two people buy free-range meat or eggs for environmental / animal welfare reasons’ (Ethical Consumer 2018).¹²³ Similar information could be put on fish products and meat substitutes: ‘one in three people are switching out meat for vegetarian and pescatarian alternatives’ (Harmston 2017).¹²⁴

Drawing on the tax experiment’s second treatment, which provided information about how *similar* people behaved, information about higher-welfare purchases could be tailored to the local area: supermarket displays could display the percentage of egg-buyers in their store who buy free-range or organic over caged eggs, for instance. Compared to the other options, however, this would likely come at a significantly higher cost to retailers, who would have to monitor and analyse the relevant data. Given that the local norms treatment was only

¹²³ Ideally, these taglines would be more precise and simpler in presentation, but due to the nature of existing research the ones proposed in this chapter are somewhat mealy-mouthed.

¹²⁴ Admittedly, more research needs to be conducted to determine the level of norm-following that will encourage conditional co-operators to comply; if rates of norm-adherence are low, consumers could be deterred from buying higher-welfare products. Framing effects may also influence consumer responses.

marginally more successful than the social norms treatment in the tax letter experiment, it is likely to be more cost-effective to use social norms rather than local norms to nudge for farm animal welfare.

There is scope to be sceptical about social norm policies in animal product markets, however. First, we may be doubtful that the policy will have much impact: in the BIT experiment, for instance, the difference between the rate of tax repayment in the control and local norms treatment was just 2.3%. Although we do not know the extent to which the so-called control represented an improvement over the pre-experiment letter, which was more technical and thus harder to understand, we might expect a social norms policy to have a relatively small effect upon consumer behaviour. In terms of animal welfare, however, these small percentages represent meaningful improvements to the lives of hundreds of thousands of animals: a 2.3% shift towards higher-welfare egg production could improve the lives of over 900,000 laying hens in the UK (DEFRA 2019a: 6), and even more in the long run.

Second, it is not clear that this use of social norms constitutes a viable *policy*. There are numerous steps, several of which could be resource-intensive, that a government would need to take to ensure the policy's success. To ensure the provision of accurate information, governments would need to conduct research into consumer behaviour in animal product markets. This research would also need to be regularly updated in order to remain relevant. The Food Standards Agency currently conducts a yearly *Food and You* survey that collects data on consumer attitudes and concerns relating to food, and so questions about animal welfare purchase habits could be included here. Legislation would also be necessary to ensure that packagers or retailers display survey results on products or product displays, and oversight and enforcement would be required to ensure that legislation is adhered to. It may therefore be more appropriate for social norms nudges to be voluntarily enacted by the private sector. In particular, social norms policies may be particularly advantageous for producers of higher-welfare goods and welfare labels: these nudges could be added to their higher-welfare products as a means of increasing sales.

8. Summary

In this chapter, I have argued that, due to the influence of choice environments and choice architecture on preferences for farm animal welfare, consumer behaviour in animal product markets frequently violates the continuity axiom. I introduced Kahneman's dual-process theory of decision-making, and argued that environmental cues in supermarkets can influence both

System 1 and System 2 decision-making. Attention plays a significant role in automatic decision-making, and so products placed at eye level are more likely to be bought than others. Additionally, automatic decision-making typically relies upon heuristics such as price and value concerns, which can be triggered by offers to drive impulse purchases that consumers would not have made upon reflection.

The role of environmental cues in System 2 decision-making is more subtle, and is best appreciated with an understanding of priming. Exposure to relevant cues triggers bundles of associated concepts, beliefs, values and attitudes, which I describe as schemas. Active schemas influence consumers' priorities, which can affect their reflective decision-making. Economic schemas are primed by exposure to economic concepts such as money, and can encourage self-interested behaviours at the expense of other-regarding ones. Supermarkets are filled with economic primes, such as commodities, prices and discounts, and so likely activate economic schemas. This means that supermarket choice environments could be suppressing ethical concern for farm animal welfare in reflective decision-making.

Behavioural nudges can encourage consumer purchases of higher-welfare goods in two ways: first, they can influence automatic behaviours by directing consumer attention and shopping heuristics towards higher-welfare goods; second, they can introduce prosocial values into reflective decision-making. This second objective is achieved through information provision, although care must be taken about what information is provided. The emotional nature of consumer engagement with animal welfare suggests that a focus on the positive aspects of higher-welfare animal agriculture – which could also be linked to a welfare labelling programme – could achieve this aim. Additionally, appeals to social norms may prime other-regarding citizen schemas that encourage prosocial behaviours, while information about the extent of these norms can also influence reflective decision-making.

This chapter, along with the previous three, has demonstrated how market failures suppress the expression of concern for farm animal welfare. We have also seen how a range of policies can address these issues. What remains to be seen is how these responses can be drawn together to construct a coherent policy framework for farm animal welfare.

Chapter 8

Conclusion: Building an Economic Farm Animal Welfare Policy Framework

In this Chapter

- An outline of John McInerney's (2004) three-tier policy framework.
- Three criticisms of this framework.
- A proposal to move away from a preference satisfaction approach to economic farm animal welfare policy and towards a values-based approach.
- A proposal for a three-tier policy framework based on public values.

1. Introduction

I began this thesis by outlining two key concepts in welfare economics, utility understood as preference satisfaction and efficiency. I showed that these concepts come together to justify improvements in farm animal welfare only where there is sufficient consumer demand, and I argued that, in the UK, there appears to be substantial willingness to pay for farm animal welfare. An effective market would be able to satisfy consumer preferences for farm animal welfare, but I contended that markets in animal products are instead characterised by market failures.

I proceeded to discuss the four forms of market failure that have significant bearing on the provision of farm animal welfare. Animal agriculture produces numerous externalities, including the satisfaction and dissatisfaction felt by people as a consequence of facts about farm animal welfare in society. Farm animal welfare also shares much in common with public goods, as people can enjoy its benefits without contributing to its provision. In addition, farm animal welfare is a credence attribute, meaning that markets are vulnerable to the effects of information asymmetries. I also proposed that consumer behaviour can be influenced by environmental cues, which, given that farm animal welfare is typically engaged with in economic environments, may hinder the expression of altruistic preferences for animal welfare.

I also considered how policy could be used to ameliorate the impact of market failures. Each type of market failure warrants different policy responses. While, for instance, public goods may

require direct government provision, information asymmetries are best dealt with by introducing effective product labelling.

In this chapter, I draw upon the work of the preceding seven to develop a framework for economic farm animal welfare policy. I propose that policy should be determined by facts about animal welfare in society and how they relate to public values. I begin with a discussion of John McInerney's (2004) three-tier policy framework. Although McInerney's framework is helpful, we will see that there are three reasons why, in light of this thesis' analysis, his interpretation is ill-suited to my ends: first, McInerney relies upon an interpretation of social utility that is untenable; second, the instability of many people's preferences for farm animal welfare limits the usefulness of efficiency as a policy goal; third, preference instability raises problems for preference satisfaction theories of welfare. In response to these criticisms, I propose that public values, rather than preference satisfaction, provide a more plausible basis for economic farm animal welfare policy, and argue that the aim of economic farm animal welfare policy should be to afford appropriate representation to different ways of valuing animal welfare. Accordingly, I develop my own three-tier framework that places public values at the heart of economic farm animal welfare policy.

2. The Need for a Policy Framework

Some writers (see Norwood and Lusk 2011: 310-11 and Sumner et al. 2011: 249 for example) do not explicitly recognise that policymakers' responsibilities towards farm animal welfare may change according to its current level of provision. This is an inadequate stance to take, however, and its flaws are best exposed by considering the provision of public goods. The provision of public goods is often regarded as a government responsibility, but it would be absurd to suggest that governments are obliged to provide public goods *ad infinitum*. Instead, they should do so insofar as the felt benefits outweigh the costs of provision (McInerney 2004: 45).

In relation to farm animal welfare, this means that – from an economic policy standpoint – it is inappropriate to treat animal welfare equally at all levels of provision. There will exist some level of supply at which the overwhelming majority of the public will be largely satisfied with the provision of farm animal welfare and public demand for further welfare improvements will be muted. It might therefore be inappropriate for governments to directly supply extremely high levels of animal welfare, as to do so would impose costs upon the public but provide few benefits. Consequently, policy ought to be influenced by both facts about animal welfare in society and public attitudes towards farm animal welfare. As farm animal welfare improves in a society,

public preferences for welfare are likely to be satisfied and, from a policy perspective, there is a less urgent need to make further welfare improvements.

For this reason, it may also be useful for policymakers to change the economic classification of farm animal welfare according to facts about its provision in society. This claim lies at the heart of McNerney's (2004) three-tier policy classification. At the lowest levels of provision, McNerney argues, we should treat farm animal welfare as a *public good*. The lowest levels of animal welfare are characterised by neglect and uncaring management and will be universally regarded as unacceptable to society – as stated in Chapter 5, these husbandry methods, defined in terms of their almost universal public opposition, can be described as *cruel practices*.¹²⁵ Although McNerney does not make this claim, it is implicit in his thought that, because cruel practice is almost universally regarded as unacceptable, people will be willing to pay to prohibit it. For this reason, policymakers possess a mandate for direct action: as, by definition, almost everyone prefers there to be no animal cruelty in society, and people are willing to pay to prohibit cruel practice, the costs of government action to directly protect farm animals from cruel practice will be outweighed by the benefits.

Although we have seen that farm animal welfare is not best understood as a public good in a technical sense, from a policy perspective it may be appropriate to treat it as one at this level of provision. Policymakers are justified to take direct action against animal cruelty as doing so satisfies nearly everyone's preference to live in a society free of animal cruelty and markets may not be an effective way of satisfying these preferences. The direct provision of a good is a hallmark of public good policy, and so the use of public good policies is appropriate here too (McNerney 2004: 44; Fearing and Matheny 2007: 164).

It would be inappropriate, however, to treat farm animal welfare as a public good in contexts where animal cruelty does not exist. Because cruelty is defined in terms of almost universal public opposition to a practice or set of practices, the provision of farm animal welfare just above this level will not be met with the same degree of public opposition: there will be significant groups of people whose preferences for farm animal welfare will be satisfied at the point where cruelty is eradicated. While policymakers possess a mandate for direct intervention against animal cruelty, their mandate for intervention above this level is weaker: the costs of direct interventions will be borne by everyone in society but significant groups of people will not derive any benefit.

¹²⁵ It is a policymaker's responsibility to create a social definition of cruel practice by drawing together the range of public opinion.

Nevertheless, policymakers might think that the further provision of farm animal welfare beyond levels of cruel practice is desirable. This could be because of beliefs held by opinion formers or a growing public interest in farm animal welfare, or the belief that a majority of people have preferences for farm animal welfare that have not yet been satisfied. For these reasons, the further provision of farm animal welfare could be viewed as being of overall benefit to society, despite the apathy felt by significant groups of people.

At these levels, policymakers may therefore view farm animal welfare as a *merit good*, a good that is beneficial to society beyond its aggregate demand (McInerney 2004: 44; Fearing and Matheny 2007: 164; Musgrave 1959: 13). Because the provision of merit goods is to society's benefit even if some people do not hold preferences for them, and because policymakers lack a mandate for directly providing farm animal welfare at merit good levels, the aim of policy here is to *encourage* people to support higher farm animal welfare and deter them from producing and consuming lower-welfare products. These policy aims have much in common with externality policy aims, and for good reason: as with merit goods, goods that bear positive externalities provide benefits to society beyond their aggregate demand.

At the highest levels of farm animal welfare, McInerney contends, we will find that most preferences for farm animal welfare will be satisfied and further government intervention, either through public good or externality policies, will come at an overall cost to society. Public good and externality policies are costly, and, when most people's preferences for farm animal welfare have been satisfied, the benefits of further intervention are limited. Nevertheless, some people will still prefer higher levels of farm animal welfare. McInerney states that there is no role for policymakers here, as farm animal welfare should be treated as a *private good*: if people continue to prefer higher levels of farm animal welfare, they should act on these preferences in the market (McInerney 2004: 46; Vetter et al. 2014: 120).

In short, McInerney's policy framework proposes that policymakers' responsibilities towards farm animal welfare are influenced by both facts about animal welfare in society and consumer preferences for animal welfare. Where practice is deemed to be cruel, public opposition – and willingness to pay to eradicate the practice – will be almost universal, and policymakers can use public good policies to protect farm animal welfare. Just above the level of cruel practice, significant numbers of people will be satisfied with the provision of farm animal welfare, but further provision may nonetheless be of overall benefit to society. Policymakers lack a mandate to directly protect farm animal welfare at this level but can use externality policies to encourage markets to supply higher levels of farm animal welfare. At the highest levels of provision,

however, further market intervention by policymakers comes at an overall cost to society, and further contributions to farm animal welfare should be left to consumers.

The analysis I have conducted in this thesis reveals three flaws in the foundations of McNerney's framework, however. First, his justification of merit good interventions relies upon an account of social utility which, in Chapter 1 (p.37), I argued is untenable. Second, the instability of many people's preferences for farm animal welfare, as discussed in the previous chapter, creates problems for policymakers hoping to ground economic farm animal welfare policy in theories of efficiency. Third, McNerney incorporates preference satisfaction at the heart of his framework, which, in light of the analysis conducted in the previous chapter, I propose provides a theory of welfare that is ill-suited to conducting economic farm animal welfare policy. Fortunately, by moving to a values-based policy framework, all three concerns can be overcome.

3. Cracks in the Foundations

a. *Social Utility*

McNerney claims (2004: 47) that interventions at a merit good level, where significant groups of people are uninterested in the further provision of farm animal welfare, can be justified as being in the public interest. Because some people are uninterested in the further provision of welfare and some merit good policies are costly, it is likely that merit good policies will make some people better-off and others worse-off. By stating that these policies are in the public interest, McNerney implicitly invokes a conception of social utility that permits interpersonal utility comparisons: he is claiming that the costs incurred by those who are uninterested in the further provision of farm animal welfare are outweighed by the benefits received by others.

Part of why McNerney's framework produces this outcome is because it relies upon the concept of merit goods, those goods whose further provision is deemed advantageous to society beyond the public's demand for them.¹²⁶ The concept of merit goods was first proposed by R.A. Musgrave in his 1959 *The Theory of Public Finance*, where he provided two reasons why governments might intervene in markets producing merit goods. First, merit goods produce positive externalities, meaning their social benefits are greater than their private benefits. Consequently, Musgrave argued that intervention in merit good markets could be advantageous to social utility (Musgrave 1959: 13; Pulsipher 1971: 267).

¹²⁶ The counterpart to the merit good is the *demerit good*, one whose further provision is viewed as detrimental to society, irrespective of preferences for it. If good animal welfare is viewed as a merit good, then poor animal welfare would be the corresponding demerit good.

This justification for merit good interventions, however, is not compatible with the arguments presented in this thesis. In Chapter 1, I outlined the theory of preference utilitarianism and showed how it underpins the concept of efficiency. In short, people can place a subjective monetary value on utility gains and losses, and, if policy losers are compensated to the extent that nobody is made worse-off by a reallocation of resources and somebody is made better-off, we can identify an improvement in efficiency and social utility. Due to different marginal utilities of wealth, *compensation must be paid*, otherwise we cannot be certain of the overall effects of a reallocation upon efficiency.

By the concept's definition, there are significant groups of people who are uninterested in the further provision of merit goods. These people will be policy losers when policymakers encourage markets to provide more farm animal welfare: they bear the costs of these interventions but experience no benefit.¹²⁷ When compensation is not paid, the effects of merit good policies upon efficiency and social utility remain uncertain. It is also surely unfeasible to pay compensation to policy losers in these cases, due to the range of informational difficulties outlined in Chapter 4 (p.81): it is difficult to identify policy winners and losers, establish reliable figures for WTP and WTA and extract payment from policy winners. Given the interpretation of Kaldor-Hicks efficiency adopted in this thesis, an appeal to social utility cannot justify merit good interventions in animal product markets.

Musgrave provides a second argument for justifying merit good interventions, one which moves away from social utility and instead focuses upon the individual's utility. He states (1959: 14) that preferences can be based upon incomplete information, or unfairly swayed by advertisers and generally distorted by modern society. Consequently, people may not recognise the full value of a good until after they have consumed it. Government intervention to encourage the supply of merit goods might therefore be justified from the perspective of individual utility: it is in each person's interest to be encouraged to consume merit goods and avoid demerit goods. For this argument to provide a basis for merit good interventions in animal product markets, it could be argued that, if people were more aware of the lives lived by farm animals in low-welfare production, their preferences would change to be more supportive of higher-welfare production and less supportive of low-welfare production. Consequently, it could be claimed that any policy that sought to achieve the level of supply demanded by these informed preferences would be in

¹²⁷ There are rare exceptions: a tax on the demerit good of animal harm, levied on lower-welfare animal products, will impose costs only upon market participants. An animal abolitionist such as Gary Francione, for instance, would not participate in such markets, and may not be interested in the provision of farm animal welfare due to his opposition to all animal ownership.

an individual's interests. The level of paternalism in this argument, however, as well as the assumption that consumers of low-welfare goods must be uninformed, surely renders this claim problematic.

With a minor adjustment, Stefan Mann (2005b: 138) removes the paternalistic element of this argument while preserving its core assertion. He proposes that people have different types of preferences that are operative in different environments. Mann suggests that people have market preferences, which can be discerned by observing consumer behaviour, and reflective preferences, which reflect their moral attitudes, in keeping with the consumer and citizen schemas discussed in the previous chapter (p.169). These preferences can conflict, and, as we saw in Chapter 1's discussion of the meat paradox (p.30), markets in animal products facilitate such a clash. Mann suggests that some people will be able to improve their utility by voting for governments that promise welfare-motivated merit good interventions in animal product markets. Such an argument loses its paternalistic overtones, but is applicable only to those people who hold reflective preferences for farm animal welfare above the level of cruelty. It remains possible that there are significant groups of people whose reflective and market preferences are aligned and are simply not interested in farm animal welfare beyond the point of cruelty, who will therefore still be left worse-off by merit good interventions. Consequently, preference-based justifications for merit good interventions appear largely incompatible with the account of social utility held in this thesis.

b. Efficiency

A second problem emerges when we recognise that preferences for farm animal welfare are often context-sensitive: this undermines the notion that markets are efficient institutions that enable consumers to satisfy their preferences for farm animal welfare. In the previous chapter (p.171), I argued that economic choice environments such as supermarkets are likely to prime people to view farm animal welfare through a consumer schema, which encourages self-interested behaviours at the expense of other-regarding ones. In other types of choice environment, citizen schemas may be activated, encouraging other-regarding behaviours at the expense of self-interested ones.

When preferences are context-sensitive, we may struggle to make reliable claims about efficiency. This is because a status quo reference point against which we assess potential states of the world lies at the heart of the Pareto and Kaldor-Hicks theories, but we are denied this by the context-sensitive nature of many people's preferences for farm animal welfare: the status

quo changes according to the environment in which people are asked to value farm animal welfare. This gives us good reason to reject Pareto and Kaldor-Hicks efficiency as policy goals in this area. It is surely meaningless to talk of aiming for a state of the world in which someone is made better-off and nobody is made worse-off when preferences are built on unstable sands – each person’s interpretation of whether a policy has made them better or worse off is influenced by the environment in which they make their assessment.

Similarly, the usefulness of willingness-to-pay studies is reduced in cases where answers are influenced by decision-making environments: their findings will have less relevance in other types of decision-making environments. Basing policy upon such studies could therefore be inappropriate: willingness to pay in a survey context may not be reflective of willingness to pay in a purchase environment, and we do not have good reason to afford greater significance to either citizen or consumer responses.¹²⁸

c. Welfare

The third reason that McNerney’s framework is ill-suited to this thesis’ project is that it is at least partially grounded in preference satisfaction theories of welfare. This is explicit in McNerney’s discussion of farm animal welfare as a private good, where he argues that welfare standards will not figure in the ‘preference functions of the majority’ (McNerney 2004: 46).¹²⁹ The argument that many people’s preferences for farm animal welfare are context-sensitive renders preference satisfaction an unsuitable theory of welfare upon which to ground economic farm animal welfare policy.

One of the arguments presented in favour of preference satisfaction theories of welfare is that they afford respect to the individual and what matters to them (Nussbaum 2001: 70). Because preferences for farm animal welfare can be unstable and their expression influenced by choice environments, however, it is unlikely that the preferences expressed for farm animal welfare in a supermarket environment will fully capture everything that matters to consumers. Instead, the economic choice environments that dominate interactions with farm animal welfare subtly prime self-interested preferences. This is likely to lead to the ‘crowding out’ of ethical concern

¹²⁸ As noted in Chapter 2 (p.59), however, the use of a ‘cheap talk script’, intended to remind respondents of the discrepancies often found between stated and revealed preference studies, can be sufficient to eliminate or at least greatly reduce these discrepancies. The script may be enough to spur System 2 processes into action and encourage participants to afford greater scrutiny to their answers.

¹²⁹ McNerney is more ambiguous in his discussion of public and merit goods, arguing that intervention at these levels is justified due to feelings of discomfort created by animal welfare (McNerney 2004: 46-47).

for farm animal welfare in quotidian interactions with animal products (Sandel 2013: 9), with implications for consumers' welfare: what satisfies consumers' preferences in the supermarket may leave them feeling worse-off elsewhere.

In Chapter 2, I argued that there exists both substantial concern for farm animal welfare and willingness to pay to improve welfare. In the previous chapter (p.172), we saw that, in a survey environment, there is substantial altruistic concern for farm animal welfare. Consequently, we may be concerned that markets encourage the expression of self-interested preferences at the expense of altruistic preferences. Because markets in animal products are liable to thwart the expression of altruistic concern for farm animal welfare, interventions seeking to improve animal welfare may be welfare-increasing for both consumers and farm animals. Despite markets satisfying consumer preferences, interventions may nevertheless be welfare-increasing.

4. Creating a Values-Based Policy Framework

Because of these concerns about preference satisfaction theories of utility and welfare, it is desirable to develop a policy framework that moves away from preference satisfaction as a policy goal. If preferences do not offer a firm foundation for conducting economic farm animal welfare policy, we need to reconsider what policy should be grounded upon.

Despite their flaws leaving them ill-suited to this thesis' project, preference satisfaction theories of utility and welfare are valuable, as they put what matters to people at the heart of policy analysis (Norwood and Lusk 2011: 198). We can maintain this commitment to human-centred policy by focusing instead on people's *values*. This may not appear obvious at first: as with preferences, the expression of values is also influenced by choice environments. The difference, however, is that choice environments seem to affect the *weighting* of values, rather than their nature. It is plausible to believe that the values themselves are not reversed as can be the case with preferences, but are instead afforded greater or lesser significance in decision-making. Policymakers may therefore find that values provide an appropriate basis for economic farm animal welfare policy. If so, policy ought to ensure that different ways of valuing farm animal welfare are afforded appropriate representation in society. When market norms are the dominant form of interaction with farm animal welfare, and ethical concern is suppressed, we can argue that ethical values are under-represented in society. In these cases, the role of policy is to afford greater representation to ethical ways of valuing farm animal welfare.

Significantly, this argument does not assume any normative position regarding the moral status of farm animals and does not attribute special significance to ethical values. Even if we believe that human interests take priority over all else and ethical values possess no special significance relative to other types of value, we can still be concerned that the preponderance of market interactions with farm animal welfare has crowded out the expression of ethical preferences; a consequence of this crowding out may be that human-farm animal interactions are not reflective of the range of ways that farm animal welfare matters to people. While a traditional preference satisfaction analysis struggles to acknowledge that this is problematic, a values-based approach to economic farm animal welfare policy recognises that market valuations of farm animal welfare are unlikely to fully-capture public concern for farm animals. This values-based approach might therefore be of benefit to society in the sense that it advocates outcomes that reflect what matters to the public in both market and non-market contexts.

That said, there may be concerns about how we should, or even can, measure public values. First, different people value farm animal welfare in different ways and to different extents. In some respects, this presents a challenge to policymakers seeking to determine what the public's values are, although, as I argue below, there are likely to be cases where there is a general consensus that the provision of farm animal welfare is either unacceptable or sufficient. As with preference satisfaction theories, the challenge is to identify appropriate policy goals in situations where there is no consensus, although I argue below that a values-based framework may be better suited to this task than a preference satisfaction framework.

Second, and perhaps more significantly, it is not clear how policymakers can determine what an individual's values, let alone those of the public, are. Observing behaviour in an economic choice environment might provide an indication of a person's consumer valuation of farm animal welfare, while surveys might provide an indication of their citizen valuation of farm animal welfare, but more work is needed to reconcile these two conflicting modes of thinking. Should policymakers weigh the citizen valuation more heavily than the consumer valuation or vice versa, or should they balance the two equally and seek a middle ground? There is no obvious answer, and it is not clear that a democratic deliberative process, whereby the public gets to decide the relative weights afforded to economic and non-economic concerns, would yield a satisfactory result: such a process is likely to afford undue significance to non-economic concerns.

Nevertheless, a values-based framework still allows us to make some broad claims, and thus is still of use. We can recognise that the primarily economic nature of interactions with farm

animal welfare can deny people the opportunity to fully-express their concern for animal welfare. It is therefore likely that other-regarding concerns are being underrepresented in both regular interactions with farm animal welfare and its overall provision in society. Although further work is needed to determine how far non-economic valuations should be integrated into economic farm animal welfare policy, it appears that the prevailing preference satisfaction model does not afford due significance to non-economic values, which might justify at least some tentative steps towards some of the policies advocated in this thesis.

With this in mind, it is appropriate to revisit the three-tier framework with the intention of producing a policy structure that does justice to the multiplicity of ways that people value farm animal welfare.

a. Public Goods

As McNerney recognises, the lowest levels of farm animal welfare – which are characterised by cruelty – are deemed unacceptable by almost everyone in society. While McNerney frames this in terms of the discomfort felt by people after becoming aware of animal cruelty in society, a values-based framework ventures deeper, and suggests that one of the causes of this discomfort is that animal cruelty is incompatible with how animal welfare is valued in a society. Although different people value animal welfare in different ways and to different extents, almost all agree that certain practices and production methods, defined as cruel practice, are unacceptable and thus incompatible with their values.

From a policy perspective, nothing changes: because almost everyone in society is united in opposition to cruel practice and is willing to pay to eradicate it, and because markets are ineffective institutions for reflecting concern for farm animal welfare, policymakers have a mandate for direct intervention. Although farm animal welfare is not a public good in the technical sense, the aims of policy at this level – provision and protection – are shared with public good policies.

As discussed in Chapter 5, policymakers may seek to achieve public good objectives through the use of legislation and taxes. Several factors must be considered when drafting effective animal welfare legislation. Ideally, welfare requirements will be applicable to both domestic and imported products. Failure to achieve this risks undermining the domestic industry's competitiveness against cheaper, lower-welfare imports, and could effectively export cruel practices abroad. Alternatively, an effective labelling scheme could mitigate the impact of

imports by making it clear to consumers that domestic goods have been produced to higher standards of welfare.¹³⁰ Additionally, the costs of making welfare improvements could be mitigated through public good taxes.

There is also scope for animal welfare legislation to be supplemented with externality policies. In line with giving farm animals ownership of their own welfare, legislation could make a life worth living the minimum standard of acceptable welfare, thus putting animal interests at the heart of animal agriculture. Cap and trade could be used to phase out low-welfare practices in the run-up to a full ban, to ensure greater adherence to the law upon its introduction.

b. Merit Goods

A values-based framework differs significantly from McInerney at the merit good level. As discussed earlier, McInerney frames the justification for merit good interventions in terms of social utility by stating that merit good interventions are in the public interest. We can avoid McInerney's problematic reading of social utility by justifying merit good interventions with an appeal to public values.

A values-based framework shifts the focus away from satisfying preferences and towards the appropriate representation of different ways of valuing farm animal welfare in society. There are several reasons to contend that other-regarding values relating to farm animal welfare are underrepresented in society, all of which stem from the fact that the majority of interactions with farm animal welfare are facilitated by markets for animal products. Those who abstain from these markets, such as vegetarians and vegans, have limited opportunities to express other-regarding values for farm animal welfare (McVittie et al. 2006: 6; Lusk 2011: 564). Those who care about animal welfare but cannot afford higher-welfare goods, which are typically more expensive, may also have limited opportunities to express their values (Fearing and Matheny 2007: 165). Additionally, people who do participate in the market and have strong other-regarding values relating to farm animal welfare are limited in the extent to which they can express these values by the quantity of animal products they can consume. Last, as proposed in Chapter 7 (p.171), markets discourage the expression of other-regarding values at the expense of self-interested ones.

¹³⁰ Requiring imports to display such a label may violate WTO rules, but it may be enough to require only domestic products to display this label. By focusing on positive elements of animal welfare, such a label may also successfully tap into consumers' emotional engagement with farm animal welfare.

Intervention is therefore warranted because the market is failing in its aims. The market is feted as a tool that enables people to get what they want, but often fails to do so in the case of farm animal welfare. The market instead denies some people the opportunity to express concern for farm animal welfare and distorts the preferences of others. Intervention is thus justified at the merit good level because farm animal welfare is important to broad swathes of society for a range of reasons and markets currently fail to respond appropriately to this: in particular, other-regarding values are underrepresented. Consequently, merit good interventions in animal product markets should increase the representation of other-regarding values.

It must be recognised that some merit good policies have the potential to leave some people worse-off. A person who is indifferent to farm animal welfare improvements at the merit good level will be made worse off if taxes on low-welfare goods make his weekly shop more expensive. This is not as serious an issue for a values-based framework as it is for a preference-based framework, however: in a values-based framework, it is more acceptable for policy to make some people worse-off in pursuit of an appropriate representation of society's values, while, in a preference-based framework, the problem of interpersonal utility comparisons makes it hard to justify making some people worse off and others better-off in the name of greater efficiency. That said, other public values, namely those connected to justice and food justice, must also be taken seriously at this level. Policies that leave people deprived of food or unable to participate in their food culture are unlikely to be compatible with public values, and so policymakers must be careful to recognise and balance considerations relating to a range of public values.

This end can be achieved with complementary policies. Externality taxes on low-welfare goods can be combined with subsidies for higher-welfare goods to limit how these policies impact those on lower incomes. System 1 nudges, those which channel automatic behaviours towards higher-welfare purchases and away from low-welfare ones, might also be useful: a subsidy on higher-welfare goods could be advertised as a discount, which might appeal to consumers operating on price or value heuristics. Alternatively, System 1 nudges that placed higher-welfare goods in more prominent locations in supermarkets could guide consumer attention to these products. This would likely increase sales of higher-welfare goods but would not impose a financial burden upon consumers as prices and the availability of goods would remain unchanged.

Similarly, a market for animal well-being units, which allows producers to sell credits received for implementing welfare improvements, could subsidise higher-welfare production without imposing costs upon society. Such a policy, however, cannot hope to internalise all

psychological externalities in this area: animal abolitionists will be unwilling to support policies that implicitly condone the continued ownership of farm animals and their use as mere means to human ends. This highlights the importance of tax relief for charitable donations: this policy allows people who care about farm animals but are unwilling to support institutions that condone animal agriculture to express their values.

Of perhaps greater concern are the informational issues that prevent policymakers from identifying a socially-optimal level of farm animal welfare. Just as this issue plagued a preference-based policy framework, so it hinders the effectiveness of a values-based one. Nevertheless, policymakers can still use externality policies to seek a greater balance between self-interested and altruistic values, without being certain that a perfect balance has been achieved.

c. Private Goods

Finally, I depart from McInerney at the private good level of provision. McInerney states (2004: 48) that animal welfare at the private good level is ‘not of relevance for public policy’, but I propose that policy has an important role to play here.

My values-based framework supports the assertion that the highest levels of farm animal welfare either accord with or go beyond what is required by most people’s values. Direct provision through public good policies or market interventions through merit good policies are costly and therefore inappropriate. Further intervention may in fact deny appropriate representation to the public’s self-interested values, and so is inappropriate. Instead, those whose values are incompatible with the prevailing level of provision should be empowered to express their values through the market.

Despite McInerney’s contention that there is no role for public policy at private good levels of provision, I disagree: policymakers should enable and empower consumers to reliably express concern for farm animal welfare. Two policies support this goal: labelling schemes allow consumers to reliably differentiate between higher- and lower-welfare goods; System 2 nudges, which influence the building blocks of reflective analysis, can encourage consumers to afford greater weight to their other-regarding values at points of purchase. The two policies can work in conjunction: appropriate in-store information provision could nudge consumers to seek out animal welfare labels.

5. Conclusion

To conclude, when developing economic farm animal welfare policy, policymakers should consider both public values and facts about the provision of animal welfare in society. When almost everyone is opposed to prevailing welfare conditions, there is scope for direct intervention to provide and protect farm animal welfare, either through legislative prohibitions or publicly funded improvements. Where significant groups are uninterested in the further provision of animal welfare, policy may be used to encourage markets to supply higher levels of welfare in recognition of the facts that (1) animal product markets provide the primary mode of interaction with farm animal welfare and (2) many people cannot or do not fully express their values relating to farm animal welfare in these markets. Policies such as externality taxes and subsidies, market creation and behavioural nudges enable stakeholders to have their voices heard and their values represented in society. Last, when public values are largely compatible with a society's prevailing levels of farm animal welfare, policy should enable and empower dissatisfied consumers to express their concern in the market through effective welfare labelling and nudges that encourage consumers to consider their other-regarding values in economic choice environments.

Appendices

Appendix I – A Comparison of Prices of Free Range and Caged or Barn Eggs

Appendix II – A Graphical Explanation of Externalities

Appendix III – Externality Taxes and Subsidies

Appendix IV – Nudges in the People’s Supermarket

Appendix I

Appendix I – A Comparison of Prices of Free Range and Caged or Barn Eggs

Free Range Product	Cost (pence)	Cost per egg (pence)	Caged / Barn Product	Cost (pence)	Cost per egg (pence)	% Difference (2 d.p.)
Tesco Mixed Sized Free Range Eggs 15 Pack	200	13.3	Tesco Mixed Sized Everyday Value Eggs 15 Pack	125	8.3	62.50
Sainsbury's Free Ranged Mixed Weight Eggs x15	200	13.3	Sainsbury's Barn Eggs, Basics x15	125	8.3	62.50
Morrisons Free Range Eggs Mixed Weight 15 per pack	200	13.3	M savers Eggs 18 per pack (Morrisons)	143	7.9	59.58
ASDA 15 Mixed Weight Free Range Eggs	200	13.3	ASDA Smartprice 15 Eggs	119	7.9	59.50
15 Pack Medium Free Range Eggs (Aldi)	189	12.6	15 British Eggs (Aldi)	119	7.9	62.96
Iceland 12 Large Free Range Eggs	150	12.5	Iceland Class 'A' 15 Mixed Sized Fresh Eggs	125	8.3	66.67


Mean % difference in price between free range and cage / barn eggs (2 d.p.): 62.29

(Sources accessed 05/12/16)

Appendix II – A Graphical Explanation of Externalities

a. Negative Externalities

(Adapted from Parkin 2000: 438)




This graph shows the supply (S) and demand (D) for Y , a good whose production creates a negative externality. The market equilibrium for Y is represented by quantity Q_0 and price P_0 . However, because Y imposes a negative externality upon society, there are costs associated with it that are not factored into its market price. This means that the social costs of X are greater than its market price at equilibrium. The social cost of the good is represented by the marginal social cost curve (MSC), which is the sum of the market price and the value of the externality. The value of the externality can therefore be found by subtracting P_0 (the private cost of the good, represented by its market price) from SC_0 (the social cost); in the graph this is represented by the distance between A and B . When considering the full social costs, it is clear that market forces will cause Y to be oversupplied; because the externality does not have to be factored into

the cost of production, firms have an incentive to use more of it than they would if it were a costly input. The socially-optimal quantity and price of Y can be found instead at Q_i and P_i , where the MSC curve intersects the demand curve. At this socially-optimal level, the graph shows that the cost of the resources used in production is C_i , and the cost of the externality can be calculated by subtracting C_i from P_i .

b. Positive Externalities

(Adapted from Begg et al. 2000: 268)



This graph illustrates the supply (S) and demand (D) of X, a good that produces a positive externality. The market equilibrium prices and quantities are found where the supply curve intersects the demand curve, in this case at (P_0, Q_0) . However, the externality effects are not factored into the market price and quantity supplied. The full benefit can be measured using the marginal social benefit (MSB) curve, with the value of the externality at the market equilibrium level of supply calculated by subtracting P_0 (the private benefit, represented by the market price) from SB_0 (the social benefit), which is represented in the graph as the distance between A and B.

It can thus be seen that markets tend to undersupply goods that bear positive externalities; at the market equilibrium price of P_0 the market provision of X is lower than the socially-optimal level. This socially-optimal level of X is found where the MSB curve intersects the supply curve, in this instance at price SB_1 and quantity Q_1 , and the value of the externality is calculated by


subtracting the private benefits of X , represented by the price paid by consumers (P_i), from the total value of X 's benefits (SB_i).

In short, welfare is maximised in externality scenarios when all the costs and benefits provided by a good are considered, rather than just private ones. This requires the use of marginal social costs and benefit curves instead of the supply and demand curves, which capture only the private costs and benefits. Individual decisionmakers, however, are likely to only consider their private costs and benefits, leading to the undersupply of goods that carry positive externalities and the oversupply of goods that bear negative externalities.

Appendix III – Externality Taxes and Subsidies

a. Externality Taxes

Taxing Negative Externalities
(Adapted from Parkin 2000: 438)




The market equilibrium for good X is found at the intersection between the supply (S) and demand (D) curves, at Q_0, P_0 . The marginal social cost (MSC) curve lies above the supply curve, however, indicating that some of X 's costs are not factored into its market price, i.e. there are negative externalities. At the market equilibrium level of supply, the marginal social cost is SC_0 , while the buyer only pays P_0 . The socially-optimal supply of X is found where the demand curve intersects the MSC curve at Q_1 ; left to its own devices, the market will therefore oversupply X . To rectify this, a purchase tax can be introduced to shift the equilibrium to the socially-optimal level.¹³¹ Calculating the optimal level of taxation, which is equivalent to the value of the externality, requires us to subtract C_1 – the marginal private cost of X at the socially-optimal equilibrium – from P_1 – the marginal social cost of X at the socially-optimal equilibrium. This

¹³¹ A production tax could be levied instead; this would shift the supply curve left, so it lies on the MSC curve.

tax causes a leftward shift in the demand curve, from D to D_1 . A price increase would typically encourage sellers to supply more of a good. Because, however, this price increase is the result of a tax, sellers are not incentivised to increase supply as the amount they receive from each sale is unaffected by the change. Consequently, the supply curve remains unchanged, and the new market equilibrium is found where D_1 intersects S , which corresponds with the intersection between D and MSC at Q_1 (Schotter 1997: 564; Snyder and Nicholson 2012: 626; Nechyba 2017: 747).

b. Externality Subsidies

Subsidising Positive Externalities



The market equilibrium for good Y occurs where the supply (S) and demand (D) curves intersect at (Q_0, P_0) . Y produces benefits which are not factored into its price: we know this because Y 's marginal social benefit (MSB) curve, which captures the its full social benefits, lies above the demand curve, which represents Y 's private benefits. While firms receive the market equilibrium price of P_0 , the full value of Y 's benefits is greater, represented by SB_0 . The socially-optimal level of provision is found at Q_1 , where the MSB curve intersects the S curve, but, because sellers cannot capture payment for all of Y 's benefits, they supply a lower amount.

A subsidy, given to either sellers or buyers, can resolve this undersupply. This graph illustrates a seller subsidy; a buyer subsidy would simply shift the demand curve right to the MSB curve. The optimal level of subsidy is equal to the marginal value of the externality at the socially-optimal level of production, which is represented by the distance between A (the marginal value of Y at the socially-optimal level of supply) and B (the marginal value of Y 's private benefits at the socially-optimal level of supply). In essence, the subsidy broaches the gap between Y 's private benefits and its social benefits at the optimal level of supply. A seller subsidy lowers

production costs, which encourages sellers to increase supply. The subsidy's benefits are also enjoyed by buyers, as lower production costs translate into a lower price (P_1), which increases demand for Y to Q_1 (Nechyba 2017: 753).

Appendix IV – Nudges in the People’s Supermarket

Using Eye-Level Product Placement to Nudge Healthier Behaviours



Using End-of-Aisle Displays to Nudge Healthier Behaviours



Information Provision through 'Nudge Points'



Encouraging Healthy Behaviours through 'Nudge Points'



Nutritional Information Provision

Eatwell Guide

Use the Eatwell Guide to help you get a balance of healthier and more sustainable food. It shows how much of what you eat overall should come from each food group.

Check the label on packaged foods

Each serving (150g) contains

Energy	Fat	Saturated fat	Sugars	Salt
1569kJ 369kcal	3.0g	1.3g	34g	0.9g
	LOW	LOW	HIGH	MED
	13%	4%	7%	35%

of an adult's reference intake
Typical values (as sold) per 100g: 697kJ/167kcal

Choose foods lower in fat, salt and sugars

Fruit and vegetables

Eat at least 5 portions of a variety of fruit and vegetables every day

Choose wholegrain or higher fibre versions with less added fat, salt or added sugar

Potatoes, bread, rice, pasta and other starchy carbohydrates

Beans, pulses, fish, eggs, meat and other proteins

Eat more beans and pulses, 2 portions of sustainably sourced fish per week, one of which is oily. Eat less red and processed meat

Dairy and alternatives

Choose lower fat and lower sugar options

Oil & spreads
Choose unsaturated oils and use in small amounts

Per day 2000kcal 2500kcal = ALL FOOD + ALL DRINKS

Source: Public Health England in association with the Welsh Government, Food Standards Scotland and the Food Standards Agency in Northern Ireland. © Crown copyright 2018

Bibliography

- Abbate, C.S., Ruggieri, S. and Boca, S. 2013a. Automatic influences of priming on prosocial behavior. *Europe's Journal of Psychology*. 9: 479-92.
- Abbate, C.S., Ruggieri, S. and Boca, S. 2013b. The effect of prosocial priming in the presence of bystanders. *The Journal of Social Psychology*. 153: 619-22.
- Adler, M., and Posner, E. 2006. *New Foundations of Cost-Benefit Analysis*. Cambridge, MA: Harvard University Press.
- Agriculture and Horticulture Development Board (AHDB). 2019a. *Deadweight Cattle Prices*. 3 August. <Accessed 10/08/19 at <http://beefandlamb.ahdb.org.uk/markets/deadweight-price-reports/deadweight-cattle-prices/>>.
- AHDB. 2019b. *Pig holdings in the UK*. <Accessed 13/07/19 at <https://pork.ahdb.org.uk/prices-stats/industry-structure/pig-holdings-in-the-uk/>>.
- AHDB. 2019c. *UK pig herd*. <Accessed 13/07/19 at <https://pork.ahdb.org.uk/prices-stats/industry-structure/uk-pig-herd/>>.
- AHDB. 2018. *UK pig meat supplies forecast*. 31 July.
- AHDB. 2017. *Brexit Scenarios: An Impact Assessment*. Warwickshire: Agriculture and Horticulture Development Board.
- Akaichi, F. and Revoredo-Giha, C. 2016. Consumers demand for products with animal welfare attributes: evidence from homescan data for Scotland. *British Food Journal*. 118: 1682-1711.
- Akerlof, G.A. 1970. The market for "lemons": quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*. 84: 488-500.
- Almeida Pez, I.C.L., Mendes, A.A., Balog, A., Vulcano, L.C., Ballarin, A.W., Almeida, I.C.L., Takahashi, S.E., Komiyama, C.M., Silva, M.C. and Cardoso, K.F.G. 2008. Study on the bone mineral density of broiler suffering femoral joint degenerative lesions. *Brazilian Journal of Poultry Science*. 10: 103-8.
- Ancev, T. 2011. Policy considerations for mandating agriculture in a greenhouse gas emissions trading scheme. *Applied Economic Perspectives and Policy*. 33: 99-115.
- Andreoni, J. 1990. Impure altruism and donations to public goods: a theory of warm-glow giving. *The Economic Journal*. 100: 464-77.
- Andreoni, J. 1989. Giving with impure altruism: applications to charity and Ricardian equivalence. *Journal of Political Economy*. 97: 1447-58.
- Andreoni, J. 1988a. Privately provided public goods in a large economy: the limits of altruism. *Journal of Public Economics*. 35: 57-73.
- Andreoni, J. 1988b. Why free ride? Strategies and learning in public good experiments. *Journal of Public Economics*. 37: 291-304.
- Animal Equality. 2019. *High-Rise of Hell*. 24 June.
- Appleby, M. 1999. *What Should We Do About Animal Welfare?* London: Blackwell Science.
- Ariely, D. 2008. *Predictably Irrational: The Hidden Forces that Shape Our Decisions*. London: HarperCollins.

- Arrow, K.J. 1963. *Social Choice and Individual Values*. New York: Cowles Foundation.
- Bagus, P. 2011. The Eurosystem: costs and tragedies. In *Institutes in Crisis: European Perspectives on the Recession*, ed. D. Howden, 117-41. Cheltenham: Edward Elgar Publishing.
- Bailey, R., Froggatt, A. and Wellesley, L. 2014. *Livestock – Climate Change’s Forgotten Sector: Global Public Opinion on Meat and Dairy Consumption*. London: Chatham House, The Royal Institute of International Affairs.
- Bannock, G., Baxter, R.E. and Rees, R. 1984. *The Penguin Dictionary of Economics*. Middlesex: Penguin Books.
- Bargh, J.A. and Ferguson, M.J. 2000. Beyond behaviorism: on the automaticity of higher mental processes. *Psychological Bulletin*. 126: 925-45.
- Bargh, J.A., Chen, M. and Burrows, L. 1996. Automaticity of social behavior: direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology*. 71: 230-44.
- BBC. 2019. *Chickens 'left to rot' at major supplier Moy Park's farms*. 25 June.
- BBC. 2012. *UK battery farms break EU rules*. 13 January.
- BBC Newsday. 2019. *Could low emission cows help tackle climate change?* 19 April.
- BBC Newsround. 2012. *30 UK farms still keeping hens in battery cages*. 13 January.
- Begg, D., Fischer, S. and Dornbusch, R. 2000. *Economics*. Berkshire: McGraw-Hill Publishing Company.
- Bennett, R.M. 1997. Farm animal welfare and food policy. *Food Policy*. 22: 281-88.
- Bennett, R.M., Anderson, J. and Blaney, R.J.P. 2002. Moral intensity and willingness to pay concerning farm animal welfare issues and the implications for agricultural policy. *Journal of Agricultural and Environmental Ethics*. 15: 187-202.
- Bennett, R.M. and Blaney, R.J.P. 2003. Estimating the benefits of farm animal welfare legislation using the contingent valuation model. *Agricultural Economics*. 29: 85-98.
- Bennett, R. and Blaney, R. 2002. Social consensus, moral intensity and willingness to pay to address a farm animal welfare issue. *Journal of Economic Psychology*. 23: 501-20.
- Bentham, J. 2004. From an introduction to the principles of morals and legislation by Jeremy Bentham. In *Utilitarianism and Other Essays*, ed. A. Ryan, 65-111. London: Penguin Classics.
- Bergstrom, T., Blume, L. and Varian, H. 1986. On the private provision of public goods. *Journal of Public Economics*. 29: 25-49.
- Berndsen, M. and van der Pligt, J. 2004. Ambivalence towards meat. *Appetite*. 42: 71-78.
- Binmore, K.G. 2009. Interpersonal comparisons of utility. In *The Oxford Handbook of Philosophy of Economics*, ed. H. Kincaid and D. Ross, 540-559. New York: Oxford University Press.
- Binnekamp, M.H.A. and Ingenbleek, P.T.M. 2006. Market barriers for welfare product innovations. *NJAS – Wageningen Journal of Life Sciences*. 54: 169-78.
- Blackburn, S. 2008. *The Oxford Dictionary of Philosophy*. Oxford: Oxford University Press.

- Blair, J.P. 1974. Relative position, utility, and the compensation principle – a comment. *The American Economist*. 18: 124.
- Blandford, D. and Fulponi, L. 1999. Emerging public concerns in agriculture: domestic policies and international trade commitments. *European Review of Agricultural Economics*. 26: 409-24.
- Blandford, D. and Harvey, D. 2014. Economics of animal welfare standards: transatlantic perspectives. *EuroChoices*. 13: 35-40.
- Blaug, M. 2007. The fundamental theorems of modern welfare economics, historically contemplated. *History of Political Economy*. 39: 185-207.
- Blythman, J. 2015. *Swallow This: Serving up the Food Industry's Darkest Secrets*. London: Fourth Estate.
- Bollard, L. 2016. Initial grants to support corporate cage-free reforms. 31 March. *Open Philanthropy Project*.
- Bouvard, V., Loomis, D., Guyton, K.Z., Grosse, Y., El Ghissassi, F., Benbrahim-Tallaa, L., Guha, N., Mattock, H. and Straif, K. 2015. Carcinogenicity of consumption of red and processed meat. *The Lancet Oncology*. 16: 1599-1600.
- Bowles, D., Williams, R. and Robinson, C. 2017. *What are the WTO rules that affect animal welfare? Can you have trade bans?* <Accessed 28/07/19 at <http://politicalanimal.org.uk/wp-content/uploads/2017/03/BrexitWTOtradebans.pdf>>.
- Bracke, M.B.M., Spruijt, B.M., Metz, J.H.M. and Schouten, W.G.P. 2002. Decision support system for overall welfare assessment in pregnant sows A: model structure and weighting procedure. *Journal of Animal Science*. 80: 1819-34.
- British Hen Welfare Trust (BHWT). 2019. *About us*. <Accessed 11/12/19 at <https://www.bhwt.org.uk/about-us/>>.
- Brook Lyndhurst. 2012. *The Animal Welfare Provenance of Food – Communicating and Engaging with Consumers: A Review of Evidence and Interventions: A Report for Defra*. <Accessed 07/11/18 at https://www.google.co.uk/url?sa=t&rct=j&q=&esrc=s&source=web&cd=12&cad=rja&uact=8&ved=2ahUKewjjwImk8_rkAhWPfMAKHaztBYcQFjALegQIBBAC&url=http%3A%2F%2Frandd.defra.gov.uk%2FDocument.aspx%3FDocument%3D10455_121115-AnimalwelfareREAreportFINAL.pdf&usq=AOvVaw3ay7U9DAOBhMpaPCj8W-5b>.
- Brook Lyndhurst. 2010. *Are Labels the Answer? Barriers to Buying Higher Animal Welfare Products: A Report for Defra*. <Accessed 28/06/19 at http://randd.defra.gov.uk/Document.aspx?Document=Cl0102Animal_welfare_final_report.pdf>.
- Broome, J. 1991. Utility. *Economics and Philosophy*: 7: 1-12.
- Bullock, D. 2012. Emissions trading in New Zealand: development, challenges and design. *Environmental Politics*. 21: 657-75.
- Busby, M. 2019. EU imposes hen welfare standards on egg imports for first time. *The Guardian*. 2 October.
- Callaway, T.R., Morrow, J.L., Edrington, T.S., Genovese, K.J., Dowd, S., Carroll, J., Dailey, J.W., Harvey, R.B., Poole, T.L., Anderson, R.C. and Nisbet, D.J. 2006. Social stress increases fecal shedding of *Salmonella* Typhimurium by early weaned piglets. *Current Issues in Intestinal Microbiology*. 7: 65-72.

- Carattini, S., Carvalho, M. and Fankhauser, S. 2017. *How to Make Carbon Taxes More Acceptable: Policy Report*. London: Grantham Research Institute on Climate Change and the Environment and Centre for Climate Change Economics and Policy, London School of Economics and Political Science.
- Cards Against Humanity. 2013. Cards Against Humanity's "\$5 More" Black Friday Sale. *Tumblr*. 3 December. <https://cah.tumblr.com/>.
- Carrington, D. 2018. Huge reduction in meat-eating 'essential' to avoid climate breakdown. *The Guardian*. 10 October.
- Carrington, D. 2017. Meat tax 'inevitable' to beat climate and health crises, says report. *The Guardian*. 11 December.
- Cartwright, E. and Patel, A. 2010. Public goods, social norms, and naïve beliefs. *Journal of Public Economic Theory*. 12: 199-223.
- Carver, C.S., Ganellen, R.J., Froming, W.J. and Chambers, W. 1983. Modeling: an analysis in terms of category accessibility. *Journal of Experimental Social Psychology*. 19: 403-21.
- Caswell, J.A. and Mojduszka, E.M. 1996. Using informational labeling to influence the market for quality in food products. *American Journal of Agricultural Economics*. 78: 1248-53.
- Caswell, J.A. and Padberg, D.I. 1992. Toward a more comprehensive theory of food labels. *American Journal of Agricultural Economics*. 74: 460-68.
- Charities Aid Foundation (CAF). 2018. *CAF UK Giving 2018: An Overview of Charitable Giving in the UK*. <Accessed 25/11/18 at <https://www.cafonline.org/docs/default-source/about-us-publications/caf-uk-giving-2018-report.pdf>>.
- Cialdini, R.B. 2016. *Pre-Suasion: A Revolutionary Way to Influence and Persuade*. London: Random House Books.
- Cialdini, R.B. 2007. *Influence: The Psychology of Persuasion*. New York: Collins.
- Clark, B., Stewart, G.B., Panzone, L.A., Kyriazakis, I. and Frewer, L.J. 2017. Citizens, consumers and farm animal welfare: a meta-analysis of willingness-to-pay studies. *Food Policy*. 68: 112-27.
- Clark, C.F., Kotchen, M.J. and Moore, M.R. 2003. Internal and external influences on pro-environmental behavior: participation in a green electricity program. *Journal of Environmental Psychology*. 23: 237-46.
- Coase, R.H. 1960. The problem of social cost. *Journal of Law and Economics*. 3: 1-44.
- Codron, J., Giraud-Héraud, E. and Soler, L. 2005. Minimum quality standards, premium private labels, and European meat and fresh produce retailing. *Food Policy*. 30: 270-83.
- Colander, D. 2007. Retrospectives: Edgeworth's hedonimeter and the quest to measure utility. *Journal of Economic Perspectives*. 21: 215-25.
- Coleman, J.S. 1994. *Foundations of Social Theory*. Cambridge, MA: Harvard University Press.
- Commission Regulation (EC) No 543/2008 of 16 June 2008 laying down detailed rules for the application of Council Regulation (EC) No 1234/2007 as regards the marketing standards for poultrymeat.
- Commission Regulation (EC) No. 589/2008 of 23 June 2008 laying down detailed rules for implementing Council Regulation (EC) No 1234/2007 as regards marketing standards for eggs.

- Compassion in World Farming (CIWF). 2019. *Pig Welfare*. <Accessed 28/07/19 at <https://www.ciwf.org.uk/farm-animals/pigs/pig-welfare/>>.
- CIWF. 2016. *Sowing Fresh Seeds Post Brexit: Charter for a Caring Food Policy that Nourishes our Health, the Environment and Animal Welfare*. <Accessed 25/09/18 at <https://www.ciwf.org.uk/media/7428475/charter-post-brexit-charter-august-2016-final-referenced-version.pdf>>.
- CIWF. 2013. *Statistics: Laying Hens*. <Accessed 03/09/19 at <https://www.ciwf.org.uk/media/5235021/Statistics-Laying-hens.pdf>>.
- CIWF. 2012. *Nutritional Benefits of Higher Welfare Animal Products*. <Accessed 18/03/19 at <https://www.ciwf.org.uk/media/5234769/Nutritional-benefits-of-higher-welfare-animal-products-June-2012.pdf>>.
- Conner, D.S. 2004. Expressing values in agricultural markets: an economic policy perspective. *Agriculture and Human Values*. 21: 27-35.
- Conniff, R. 2016. Back in fashion. *National Geographic*. September: 96-113.
- Cotra, A. 2017. Initial grants to support corporate cage-free reforms. 15 September. *Open Philanthropy Project*.
- Council Directive 1999/74/EC of 19 July 1999 concerning laying down minimum standards for the protection of laying hens*.
- Council Regulation (EC) No 1099/2009 of 24 September 2009 on the protection of animals at the time of killing*.
- Council Regulation (EC) No 834/2007 of 28 June 2007 on organic production and labelling of organic products and repealing Regulation (EEC) No 2092/91*.
- Cowen, T. 2006. Market failure for the treatment of animals. *Society*. 43: 39-44.
- Creaghan, C. 2019. Thousands of dead chickens dumped outside intensive farm. *The Lincolnite*. 26 July.
- Curtis, S.E. 2007. Commentary: performance indicates animal state of being: a Cinderella axiom? *The Professional Animal Scientist*. 23: 573-83.
- Curtis, S.E. 1989. The case for intensive farming of food animals. In *Animal Rights and Human Obligations*, ed. T. Regan and P. Singer, 169-175. Englewood Cliffs, NJ: Prentice Hall.
- D'Silva, J. 2008. The urgency of change: a view from a campaigning organization. In *The Future of Animal Farming: Renewing the Ancient Contract*, ed. M.S. Dawkins and R. Bonney, 33-44. Oxford: Blackwell.
- Daniel, N., Casadevall, N., Sun, P., Sugden, D. And Aldin, V. 2020. *The Burden of Foodborne Disease in the UK 2018*. London: Food Standards Agency.
- Darby, M.R. and Karni, E. 1973. Free competition and the optimal amount of fraud. *The Journal of Law & Economics*. 16: 67-88.
- Darnton, A. 2016. *Our Food Future*. London: Food Standards Agency.
- Dasgupta, A.K. and Pearce, D.W. 1972. *Cost-Benefit Analysis: Theory and Practice*. London: The Macmillan Press Ltd.

- Dasgupta, P. 2005. What do economists analyze and why: values or facts? *Economics and Philosophy*. 21: 221-78.
- Datar, I., Kim, E. and d'Origny, G. 2016. New Harvest: building the cellular agriculture economy. In *The Future of Meat without Animals*, ed. B. Donaldson, and C. Carter, 121-31. London: Rowman & Littlefield International.
- Davidson, D., McKinsey, J.C.C. and Suppes, P. 1955. Outlines of a formal theory of value, I. *Philosophy of Science*. 22: 140-60.
- Davis, K.A. 1949. *Human Society*. New York: Macmillan.
- Dawes, R.M. and Thaler, R.H. 1988. Anomalies: cooperation. *Journal of Economic Perspectives*. 2: 187-97.
- Dawkins, M.S. 1980. *Animal Suffering: The Science of Animal Welfare*. London: Chapman and Hall.
- De Backer, C.J.S. and Hudders, L. 2015. Meat morals: relationship between meat consumption consumer attitudes towards human and animal welfare and moral behaviour. *Meat Science*. 99: 68-74.
- De Mol, R.M., Schouten, W.G.P., Evers, E., Drost, H., Houwers, H.W.J. and Smits, A.C. 2006. A computer model for welfare assessment of poultry production systems for laying hens. *Netherlands Journal of Agricultural Science*. 54: 157-68.
- De Scitovszky, T. 1941. A note on welfare propositions in economics. *The Review of Economic Studies*. 9: 77-88.
- Deci, E.L. 1971. Effects of externally mediated rewards on intrinsic motivation. *Journal of Personality and Social Psychology*. 18: 105-15.
- Department for Environment, Food & Rural Affairs (DEFRA). 2019a. *Agriculture in the United Kingdom 2018*. <Accessed 10/08/19 at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/815303/AUK_2018_09jul19.pdf>.
- DEFRA. 2019b. *United Kingdom Poultry and Poultry Meat Statistics – June 2019*. <Accessed 03/09/19 at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/826688/poultry-statsnotice-18jul19.pdf>.
- DEFRA. 2017a. *Farm Business Income by Type of Farm in England, 2016/17*. <Accessed 27/09/18 at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/655060/fbs-businessincome-statsnotice-26oct17.pdf>.
- DEFRA. 2017b. *United Kingdom Poultry and Poultry Meat Statistics – November & December 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/594026/poultry-statsnotice-26jan17.pdf>.
- DEFRA. 2017c. *United Kingdom Slaughter Statistics – December 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/591983/slaughtsl-statsnotice-19jan17.pdf>.
- DEFRA. 2016a. *United Kingdom Poultry and Poultry Meat Statistics – January 2016*. <Accessed 15/08/17 at

Bibliography

https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/510402/poultry-statsnotice-25feb16.pdf>.

DEFRA. 2016b. *United Kingdom Poultry and Poultry Meat Statistics – February 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/517884/poultry-statsnotice-24mar16.pdf>.

DEFRA. 2016c. *United Kingdom Poultry and Poultry Meat Statistics – March 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/525685/poultry-statsnotice-21apr16.pdf>.

DEFRA. 2016d. *United Kingdom Poultry and Poultry Meat Statistics – April 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/530753/poultry-statsnotice-26may16.pdf>.

DEFRA. 2016e. *United Kingdom Poultry and Poultry Meat Statistics – May 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/539525/poultry-statsnotice-23jun16.pdf>.

DEFRA. 2016f. *United Kingdom Poultry and Poultry Meat Statistics – June 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/548002/poultry-statsnotice-21jul16.pdf>.

DEFRA. 2016g. *United Kingdom Poultry and Poultry Meat Statistics – July 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/554414/poultry-statsnotice-25aug16.pdf>.

DEFRA. 2016h. *United Kingdom Poultry and Poultry Meat Statistics – August 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/561335/poultry-statsnotice-22sep16b.pdf>.

DEFRA. 2016i. *United Kingdom Poultry and Poultry Meat Statistics – September 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/572041/poultry-statsnotice-20oct16.pdf>.

DEFRA. 2016j. *United Kingdom Poultry and Poultry Meat Statistics – October 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/586020/poultry-statsnotice-24nov16.pdf>.

DEFRA. 2016k. *United Kingdom Slaughter Statistics – January 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/508324/slaughter-statsnotice-18feb16.pdf>.

DEFRA. 2016l. *United Kingdom Slaughter Statistics – February 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/516381/slaughter-statsnotice-17mar16a.pdf>.

- DEFRA. 2016m. *United Kingdom Slaughter Statistics – March 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/524102/slaughter-statsnotice-14apr16.pdf>.
- DEFRA. 2016n. *United Kingdom Slaughter Statistics – April 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/529609/slaughter-statsnotice-19may16.pdf>.
- DEFRA. 2016o. *United Kingdom Slaughter Statistics – May 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/537583/slaughter-statsnotice-16juni16.pdf>.
- DEFRA. 2016p. *United Kingdom Slaughter Statistics – June 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/546546/slaughter-statsnotice-14jul16.pdf>.
- DEFRA. 2016q. *United Kingdom Slaughter Statistics – July 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/552912/slaughter-statsnotice-18aug16.pdf>.
- DEFRA. 2016r. *United Kingdom Slaughter Statistics – August 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/559825/slaughter-statsnotice-15sep16.pdf>.
- DEFRA. 2016s. *United Kingdom Slaughter Statistics – September 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/570060/slaughter-statsnotice-13oct16.pdf>.
- DEFRA. 2016t. *United Kingdom Slaughter Statistics – October 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/577664/slaughter-statsnotice-17nov16.pdf>.
- DEFRA. 2016u. *United Kingdom Slaughter Statistics – November 2016*. <Accessed 15/08/17 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/584323/slaughter-statsnotice-15dec16.pdf>.
- DEFRA. 2013. *UK Packing Station Throughput and Prices*. <Accessed 03/09/19 at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/193615/eggs-packers-01aug13.xls>.
- DEFRA. 2011. *Greenhouse Gas Emission Projections for UK Agriculture to 2030*. <Accessed 20/02/19 at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/69225/pb13622-ghg-emission-projections.pdf>.
- Department for Business, Energy and Industrial Strategy (DBEIS). 2019. *2017 UK Greenhouse Gas Emissions*. <Accessed 11/12/19 at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/776083/2017_Final_emissions_statistics_one_page_summary.pdf>.
- Desvousges, W.H., Johnson, F.R., Dunford, R.W., Hudson, S.P., Wilson, K.N. and Boyle, K.J. 1993. Measuring natural resource damages with contingent valuation: tests of validity and reliability. In *Contributions to Economic Analysis: Volume 220*, ed. J.A. Hausman, 91-164. Oxford: North-Holland.
- Dicker, R. 2013. Cards Against Humanity Actually RAISED Prices On Black Friday. 3 December. *Huffington Post*.

- Doonan, G., Appelt, M. and Inch, C. 2009. Role of legislation in support of animal welfare. *The Canadian Veterinary Journal*. 50: 233-34.
- Dransfield, E., Ngapo, T.M., Nielsen, N.A., Bredahl, L., Sjöden, P.O., Magnusson, M., Campo, M.M. and Nute, G.R. 2005. Consumer choice and suggested price for pork as influenced by its appearance, taste and information concerning country of origin and organic pig production. *Meat Science*. 69: 61-70.
- Druce, C. and Lymbery, P. 2006. Outlawed in Europe. In *In Defense of Animals: The Second Wave*, ed. P. Singer, 123-31. Oxford: Blackwell.
- Duffy, R. and Fearne, A. 2009. Value perceptions of farm assurance in the red meat supply chain. *British Food Journal*. 111: 669-85.
- Dupuit, J. 1969. On the measurement of the utility of public works. In *Readings in Welfare Economics*, ed. K.J. Arrow and T. Scitovsky, 255-83. London: Allen and Unwin.
- Durning, A.T. and Brough, H.B. 1995. Animal farming and the environment. In *Just Environments: Intergenerational, International and Interspecies Issues*, ed. D.E. Cooper and J.A. Palmer, 149-64. London: Routledge.
- EAT-Lancet Commission. 2019. *Food Planet Health: Healthy Diets From Sustainable Food Systems*. <Accessed 29/02/20 at https://eatforum.org/content/uploads/2019/01/EAT-Lancet_Commission_Summary_Report.pdf>.
- Ecosia. 2019. How does the personal counter work? <Accessed 29/10/19 at <https://ecosa.zendesk.com/hc/en-us/articles/201657341-How-does-the-personal-counter-work->>>.
- Edgeworth, F.Y. 1881. *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*. London: C. Kegan Paul & Co.
- Edwards, W. 1954. The theory of decision making. *Psychological Bulletin*. 51: 380-417.
- Egg Info. 2019. *Industry data*. <Accessed 11/12/19 at <https://www.egginfo.co.uk/egg-facts-and-figures/industry-information/data>>.
- Engel, C. 2011. Dictator games: a meta-study. *Experimental Economics*. 14: 583-610.
- Epley, N. and Gilovich, T. 1999. Just going along: nonconscious priming and conformity to social pressure. *Journal of Experimental Social Psychology*. 35: 578-89.
- Ethical Consumer. 2018. *Markets Report 2017*. <Accessed 09/10/18 at <https://www.ethicalconsumer.org/sites/default/files/inline-files/ec%20markets%20report%202017.pdf>>.
- Eurobarometer. 2005. *Attitudes of Consumers towards the Welfare of Farmed Animals, Special Eurobarometer 229/Wave 63.2*. Brussels: European Commission.
- Eurobarometer. 2007. *Attitudes of EU citizens towards Animal Welfare, Special Eurobarometer 270/Wave 66.1*. Brussels: European Commission.
- Eurobarometer. 2016. *Attitudes of Europeans towards Animal Welfare, Special Eurobarometer 442/Wave EB84.4*. Brussels: European Commission.
- Evans, J.St.B.T. and Stanovich, K.E. 2013. Dual-process theories of higher cognition: advancing the debate. *Perspectives on Psychological Science*. 8: 223-41.

- Farm Animal Investment Risk and Return Initiative (FAIRR). 2017. *The Livestock Levy: Are Regulators Considering Meat Taxes? (White Paper)*. London: The FAIRR Initiative.
- Farm Animal Welfare Council (FAWC). 2011. *Economics and Farm Animal Welfare*. London: DEFRA.
- FAWC. 2009. *Farm Animal Welfare in Great Britain: Past, Present and Future*. London: DEFRA.
- FAWC. 2008. *Opinion on Policy Instruments for Protecting and Improving Farm Animal Welfare*. London: DEFRA.
- FAWC. 2006. *Report on Welfare Labelling*. London: DEFRA.
- Farrell, S. 2019. Hargreaves Lansdown faces questions over Neil Woodford fund. *The Guardian*. 5 August.
- Fazio, R.H., Sanbonmatsu, D.M., Powell, M.C. and Kardes, F.R. 1986. On the automatic activation of attitudes. *Journal of Personality and Social Psychology*. 50: 229-38.
- Fearing, J. and Matheny, G. 2007. The role of economics in achieving welfare gains for animals. In *The State of the Animals IV*, ed. D.J. Salem and A.N. Rowan, 159-73. Washington, DC: Humane Society Press.
- Fearne, A. and Lavelle, D. 1996. Segmenting the UK egg market: results of a survey of consumer attitudes and perceptions. *British Food Journal*. 98: 7-12.
- Fehr, E. and Gächter, S. 2000. Cooperation and punishment in public goods experiments. *The American Economic Review*. 90: 980-94.
- Fehr, E. and Schmidt, K.M. 1999. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*. 114: 817-68.
- Feldman, A.M. and Serrano, R. 2010. *Welfare Economics and Social Choice Theory*. New York: Springer.
- Festinger, L. 1957. *A Theory of Cognitive Dissonance*. Stanford, CA: Stanford University Press.
- Fine, C. 2012. *Delusions of Gender: The Real Science Behind Sex Differences*. London: Icon Books Ltd.
- Fishburn, P.C. 1968. Utility theory. *Management Science*. 14: 335-78.
- Fiske, S.T. and Taylor, S.E. 1991. *Social Cognition*. New York: McGraw-Hill.
- Flaherty, K.E. and Mowen, J.C. 2010. The role of schemas in consumer behavior research. In *Wiley International Encyclopedia of Marketing, Vol. 3*, ed. J.N. Sheth and N. Malhotra, 255-62. Chichester: Wiley-Blackwell.
- Fletcher, G. 2016. *The Philosophy of Well-Being: An Introduction*. London: Routledge.
- Food and Agriculture Organization of the United Nations (FAO). 2001. *Guidelines for Human Handling, Transport and Slaughter of Livestock*. <Accessed 13/07/19 at <http://www.fao.org/3/a-x6909e.pdf>>.
- Food Ethics Council and Pickett, H. 2014. *Farm Animal Welfare: Past, Present and Future*. <Accessed 24/01/19 at https://www.rspcaassured.org.uk/media/1041/summary_report_aug26_low-res.pdf>.

Food Standards Agency (FSA). 2019. *Public Attitudes Tracker*. <Accessed 07/10/19 at https://www.food.gov.uk/sites/default/files/media/document/public-attitudes-tracker-wave-18-final_o.pdf>.

FSA. 2018. Packaging and labelling. <Accessed 03/09/19 at <https://www.food.gov.uk/business-guidance/packaging-and-labelling>>.

Foresight. 2011. *The Future of Food and Farming*. Final Project Report. London: The Government Office for Science.

Francione, G.L. 2018. Should animal rights advocates promote “in-vitro” or “cultured” meat? <Accessed 01/12/19 at <https://www.abolitionistapproach.com/should-animal-rights-advocates-promote-in-vitro-or-cultured-meat/>>.

Francione, G.L. and Charlton, A. 2015. *Animal Rights: The Abolitionist Approach*. Exempla Press.

Frank, R.H., Gilovich, T. and Regan, D.T. 1993. Does studying economics inhibit cooperation? *Journal of Economic Perspectives*. 7: 159-71.

Frankfurt, H.G. 1971. Freedom of the will and the concept of a person. *The Journal of Philosophy*. 68: 5-20.

Fraser, D. 2012. Animal ethics and food production in the twenty-first century. In *The Philosophy of Food*, ed. D.M. Kaplan, 190-213. London: University of California Press.

Global Animal Partnership (GAP). 2019a. For shoppers. <Accessed 08/11/19 at <https://globalanimalpartnership.org/shoppers/>>.

GAP. 2019b. G.A.P. label quick reference guide. <Accessed 08/11/19 at <https://globalanimalpartnership.org/certified-gap>>.

GAP. 2019c. Laying hen standards. <Accessed 08/11/19 at <https://globalanimalpartnership.org/standards/laying-hen/>>.

Garnett, T. 2009. Livestock-related greenhouse gas emissions: impacts and options for policy makers. *Environmental Science & Policy*. 12: 491-503.

Gerber, P.J., Steinfeld, H., Henderson, B., Mottet, A., Opio, C., Dijkman, J., Falcucci, A. and Tempio, G. 2013. *Tackling Climate Change Through Livestock: A Global Assessment of Emissions and Mitigation Opportunities*. Rome: Food and Agricultural Organisation of the United Nations (FAO).

Goodland, R. and Anhang, J. 2009. Livestock and climate change: what if the key actors in climate change are ... cows, pigs and chickens? *World Watch Magazine*. 22: 10-19.

Gordon, R., Thomas, R. and Foster, A. 2015. The health impact of selective breeding in poultry: a probable case of ‘creeper’ chicken (*Gallus gallus*) from 16th-century Chester, England. *International Journal of Paleopathology*. 9: 1-7.

Grant, H. 2019. Footage reveals chickens in ‘horrifying conditions’ at big UK producer. *The Guardian*. 25 June.

Grayot, J.D. Dual process theories in behavioral economics and neuroeconomics: a critical review. *Review of Philosophy and Psychology*.

Greitemeyer, T. 2009. Effects of songs with prosocial lyrics on prosocial behavior: further evidence and a mediating mechanism. *Personality and Social Psychology Bulletin*. 35: 1500-11.

- Griffin, J. 1986. *Well-Being: Its Meaning, Measurement and Moral Importance*. Oxford: Clarendon Press.
- Grunert, K.G. 2006. Future trends and consumer lifestyles with regard to meat consumption. *Meat Science*. 74: 149-60.
- Gruzalski, B. 2004. Why it's wrong to eat animals raised and slaughtered for food. In *Food for Thought: The Debate over Eating Meat*, ed. S.F. Sapontzis, 124-37. Amherst, NY: Prometheus Books.
- Guéguen, N. 2012. The sweet smell of ... implicit helping: effects of pleasant ambient fragrance on spontaneous help in shopping malls. *The Journal of Social Psychology*. 152: 397-400.
- Gulland, A. and Wilson, J. 2019. Eat less red meat to reduce climate change, UN urges. *The Telegraph*. 8 August.
- Güth, W., Schmittberger, R. and Schwarze, B. 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization*. 3: 367-88.
- Halpern, D. 2015. *Inside the Nudge Unit: How Small Changes can make a Big Difference*. London: WH Allen.
- Hammond, P.J. 1991. Interpersonal comparisons of utility: why and how they are and should be made. In *Interpersonal Comparisons of Well-Being*, ed. J. Elster and J.E. Roemer, 200-54. Cambridge: Cambridge University Press.
- Häne, M., Huber-Eicher, B. and Fröhlich, E. 2000. Survey of laying hen husbandry in Switzerland. *World's Poultry Science Journal*. 56: 21-31.
- Hansson, S.O. and Grüne-Yanoff, T. 2018. Preferences. In *The Stanford Encyclopedia of Philosophy (Summer 2018 Edition)*, ed. E.N. Zalta. <Accessed 05/09/19 at <https://plato.stanford.edu/entries/preferences/>>.
- Hare, R.M. 1981. *Moral Thinking: Its Levels, Method and Point*. Oxford: Oxford University Press.
- Harmston, S. 2017. Over half happy to have meat-free meals. *YouGov Reports*. April 6.
- Harper, G. and Henson, S. 1999. *Consumer Concerns about Animal Welfare and the Impact on Food Choice*. Comparative Focus Groups Report EU Fair CT98-3678. University of Reading.
- Harper, G. and Henson, S. 2001. *Consumer Concerns about Animal Welfare and the Impact on Food Choice*. Final Report EU Fair CT98-3678. Centre for Food Economics Research. University of Reading.
- Harsanyi, J.C. 1986. Individual utilities and utilitarian ethics. In *Paradoxical Effects of Social Behavior: Essays in Honor of Anatol Rapoport*, ed. M. Diekmann and P. Mitter, 1-12. Heidelberg: Physica-Verlag Heidelberg Wien.
- Harsanyi, J.C. 1977a. Morality and the theory of rational behaviour. *Social Research*. 44: 623-56.
- Harsanyi, J.C. 1977b. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge: Cambridge University Press.
- Harvey, D. and Hubbard, C. 2013. Reconsidering the political economy of farm animal welfare: An anatomy of market failure. *Food Policy*. 38: 105-14.

- Harwatt, H., Ripple, W.J., Chaudhary, A., Betts, M.G. and Hayek, M.N. 2019. Scientists call for renewed Paris pledges to transform agriculture. *The Lancet Planetary Health*. <Accessed 12/12/19 at <https://www.thelancet.com/action/showPdf?pii=S2542-5196%2819%2930245-1>>.
- Haslam, S., Kestin, S. and Corry, J. 2004. Reduced welfare in broilers may be associated with increased risk of colonization with *Campylobacter* species. In *Measuring and Auditing Broiler Welfare*, ed. C.A. Weeks and A. Butterworth, 293-94. Oxford: CABI Publishing.
- Haspel, T. 2010. Backyard eggs vs. store-bought: they taste the same. *The Washington Post*. June 2.
- Hausman, D.M. 2012. *Preference, Value, Choice, and Welfare*. Cambridge: Cambridge University Press.
- Hausman, D.M. 2010. Hedonism and welfare economics. *Economics and Philosophy*. 26: 321-44.
- Hausman, D.M. 1995. The impossibility of interpersonal utility comparisons. *Mind*. 104: 473-90.
- Hausman, D.M. and McPherson, M.S. 2010. *Economic Analysis, Moral Philosophy, and Public Policy*. Cambridge: Cambridge University Press.
- Hausman, D.M. and McPherson, M.S. 2009. Preference satisfaction and welfare economics. *Economics and Philosophy*. 25: 1-25.
- Hausman, D.M., McPherson, M.S. and Satz, D. 2017. *Economic Analysis, Moral Philosophy, and Public Policy*. Cambridge: Cambridge University Press.
- Heilbroner, R. L. 2000. *The Worldly Philosophers: The Lives, Times, and Ideas of the Great Economic Thinkers*. London: Penguin Books.
- Heng, Y., Peterson, H.H. and Li, X. 2013. Consumer attitudes towards farm-animal welfare: the case of laying hens. *Journal of Agricultural and Resource Economics*. 38: 418-34.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H. and McElreath, R. 2001. In search of Homo Economicus: behavioral experiments in 15 small-scale societies. *The American Economic Review*. 91: 73-78.
- Henson, R.G. 1971. Utilitarianism and the wrongness of killing. *The Philosophical Review*. 80: 320-37.
- Hicks, J.R. 1940. The valuation of social income. *Economica*. 105-24.
- Hicks, J.R. and Allen, R.G.D. 1934a. A reconsideration of the theory of value. Part I. *Economica*. 1: 52-76.
- Hicks, J.R. and Allen, R.G.D. 1934b. A reconsideration of the theory of value. Part II. *Economica*. 1: 196-219.
- Hidayatno, A. and Komarudin, K. 2009. Eye-tracking research methodology for the planograms of supermarket shelf that draw shopper's visual attention: case study in shampoo packaging. *Proceeding, International Seminar on Industrial Engineering and Management*.
- Higgins, E.T. 1996. Knowledge activation: accessibility, applicability, and salience. In *Social Psychology: Handbook of Basic Principles*, ed. E.T. Higgins and A.W. Kruglanski, 133-68. New York: The Guildford Press.
- Higgins, E.T., Bargh, J.A. and Lombardi, W.J. 1985. Nature of priming effects on categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 11: 59-69.

- High, J. and Bloch, H. 1989. On the history of ordinal utility theory: 1900-1932. *History of Political Economy*. 21: 351-65.
- Hill, R. and Myatt, A. 2007. Overemphasis on perfectly competitive markets in microeconomics principles textbooks. *The Journal of Economic Education*. 38: 58-76.
- Holcombe, R.G. and Sobel, R.S. 2000. Consumption externalities and economic welfare. *Eastern Economic Journal*. 26: 157-170.
- Holländer, H. 1990. A social exchange approach to voluntary cooperation. *The American Economic Review*. 80: 1157-67.
- Hoogland, C.T., de Boer, J. and Boersema, J.J. 2007. Food and sustainability: Do consumers recognize, understand and value on-package information on production standards? *Appetite*. 49: 47-57.
- Hoogland, C.T., de Boer, J. and Boersema, J.J. 2005. Transparency of the meat chain in the light of food culture and history. *Appetite*. 45: 15-23.
- Horsley, J.A., Absalom, K.A.R., Akiens, E.M., Dunk, R.J. and Ferguson, A.M. 2014. The proportion of unhealthy foodstuffs children are exposed to at the checkout of convenience supermarkets. *Public Health Nutrition*. 17: 2453-58.
- Huber, J., Payne, J.W. and Puto, C. 1982. Adding asymmetrically dominated alternatives: violations of regularity and the similarity hypothesis. *Journal of Consumer Research*. 9: 90-98.
- Humane Slaughter Association (HSA). 2014. *A Fowl End?* <Accessed 11/12/19 at <https://www.hsa.org.uk/downloads/related-items/poultry-slaughter.pdf>>.
- IGD. 2007. *Consumer Attitudes to Animal Welfare*. <Accessed 07/11/18 at <https://faunalytics.org/wp-content/uploads/2015/05/Citation374.pdf>>.
- Ince, J. 2010. Food Standards Agency research into consumers' use of country of origin labelling on food. *New Food Magazine*. 30 June.
- Interagency Working Group on Social Cost of Greenhouse Gases, United States Government. 2016. *Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis – Under Executive Order 12866*. <Accessed 19/04/19 at https://www.epa.gov/sites/production/files/2016-12/documents/sc_co2_tsd_august_2016.pdf>.
- International Agency for Research on Cancer (IARC). 2015. IARC Monographs evaluate consumption of red meat and processed meat [press release]. 26 October. <Accessed 10/02/19 at https://www.iarc.fr/wp-content/uploads/2018/07/pr240_E.pdf>.
- Isaac, R.M., McCue, K.F. and Plott, C.R. 1985. Public goods provision in an experimental environment. *Journal of Public Economics*. 26: 51-74.
- Isaac, R.M., Walker, J.M. and Thomas, S.H. 1984. Divergent evidence on free riding: an experimental examination of possible explanations. *Public Choice*. 43: 113-49.
- Isaac, R.M., Walker, J.M. and Williams, A.W. 1994. Group size and the voluntary provision of public goods: experimental evidence utilizing large groups. *Journal of Public Economics*. 54: 1-36.
- Jago, J., Fisher, A. and Le Neindre, P. 2000. Animal welfare and product quality. In *Biological Resource Management Connecting Science and Policy*, ed. E. Balázs, E. Galante, J.M. Lynch, J.S. Schepers, J. Toutant, D. Werner, and P.A.T.J. Werry, 163-71. Berlin: Springer.

- Jansson-Boyd, C. 2018. The Ikea effect: how Ingvar Kamprad's company changed the way we shop. *The Conversation*. 30 January.
- Jiang, Y. Chen, Z. and Wyer Jr., R.S. 2014. Impact of money on emotional expression. *Journal of Experimental Social Psychology*. 55: 228-233.
- Jonsäll, A., Johansson, L., Lundström, K. Andersson, K.H., Nilsen, A.N. and Risvik, E. 2002. Effects of genotype and rearing system on sensory characteristics and preference for pork (*M. Longissimus dorsi*). *Food Quality and Preference*. 13: 73-80.
- Just, D.R. 2011. Behavioral economics and the food consumer. In *The Oxford Handbook of the Economics of Food Consumption and Policy*, ed. J.L. Lusk, J. Roosen and J.F. Shogren, 99-118. Oxford: Oxford University Press.
- Kahneman, D. 2012. *Thinking, Fast and Slow*. London: Penguin Books.
- Kahneman, D., Knetsch, J.L and Thaler, R.H. 1991. Anomalies: the endowment effect, loss aversion, and status quo bias. *The Journal of Economic Perspectives*. 5: 193-206.
- Kahneman, D., Knetsch, J.L. and Thaler, R.H. 1986. Fairness and the assumptions of economics. *The Journal of Business*. 59: S285-S300.
- Kaldor, N. 1939. Welfare propositions of economics and interpersonal comparisons of utility. *The Economic Journal*. 49: 549-52.
- Kaneko, M. and Wooders, M.H. 1994. Widespread externalities and perfectly competitive markets: examples. In *Imperfections and Behavior in Economic Organizations*, ed. R.P. Gilles and P.H.M. Ruys, 71-87. New York: Springer Science+Business Media.
- Kaul, I. 2012. Global public goods: explaining their underprovision. *Journal of International Economic Law*. 15: 729-50.
- Kay, A.C., Wheeler, S.C., Bargh, J.A. and Ross, L. 2004. Material priming: the influence of mundane physical objects on situational construal and competitive behavioral choice. *Organizational Behavior and Human Decision Processes*. 95: 83-96.
- Kemp, M.C. and Penzanis-Christou, P. 1999. Pareto's compensation principle. *Social Choice and Welfare*. 16: 441-444.
- Keser, C. and van Winden, F. 2000. Conditional cooperation and voluntary contributions to public goods. *The Scandinavian Journal of Economics*. 102: 23-39.
- Kjaernes, U. and Lavik, R. 2007. Farm animal welfare and food consumption practices: results from surveys in seven countries. In *Attitudes of Consumers, Retailers and Producers to Farm Animal Welfare, Welfare Quality Report 2*, ed. U. Kjaernes, M. Miele and J Roex, 1-30. Cardiff: School of City and Regional Planning.
- Knowles, T.G. 1994. Handling and transport of spent hens. *World's Poultry Science Journal*. 50: 60-61.
- Knowles, T.G. and Broom, D.M. 1990. The handling and transport of broilers and spent hens. *Applied Animal Behaviour Science*. 28: 75-91.
- Köbberling, V. 2004. Strength of preference and cardinal utility. *Economic Theory*. 27: 375-91.
- Kolstad, C.D. 2011. *Intermediate Environmental Economics*. Oxford: Oxford University Press.
- Koneswaran, G. and Nierenberg, D. 2008. Global farm animal production and global warming: impacting and mitigating climate change. *Environmental Health Perspectives*. 116: 578-582.

- Korthals, M. 2012. Two evils in food country: hunger and lack of representation. In *The Philosophy of Food*, ed. D.M. Kaplan, 103-21. London: University of California Press.
- Kouchaki, M., Smith-Crowe, K. Brief, A.P. and Sousa, C. 2013. Seeing green: mere exposure to money triggers a business decision frame and unethical outcomes. *Organizational Behavior and Human Decision Processes*. 121: 53-61.
- Koutsoyiannis, A. 1980. *Modern Microeconomics*. London: The Macmillan Press.
- Kreps, D.M. 1990. *A Course in Microeconomic Theory*. London: Harvester Wheatsheaf.
- Kunst, J.R. and Hohle, S.M. 2016. Meat eaters by dissociation: how we present, prepare and talk about meat increases willingness to eat meat by reducing empathy and disgust. *Appetite*. 105: 758-74.
- Labour Party. 2019. *Labour's Animal Welfare Manifesto*. <Accessed 23/09/19 at <http://labour.org.uk/wp-content/uploads/2019/08/Labours-Animal-Welfare-Manifesto.pdf>>.
- Lang, T. 2019. No-deal food planning in UK Brexit. *The Lancet*. 394: 814-15.
- Lawler, A. 2014. Chicken of tomorrow. *Aeon*. 5 November.
- Lawrence, F. 2016. If consumers knew how farmed chickens were raised, they might never eat meat again. *The Guardian*. 24 April.
- Ledyard, J.O. 1995. Public goods: a survey of experimental research. In *Handbook of Experimental Economics*, ed. J.H. Kagel and A.E. Roth, 111-94. Princeton: Princeton University Press.
- Leung, K. and Morris, M.W. 2015. Values, schemas and norms in the culture-behavior nexus: a situated dynamics framework. *Journal of International Business Studies*. 46: 1028-50.
- Levitt, T. 2018a. Dairy's 'dirty secret': it's still cheaper to kill male calves than to rear them. *The Guardian*. 26 March.
- Levitt, T. 2018b. Cheap bacon: how shops and shoppers let down our pigs. *The Guardian*. 27 June.
- Lewis, M. 2017. *The Undoing Project: A Friendship that Changed the World*. London: Penguin Books.
- Lieberman, V., Samuels, S.M. and Ross, L. 2004. The name of the game: predictive power of reputations versus situational labels in determining prisoner's dilemma game moves. *Personality and Social Psychology Bulletin*. 30: 1175-85.
- Lodder, P., Ong, H.H., Grasman, R.P.P.P. and Wicherts, J. 2019. A comprehensive meta-analysis of money priming. *Journal of Experimental Psychology: General*. 148: 688-712.
- López-Alt, J.K. 2010. The food lab: do 'better' eggs really taste better? *Serious Eats*. August.
- Loughnan, S., Haslam, N. and Bastian, B. 2010. The role of meat consumption in the denial of moral status and mind to meat animals. *Appetite*. 55: 156-59.
- Lusk, J.L. 2011. The market for animal welfare. *Agriculture and Human Values*. 28: 561-75.
- Lusk, J.L. and Norwood, F.B. 2012. Speciesism, altruism and the economics of animal welfare. *European Review of Agricultural Economics*. 39: 189-212.
- Lusk, J.L. and Norwood, F.B. 2011. Animal welfare economics. *Applied Economic Perspectives and Policy*. 33: 463-83.

- MacKay, A.F. 1986. Extended sympathy and interpersonal utility comparisons. *The Journal of Philosophy*. 83: 305-22.
- Macrae, C.N. and Johnston, L. 1998. Help, I need somebody: automatic action and inaction. *Social Cognition*. 16: 400-17.
- Maguire, E.R. and Monsivais, P. 2015. Socio-economic dietary inequalities in UK adults: an updated picture of key food groups and nutrients from national surveillance data. *British Journal of Nutrition*. 113: 181-89.
- Majot, J. and Kuyek, D. 2017. Big meat and big dairy's climate emissions put Exxon Mobil to shame. *The Guardian*. 7 November.
- Makortoff, K. 2019a. Neil Woodford blocks investors from pulling cash from flagship fund. *The Guardian*. 3 June.
- Makortoff, K. 2019b. Hargreaves Lansdown confronted Neil Woodford in late 2017. *The Guardian*. 19 June.
- Malone, J. 2019. Personal communication. 10 July 2019.
- Malone, T. and Lusk, J.L. 2016. Putting the chicken before the egg price: an ex post analysis of California's battery cage ban. *Journal of Agricultural and Resource Economics*. 41: 518-32.
- Mandler, M. 2006. Cardinality versus ordinality: a suggested compromise. *The American Economic Review*. 96: 1114-36.
- Mann, S. 2005a. Ethological farm programs and the "market" for animal welfare. *Journal of Agricultural and Environmental Ethics*. 18: 369-82.
- Mann, S. 2005b. On institutionalizing animal welfare. In *New Developments in Food Policy, Control and Research*, ed. A.P. Riley, 133-48. New York: Nova Science Publishers, Inc.
- Mansfield, E. 1985. *Microeconomics: Theory and Applications*. London: W.W. Norton & Company.
- Marcus, E. 2005. *Meat Market: Animals, Ethics and Money*. Boston, MA: Brio Press.
- Marwell, G. and Ames, R.E. 1981. Economists free ride, does anyone else? Experiments on the provision of public goods, IV. *Journal of Public Economics*. 15: 295-310.
- Mas-Colell, A., Whinston, M.D. and Green, J.R. 1995. *Microeconomic Theory*. New York: Oxford University Press.
- Mason, J. and Finelli, M. 2006. Brave new farm? In *In Defense of Animals: The Second Wave*, ed. P. Singer, 104-122. Oxford: Blackwell.
- Matheny, G. 2006. Utilitarianism and animals. In *In Defense of Animals: The Second Wave*, ed. P. Singer, 13-25. Oxford: Blackwell.
- Matthews, L.R. 1996. Animal welfare and sustainability of production under extensive conditions: a non-EU perspective. *Applied Animal Behaviour Science*. 49: 41-46.
- Matthews, L.R. and Ladewig, J. 1994. Environmental requirements of pigs measured by behavioural demand functions. *Animal Behaviour*. 47: 713-19.
- Mattick, C. and Allenby, B. 2012. *Cultured meat: the systemic implications of an emerging technology*. 2012 IEEE International Symposium on Sustainable Systems and Technology (ISSST).

- Mattick, C. S., Landis, A. E., Allenby, B. R. and Genovese, N. J. 2015. Anticipatory life cycle analysis of in vitro biomass cultivation for cultured meat production in the United States. *Environmental Science & Technology*. 49: 11941-49.
- May, K.O. 1954. Intransitivity, utility, and the aggregation of preference patterns. *Econometrica*. 22: 1-13.
- McClure, S.M., Li, J., Tomlin, D., Cypert, K.S., Montague, L.M. and Montague, P.R. 2004. Neural correlates of behavioral preference for culturally familiar drinks. *Neuron*. 44: 379-87.
- McConnell, K.E. 1997. Does altruism undermine existence value? *Journal of Environmental Economics and Management*. 32: 22-37.
- McEachern, M.G. and Schroeder, M.J.A. 2004. Integrating the voice of the consumer within the value chain: a focus on value-based labelling communications in the fresh-meat sector. *Journal of Consumer Marketing*. 21: 497-509.
- McEachern, M.G. and Seaman, C. 2005. Consumer perceptions of meat production: enhancing the competitiveness of British agriculture by understanding communication with the consumer. *British Food Journal*. 107: 572-93.
- McEachern, M.G. and Warnaby, G. 2004. Retail 'quality assurance' labels as a strategic marketing communication mechanism for fresh meat. *The International Review of Retail, Distribution and Consumer Research*. 14: 255-71.
- McEachern, M.G., Schröder, M.J.A., Willock, J., Whitelock, J. and Mason, R. 2007. Exploring ethical brand extensions and consumer buying behaviour: the RSPCA and the "Freedom Food" brand. *Journal of Product & Brand Management*. 16: 168-77.
- McInerney, J. 2004. *Animal Welfare, Economics, and Policy*. London: DEFRA.
- McInerney, J. 1991. A socio-economic perspective on animal welfare. *Outlook on Agriculture*. 20: 51-56.
- McInerney, J. 1998. The economics of welfare. In *Ethics, Welfare, Law and Market Forces: The Veterinary Interface*, ed. A.R. Michell and R. Ewbank, 115-32. Wheathampstead, UK: Universities Federation for Animal Welfare.
- McMullen, S. 2016a. *Animals and the Economy*. London: Palgrave Macmillan.
- McMullen, S. 2016b. An ethical consumer capitalism. In *The Future of Meat without Animals*, ed. B. Donaldson and C. Carter, 35-48. Plymouth: Rowman and Littlefield International.
- McVittie, A., Moran, D., Sandilands, V. and Spaks, N. 2006. *Estimating Non-Market Benefits of Reduced Stocking Density and Other Welfare Increasing Measures for Meat Chickens in England*. London: DEFRA.
- Meade, J.E. 1952. External economies and diseconomies in a competitive situation. *The Economic Journal*. 62: 54-67.
- Miele, M. 2010. Report concerning consumer perceptions and attitudes towards farm animal welfare. *Official Experts Report EAWP (Task 1.3)*. Uppsala, Sweden: Uppsala University.
- Milman, O. 2018. Why eating less meat is the best thing you can do for the planet in 2019. *The Guardian*. 21 December.
- Mintel. 2018. *Added Value in Dairy Drinks, Milk and Cream*. UK: Mintel.
- Mintel. 2017a. *Processed Poultry and Red Meat*. UK: Mintel.

- Mintel. 2017b. *Unprocessed Poultry and Red Meat*. UK: Mintel.
- Mintel. 2017c. *Cheese*. UK: Mintel.
- Mintel. 2017d. *Fish and Shellfish*. UK: Mintel.
- Mintel. 2017e. *Meat-Free Foods*. UK: Mintel.
- Mishan, E.J. 1988. *Cost-Benefit Analysis*. London: Unwin Hyman.
- Mitchelson, A. 2018. Ikea's global chief designer reveals why the furniture giant's stores are built like mazes. *Smart Company*. 30 August.
- Molinsky, A.L., Grant, A.M. and Margolis, J.D. 2012. The bedside manner of homo economicus: how and why priming an economic schema reduces compassion. *Organizational Behavior and Human Decision Processes*. 119: 27-37.
- Moran, D. and McVittie, A. 2008. Estimation of the value the public places on regulations to improve broiler welfare. *Animal Welfare*. 17: 43-52.
- Morris, R. and Cook, J. 2018. Would you eat slaughter-free meat? *BBC News*. 15 October.
- Moser, R., Raffaelli, R. and Thilmany-McFadden, D. 2011. Consumer preferences for fruit and vegetables with credence-based attributes: a review. *International Food and Agribusiness Management Review*. 14: 121-42.
- Moss-Racusin, C.A., Dovidio, J.F., Brescoll, V.L., Graham, M.J. and Handelsman, J. 2012. Science faculty's subtle gender biases favor male students. *PNAS*. 41: 16474-16479
- Mullally, C. and Lusk, J.L. 2018. The impact of farm animal housing restrictions on egg prices, consumer welfare, and production in California. *American Journal of Agricultural Economics*. 100: 649-69.
- Musgrave, R.A. 1959. *The Theory of Public Finance: A Study in Public Economy*. London: McGraw-Hill.
- Napolitano, F., Girolami, A. and Braghieri, A. 2010. Consumer liking and willingness to pay for high welfare animal-based products. *Trends in Food Science & Technology*. 21: 537-43.
- Narveson, J. 2003. The "invisible hand". *Journal of Business Ethics*. 46: 201-12.
- Nechyba, T.J. 2017. *Microeconomics: An Intuitive Approach with Calculus*. Boston, MA: CENGAGE Learning.
- Nelson, P. 1970. Information and consumer behavior. *Journal of Political Economy*. 78: 311-29.
- Nicholson, W. 2010. *Intermediate Microeconomics and its Applications*. Orlando, FL: Dryden Press.
- Nicholson, W. and Snyder, C. 2010. *Intermediate Microeconomics and its Applications*. Boston, MA: CENGAGE Learning.
- Nicol, C. 2018. Feather pecking and cannibalism: can we really stop beak trimming? In *Advances in Poultry Welfare*, ed. J.A. Mench, 175-97. Duxford: Woodhead Publishing.
- Nitschke, J.B., Dixon, G.E., Sarinopoulos, I., Short, S.J., Cohen, J.D., Smith, E.E., Kosslyn, S.M., Rose, R.M. and Davidson, R.J. 2006. Altering expectancy dampens neural response to aversive taste in primary taste cortex. *Nature Neuroscience*. 9: 435-42.

- Noor, P. 2019. Caroline Lucas is wrong – a meat tax would only hurt the poor. *The Guardian*. 7 January.
- Norman, D.A. 2013. *The Design of Everyday Things*. London: The MIT Press.
- Norwood, F.B. and Lusk, J.L. 2013. Animal welfare and food safety. *Food Safety Magazine*, February/March.
- Norwood, F.B. and Lusk, J.L. 2011. *Compassion, by the Pound: The Economics of Farm Animal Welfare*. Oxford: Oxford University Press.
- Nussbaum, M.C. 2001. Symposium on Amartya Sen's philosophy: 5 adaptive preferences and women's options. *Economics and Philosophy*. 17: 67-88.
- O'Dwyer, G. 2014. Swedish pig farmers call for state support. *Global Meat News*. 3 March.
- Olivola, C.Y. and Sussman, A.B. 2015. Taxes and consumer behaviour. In *The Cambridge Handbook of Consumer Psychology*, ed. M.I. Norton, D.D. Rucker and C. Lambertson, 564-88. Cambridge: Cambridge University Press.
- Oxfam. 2020. Poverty-Busting Pig. <Accessed 22/01/20 at <https://www.oxfam.org.uk/shop/oxfam-unwrapped/charity-gift-card/poverty-busting-pig-ouio61ml>>.
- Pareto, V. 2008. The maximum of utility given by free competition. *Giornale degli Economisti e Annali di Economica*. 67: 387-403.
- Parfit, D. 1987. *Reasons and Persons*. Oxford: Oxford University Press.
- Parkin, M. 2000. *Microeconomics*. Reading, MA: Addison-Wesley.
- Passantino, A., Conte, F. and Russo, M. 2008. Animal welfare labelling and the approach of the European Union: an overview on the current situation. *Journal for Consumer Protection and Food Safety*. 3: 396-99.
- Penn, A. 2011. Who enjoys shopping in IKEA? (18 Jan 2011). <Accessed 02/12/19 at https://www.youtube.com/watch?v=NkePRXxH9D4&feature=emb_title>.
- Perloff, J.M. 1999. *Microeconomics*. Reading, MA: Addison-Wesley.
- Peters, S. 2012. *The Chimp Paradox: The Mind Management Programme to Help You Achieve Success, Confidence and Happiness*. London: Vermilion.
- Pettersson, I.C., Weeks, C.A., Wilson, L.R.M. and Nicol, C.J. 2016. Consumer perceptions of free-range laying hen welfare. *British Food Journal*. 118: 1999-2013.
- Pigou, A.C. 2009. *The Economics of Welfare*. London: Transaction Publishers.
- Pindyck, R.S. and Rubinfeld, D.L. 2013. *Microeconomics*. London: Pearson.
- Pollan, M. 2011. *The Omnivore's Dilemma: The Search for a Perfect Meal in a Fast-Food World*. London: Bloomsbury.
- Posner, R.A. 1969. Natural monopoly and its regulation. *Stanford Law Review*. 21: 548-643.
- Post, M. 2012. Cultured meat from stem cells: challenges and prospects. *Meat Science*. 92: 297-301.
- Prest, A.R. and Turvey, R. 1965. Cost-benefit analysis: a survey. *The Economic Journal*. 75: 683-735.

- Primatt, H. 1776. *A Dissertation on the Duty of Mercy and Sin of Cruelty to Brute Animals*. London: R. Hett.
- Proudfoot, F.G., Hulan, H.W. and Ramey, D.R. 1979. The effect of four stocking densities on broiler carcass grade, the incidence of breast blisters, and other performance traits. *Poultry Science*. 58: 791-93.
- Pulsipher, A.G. 1971. The properties and relevancy of merit goods. *FinanzArchive / Public Finance Analysis*. 30: 266-86.
- Qizilbash, M. 1998. The concept of well-being. *Economics and Philosophy*. 14: 51-73.
- Raha, R. 2006. Animal liberation: an interview with Professor Peter Singer. *The Vegan*. Autumn.
- Railton, P. 1986. Facts and values. *Philosophical Topics*. 14: 5-31.
- Reese, J. 2018. *The End of Animal Farming: How Scientists, Entrepreneurs, and Activists are Building an Animal-Free Food System*. Boston, MA: Beacon Press.
- Regan, T. 2004. *The Case for Animal Rights*. Berkeley, L.A.: University of California Press.
- Rege, M. 2004. Social norms and private provision of public goods. *Journal of Public Economic Theory*. 6: 65-77.
- Reiss, J. 2013. *Philosophy of Economics: A Contemporary Introduction*. Oxford: Routledge.
- Ricke, K., Drouet, L., Caldeira, K. and Tavoni, M. 2018. Country-level social cost of carbon. *Nature Climate Change*. 8: 895-900.
- Rigby, D. 2019. Private communication at *Non Market Valuation: FSA Training Workshop*. Food Standards Agency. Delivered 04/07/2019.
- Ritter, T. 2017. *Don't Believe the Carbon Market Hype: Why States Should Not Pursue Carbon Markets and What They Can Do Instead*. Institute for Agriculture and Trade Policy.
- Robbins, L. 1932. *An Essay on the Nature and Significance of Economic Science*. London: MacMillan and Co.
- Roberts, R.D. 1987. Financing public goods. *Journal of Political Economy*. 95: 420-37.
- Rollin, B. 2008. The ethics of agriculture: the end of true husbandry. In *The Future of Animal Farming: Renewing the Ancient Contract*, ed. M.S. Dawkins and R. Bonney, 7-19. Oxford: Blackwell.
- Rothgerber, H. 2014. Efforts to overcome vegetarian-induced dissonance among meat eaters. *Appetite*. 79: 32-41.
- Rowlands, M. 2009. *Animal Rights: Moral Theory and Practice*. Basingstoke: Palgrave Macmillan.
- Rowlands, M. 2002. *Animals Like Us*. London: Verso.
- Royal Society for Public Health (RSPH) and Slimming World. 2019. *Health on the Shelf*. <Accessed 20/11/19 at <https://www.rsph.org.uk/uploads/assets/uploaded/5ec3d502-2e94-47c0-ad1bbeid4a74218c.pdf>>.
- RSPCA. 2018. *Into the Fold: Targeted Financial Support to Improve Farm Animal Welfare*. West Sussex: RSPCA.

- RSPCA. 2017a. *RSPCA Welfare Standards for Laying Hens*. <Accessed 25/01/19 at <https://science.rspca.org.uk/ImageLocator/LocateAsset?asset=document&assetId=1232741072801&mode=prd>>.
- RSPCA. 2017b. *RSPCA Welfare Standards for Meat Chickens*. <Accessed 23/01/19 at <https://science.rspca.org.uk/ImageLocator/LocateAsset?asset=document&assetId=1232740881600&mode=prd>>.
- RSPCA. 2006. *Everyone's a Winner: How Rearing Chickens to Higher Welfare Standards can Benefit the Chicken, Producer, Retailer and Consumer*. <Accessed 14/03/19 at <https://www.rspca.org.uk/documents/1494939/7712578/Everyone%27s+a+winner+chicken+report+%28PDF+590KB%29.pdf/d7079dbf-30a5-f240-285e-1086668fo2e9?t=1554205296757>>.
- RSPCA Assured. 2018. *About Us*. <Accessed 27/01/20 at <https://www.rspcaassured.org.uk/about-us/>>.
- RSPCA Assured. 2017a. *Annual Review 2017*. <Accessed 24/01/19 at <https://www.berpscaassured.org.uk/media/1293/rspca-assured-annual-review-2017.pdf>>.
- RSPCA Assured. 2017b. *Membership Fees*. <Accessed 10/11/19 at <https://www.berpscaassured.org.uk/media/1268/membership-fees-nov-2017.pdf>>.
- Russell, B. 2004. *History of Western Philosophy*. Oxford: Routledge.
- Ryder, R.D. 2000. *Animal Revolution: Changing Attitudes Towards Speciesism*. Oxford: Berg.
- Ryland, D. 2014. Advancing animal welfare: state, society and economy. In *Democratic and Social State under the Rule of Law: Society, Politics, Economy*, ed. A.F. Modrzewski, 851-67. Krakow: Krakow University and the Jagiellonian University.
- Saad, G. 2011. *The Consuming Instinct: What Juicy Burgers, Ferraris, Pornography, and Gift Giving Reveal about Human Nature*. New York: Prometheus Books.
- Safran Foer, J. 2010. *Eating Animals*. London: Penguin Books.
- Sagoff, M. 2001. At the shrine of our lady of Fatima, or why political questions are not all economic. In *Environmental Ethics: Readings in Theory and Application*, ed. L.P. Pojman, 467-73. Belmont, CA: Wadsworth.
- Salvatore, D. 2003. *Microeconomics: Theory and Applications*. Oxford: Oxford University Press.
- Sandel, M. 2013. *What Money Can't Buy: The Moral Limits of Markets*. London: Penguin.
- Santos, A.C., Roberts, J.A., Cook, A.J., Simons, R., Sheehan, R., Lane, C., Adak, G.K., Clifton-Hadley, F.A. and Rodrigues, L.C. 2010. Salmonella Typhimurium and Salmonella Enteritidis in England: costs to patients, their families, and primary and community health services of the NHS. *Epidemiology & Infection*. 139: 742-53.
- Saul, H. 2015. Hatched, discarded, gassed: what happens to male chicks in the UK. *The Independent*. 5 March.
- Say, J.B. 1971. *A Treatise on Political Economy*. Philadelphia: Augustus M. Kelley Publishers.
- Scanlon, T.M. 1991. The moral basis of interpersonal comparisons. In *Interpersonal Comparisons of Well-Being*, ed. J. Elster and J.E. Roemer, 17-44. Cambridge: Cambridge University Press.
- Schimmack, U., Heene, M. and Kesavan, K. 2017. *Reconstruction of a train wreck: how priming research went off the rails*. <Accessed 02/12/19 at

<https://replicationindex.com/2017/02/02/reconstruction-of-a-train-wreck-how-priming-research-went-of-the-rails/>>.

Schotter, A.R. 1997. *Microeconomics: A Modern Approach*. Reading, MA: Addison-Wesley.

Schröder, M.J.A. and McEachern, M.G. 2004. Consumer value conflicts surrounding ethical food purchase decisions: a focus on animal welfare. *International Journal of Consumer Studies*. 28: 168-77.

Schroeder, T.C. and Tonsor, G.T. 2011. Demand for meat quality attributes. In *The Oxford Handbook of the Economics of Food Consumption and Policy*, ed. J.L. Lusk, J. Roosen and J.F. Shogren, 791-810. Oxford: Oxford University Press.

Schulz, F. 2020. Could Germany impose a tax on meat? *Euractiv.com*. 10 January.

Schumpeter, J.A. 2006. *History of Economic Analysis*. London: Routledge.

Schroeder, J.E. 2004. Visual consumption in an image economy. In *Elusive Consumption*, ed. K. Ekstrom and H. Brembeck, 229-244. Oxford: Berg.

Schwartz, B. 2005. *The Paradox of Choice: Why More is Less*. New York: HarperCollins.

Sen, A. 2001. *Development as Freedom*. Oxford: Oxford University Press.

Sen, A. 1988. Freedom of choice: concept and content. *European Economic Review*. 32: 269-94.

Sen, A. 1979. Interpersonal comparisons of welfare. In *Economics and Human Welfare: Essays in Honour of Tibor Scitovsky*, ed. M.J. Boskin, 183-202. London: Academic Press.

Shankar, V., Inman, J.J., Mantrala, M., Kelley, E. and Rizley, R. 2011. Innovations in shopper marketing: current insights and future research issues. *Journal of Retailing*. 87S: S29-S42.

Shapiro, P. 2018. *Clean Meat: How Growing Meat Without Animals Will Revolutionize Dinner and the World*. New York: Gallery Books.

Singer, P. 2015a. *Animal Liberation*. London: The Bodley Head.

Singer, P. 2015b. *The Most Good You Can Do: How Effective Altruism Is Changing Ideas About Living Ethically*. London: Yale University Press.

Singer, P. 2011. *The Expanding Circle: Ethics, Evolution, and Moral Progress*. Oxford: Princeton University Press.

Singer, P. and Mason, J. 2006. *The Ethics of What We Eat: Why Our Food Choices Matter*. United States: Rodale.

Sinha, R., Cross, A.J., Graubard, B.I., Leitzmann, M.F. and Schatzkin, A. 2009. Meat intake and mortality: a prospective study of over half a million people. *Archives of Internal Medicine*. 169: 562-71.

Smith, A. 1999. *The Wealth of Nations: Books I-III*. London: Penguin Classics.

Smith, A. 2002. *The Theory of Moral Sentiments*. Cambridge: Cambridge University Press.

Snyder, C. and Nicholson, W. 2012. *Microeconomic Theory: Basic Principles and Extensions*. Mason, OH: South-Western CENGAGE Learning.

Soil Association. 2019. *Certification Fees for Food Businesses to Soil Association or EU Organic Standards*. <Accessed 10/11/19 at <https://www.soilassociation.org/media/18512/fees-for-uk-food-businesses-2019-20.pdf>>

- Solotaroff, P. 2013. Animal cruelty is the price we pay for cheap meat. *Rolling Stone*. 10 December.
- Sonesson, U., Davis, J. and Ziegler, F. 2009. *Food Production and Emissions of Greenhouse Gases: An Overview of the Climate Impact of Different Product Groups*. Swedish Institute for Food and Biotechnology.
- Springmann, M., Mason-D'Croz, D., Robinson, S., Wiebe, K., Godfray H.C.J., Rayner, M. and Scarborough, P. 2018a. Health-motivated taxes on red and processed meat: a modelling study on optimal tax levels and associated health impacts. *PLoS ONE*. 13: 1-16.
- Springmann, M., Mason-D'Croz, D., Robinson, S., Wiebe, K., Godfray H.C.J., Rayner, M. and Scarborough, P. 2018b. Health-motivated taxes on red and processed meat: a modelling study on optimal tax levels and associated health impacts. Appendix. <Accessed 10/02/19 at <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0204139>>.
- Springmann, M., Mason-D'Croz, D., Robinson, S., Wiebe, K., Godfray H.C.J., Rayner, M. and Scarborough, P. 2016a. Mitigation potential and global health impacts from emissions pricing of food commodities. *Nature Climate Change*. 7: 69-76.
- Springmann, M., Mason-D'Croz, D., Robinson, S., Wiebe, K., Godfray H.C.J., Rayner, M. and Scarborough, P. 2016b. Mitigation potential and global health impacts from emissions pricing of food commodities. Supplementary Information. <Accessed 19/04/19 at <https://www.nature.com/articles/nclimate3155#abstract>>.
- Średnicka-Tober, D., Barański, M., Seal, C., Sanderson, R., Benbrook, C., Steinshamn, H., Gromadzka-Ostrowska, J., Rembiałkowska, E., Skwarło-Sońta, K., Eyre, M., Cozzi, G., Krogh Larsen, M., Jordon, T., Niggli, U., Sakowski, T., Calder, P.C., Burdge, G.C., Sotiraki, S., Stefanakis, A., Stergiadis, S., Yolcu, H., Chatzidimitriou, E., Butler, G., Stewart, G. and Leifert, C. 2016a. Higher PUFA and η -3 PUFA, conjugated linoleic acid, α -tocopherol and iron, but lower iodine and selenium concentrations in organic milk: a systematic literature review and meta- and redundancy analyses. *British Journal of Nutrition*. 115: 1043-60.
- Średnicka-Tober, D., Barański, M., Seal, C., Sanderson, R., Benbrook, C., Steinshamn, H., Gromadzka-Ostrowska, J., Rembiałkowska, E., Skwarło-Sońta, K., Eyre, M., Cozzi, G., Krogh Larsen, M., Jordon, T., Niggli, U., Sakowski, T., Calder, P.C., Burdge, G.C., Sotiraki, S., Stefanakis, A., Yolcu, H., Stergiadis, S., Chatzidimitriou, E., Butler, G., Stewart, G. and Leifert, C. 2016b. Composition differences between organic and conventional meat: a systematic literature review and meta-analysis. *British Journal of Nutrition*. 115: 994-1011.
- Statista. 2018. *Average time spent per shopping visit in Great Britain between 2012 and 2013, by location*. <Accessed 20/11/19 at <https://www.statista.com/statistics/412514/shopping-time-spent-by-location-great-britain-uk/>>.
- Stevenson, P. 2018. *A Better Brexit for Farm Animals: What the Government must do to Protect Welfare Standards*. Food Research Collaboration Brexit Briefing.
- Stevenson, P. 2015. *The Impact of World Trade Organisation Rules on Animal Welfare*. <Accessed 28/07/19 at <https://www.ciwf.org.uk/media/7131790/the-impact-of-the-world-trade-organisation-rules-on-animal-welfare-2015.pdf>>.
- Stevenson, P. 2011. *Reviewing the Costs: The Economics of Moving to Higher Welfare Farming*. Surrey: Compassion in World Farming.
- Stiglitz, J. E. 1991. *The Invisible Hand and Modern Welfare Economics*. Cambridge, MA: National Bureau of Economic Research.

- Stiglitz, J. E., Sen, A., and Fitoussi, J. 2010. *Mismeasuring Our Lives: Why GDP Doesn't Add Up*. New York: The New Press.
- Strotz, R.H. 1953. Cardinal utility. *The American Economic Review*. 43: 384-97.
- Styron, W. 2004. *Sophie's Choice*. London: Vintage Books.
- Sumner, D.A., Gow, H., Hayes, W., Matthews, W., Norwood, B., Rosen-Molina, J.T. and Thurman, W. 2011. Economic and market issues on the sustainability of egg production in the United States: analysis of alternative production systems. *Poultry Science*. 90: 241-50.
- Tam, C.C. and O'Brien, S.J. 2016. Economic cost of *campylobacter*, norovirus and rotavirus disease in the United Kingdom. *PloS One*. 11.
- Tasker, J. 2018. Better animal welfare should be rewarded – Gove. *Farmer's Weekly*. 20 February.
- Thaler, R.H. 2016. *Misbehaving: The Making of Behavioural Economics*. London: Penguin Books.
- Thaler, R.H. and Sunstein, C.R. 2009. *Nudge: Improving Decisions about Health, Wealth and Happiness*. London: Penguin Books.
- The Welfare of Farmed Animals (England) Regulations 2007*.
- Thomas, K. 1984. *Man and the Natural World: Changing Attitudes in England 1500-1800*. London: Penguin Books.
- Thomassim, P.J. 2013. Canadian agriculture and the development of a carbon trading and offset system. *American Journal of Agricultural Economics*. 5: 1171-77.
- Tiffin, R., Balcombe, K., Salois, M. and Kehlbacher, A. 2011. *Estimating Food and Drink Elasticities*. University of Reading.
- Toma, L., Scott, A.W., Revoredo-Giha, C. and Kupiec-Teahan, B. 2012. Consumers and animal welfare. A comparison between European Union countries. *Appetite*. 58: 597-607.
- Tonsor, G.T. and Wolf, C.A. 2011. On mandatory labeling of animal welfare attributes. *Food Policy*. 36: 430-37.
- Trivers, R.L. 1971. The evolution of reciprocal altruism. *The Quarterly Review of Biology*. 46: 35-57.
- Trumbull, W.N. 1990. Who has standing in cost-benefit analysis? *Journal of Policy Analysis and Management*. 9: 201-18.
- Tucker, I. 2011. Adrift in a shopping maze: it's a successful no-exit strategy. *The Guardian*. 30 January.
- Tversky, A. 1969. Intransitivity of preferences. *Psychological Review*. 76: 31-48.
- Tversky, A. and Simonson, I. 1993. Context-dependent preferences. *Management Science*. 39: 1179-89.
- UK Parliament. 2018. *Agriculture Bill [HC]*. Bill 226. 2017-19.
- UK Parliament. 2007. *The Welfare of Farmed Animals (England) Regulations 2007 (SI 2007/2078)*.

- UK Parliament. House of Lords European Union Committee. 2017. *Brexit: Farm Animal Welfare*. HL 2017-2019 (15). London: The Stationery Office.
- United Nations Human Rights Officer of the High Commissioner. 1966. *International Covenant on Economic, Social and Cultural Rights*. <Accessed 15/05/19 at <https://www.ohchr.org/en/professionalinterest/pages/cescr.aspx>>.
- Urry, J. 2000. *Consuming Places*. London: Routledge.
- Ursinus, W.W., Schepers, F., de Mol, R.M., Bracke, M.B.M., Metz, J.H.M. and Groot Koerkamp, P.W.G. 2009. COWEL: a decision support system to assess welfare of husbandry systems for dairy cattle. *Animal Welfare*. 18: 545-52.
- van Mensvoort, K. and Grievink, H.-J. 2014. *The In Vitro Meat Cookbook*. Amsterdam: BIS Publishers.
- van Tongeren, D.R., Hibbard, R., Edwards, M., Johnson, E., Diepholz, K., Newbound, H., Shay, A., Houpt, R., Cairo, A. and Green, J.D. 2018. Heroic helping: the effects of priming superhero images on prosociality. *Frontiers in Psychology*. 9: 2243.
- Varian, H.R. 2010. *Intermediate Microeconomics*. New York: Norton.
- Varian, H.R. 1992. *Microeconomic Analysis*. New York: W.W. Norton & Company.
- Ventin, M. Forthcoming. *Animal-Free Animal Products: Cultivated Animal Product Technologies and Society*. London: Food Standards Agency.
- Verbeke, W. and Viaene, J. 1999. Consumer attitude to beef quality labeling and associations with beef quality labels. *Journal of Food & Agribusiness Marketing*. 10: 45-65.
- Verbeke, W., Pérez-Cueto, F.J.A., de Barcellos, M.D., Krystallis, A. and Grunert, K.G. 2010. European citizen and consumer attitudes and preferences regarding beef and pork. *Meat Science*. 84: 284-92.
- Verplanken, B. and van Knippenberg, A. 1998. Predicting behavior from actions in the past: repeated decision making or a matter of habit? *Journal of Applied Social Psychology*. 28: 1355-74.
- Vetter, S., Vasa, L. and Ózsvári, L. 2014. Economic aspects of animal welfare. *Acta Polytechnica Hungarica*. 11: 119-3.
- Voegler, C.S. 2017. Farm animal welfare policy in comparative perspective: determinants of cross-national differences in Austria, Germany, and Switzerland. *Contemporary Issues of Policy-Making Across Europe*. 3: 20-47.
- Vohs, K.D. 2015. Money priming can change people's thoughts, feelings, motivations, and behaviors: an update on 10 years of experiments. *Journal of Experimental Psychology: General*. 144: e86-e93.
- Vohs, K.D., Mead, N.L. and Goode, M.R. 2008. Merely activating the concept of money changes personal and interpersonal behavior. *Current Directions in Psychological Science*. 17: 208-12.
- Vohs, K.D., Mead, N.L. and Goode, M.R. 2006. The psychological consequences of money. *Science*. 314: 1154-56.
- Voogd, E.L. 2009. Does animal welfare affect food safety? *Food Safety Magazine*. February/March.

- Waitrose and Partners. 2018. *Food and Drink Report 2018-19: The Era of the Mindful Consumer*. <Accessed 22/02/19 at <https://www.waitrose.com/content/dam/waitrose/Inspiration/Waitrose%20&%20Partners%20oFood%20and%20Drink%20Report%202018.pdf>>.
- Wang, L., Malhotra, D. and Murnighan, J.K. 2011. Economics education and greed. *Academy of Management Learning & Education*. 10: 643-60.
- Wasley, A., Harvey, F., Davies, M. and Child, D. 2017. UK has nearly 800 livestock mega farms, investigation reveals. *The Guardian*. 17 July.
- Webster, J. 2016. Animal welfare: freedoms, dominions and “a life worth living”. *Animals*. 6: 35.
- Wedell, D.H. and Pettibone, J.C. 1996. Using judgements to understand decoy effects in choice. *Organizational Behavior and Human Decision Processes*. 67: 326-44.
- Wellesley, L., Happer, C. and Froggatt, A. 2015. *Changing Climate, Changing Diets: Pathways to Lower Meat Consumption*. London: Chatham House, The Royal Institute of International Affairs.
- Whole Foods Market. 2019. 5-step animal welfare rating. <Accessed 08/11/19 at <https://www.wholefoodsmarket.com/mission-values/animal-welfare/5-step-animal-welfare-rating>>.
- Wirth, F.F., Stanton, J.L. and Wiley, J.B. 2011. The relative importance of search versus credence product attributes: organic and locally grown. *Agricultural and Resource Economics Review*. 40: 48-62.
- Wollaston, W. 1725. *The Religion of Nature Delineated*. London: S. Palmer.
- World Animal Protection. 2019. *Dairy farming in the UK*. <Accessed 11/12/19 at <https://www.worldanimalprotection.org.uk/campaigns/animals-farming/dairy-farming-in-uk>>.
- World Resources Institute. 2019. *Creating a Sustainable Future: A Menu of Solutions to Feed Nearly 10 Billion People by 2050*. Washington, D.C.: World Resources Institute.
- World Society for the Protection of Animals (WSPA). 2014. *How protecting farm animal welfare can boost jobs and livelihoods*. <Accessed 30/10/19 at https://d31j74p4lp4lpxfp.cloudfront.net/sites/default/files/int_files/protect-farm-animal-welfare-boost-jobs-livelihoods.pdf>.
- World Trade Organisation (WTO). 2019. *Principles of the Trading System*. <Accessed 28/07/19 at https://www.wto.org/english/thewto_e/whatis_e/tif_e/fact2_e.htm>.
- WTO. 2014. *The WTO Agreements Series: Technical Barriers to Trade*. Geneva: World Trade Organisation.
- WTO. 1995. *Agreement on Technical Barriers to Trade*. <Accessed 23/09/19 at https://www.wto.org/english/docs_e/legal_e/17-tbt.pdf>.
- WTO. 1986. *General Agreement on Tariffs and Trade*. Geneva: World Trade Organisation.
- Xiao, S.H. and Nicholson, M. 2013. A multidisciplinary cognitive behavioural framework of impulse buying: a systematic review of the literature. *International Journal of Management Reviews*. 15: 333-56.

- Xin, H., Gates, R.S., Green, A.R., Mitloehner, F.M., Moore Jr., P.A. and Wathes, C.M. 2011. Environmental impacts and sustainability of egg production systems. *Poultry Science*. 90: 263-77.
- Yeates, J.W. 2017. How good? Ethical criteria for a 'good life' for farm animals. *Journal of Agricultural and Environmental Ethics*. 30: 23-35.
- Young, H.P. 2015. The evolution of social norms. *Annual Review of Economics*. 7: 359-87.
- Zambasri, I. 2017. The nudge – how IKEA's store layout design influences your spending. <Accessed 02/12/19 at <https://wp.nyu.edu/thoughtsonwayfinding/2017/11/19/the-nudge-how-ikeas-store-layout-design-influences-your-spending/>>.
- Zickfeld, J.H., Kunst, J.R. and Hohle, S.M. 2018. Too sweet to eat: exploring the effects of cuteness on meat consumption. *Appetite*. 120: 181-95.
- Zuev, D. 2015. Cities, visual consumption of. In *The Wiley Blackwell Encyclopedia of Consumption & Consumer Studies*, ed. D.T. Cook and J.M. Ryan, 79-80. Oxford: Wiley.