

The consequences of differing car insurance premiums

Abstract

Insurance companies use a range of factors to set car insurance premiums, some of which have unwelcome consequences. Much of the relevant philosophical and legal literature has focussed on when these factors constitute discrimination, and has sought non-consequentialist grounds to justify when discrimination is wrong. In this dissertation, I support Lippert-Rasmussen's argument for consequentialism (2014). However, I argue that we should consider all consequences of differential pricing, and not merely those that constitute discrimination. I argue that such an analysis is pragmatic, allowing us to appeal to those with consequentialist and non-consequentialist ethical perspectives, and to more effectively recognise and reduce harms. I explore a range of potential consequences of different mechanisms for pricing car insurance: increased implementation costs, susceptibility to cheating, reduction in privacy, reduction in insurance, changes in behaviour, economic inequality, stigma and fairness. I argue that these consequences should be taken into account by companies, society, and regulators.

Table of Contents

Abstract.....	1
1 Introduction	3
2 How car insurance works	4
3 Justifying a consequence-based analysis of car insurance pricing	6
4 Identification and consideration of the consequences of car insurance pricing	10
5 Responding to the consequences of differential car insurance pricing	14
6 Considering two car insurance pricing controversies	16
7 Conclusion.....	18
References	20

1 Introduction

It is commonly accepted that insurance companies offering insurance against the costs of car accidents and theft need not charge the same premium for customers with different risks (Heath, 2014, p345, Scholes, 2017, p112). Their business therefore depends on the assessment of appropriate premiums for individual customers, high enough to cover expected losses but low enough to avoid losing profitable business. However, this freedom to differentiate can disguise harmful discrimination.

In this dissertation, I consider when differential car insurance premiums are acceptable. I endorse the common intuition that some forms of differential pricing, for example those on purely racial grounds, are unacceptable, and are rightly outlawed in many jurisdictions. However, there are many cases that are less clear-cut. I focus on two real world examples in greater detail later. One is the question of charging some drivers more for their insurance on the basis of them being male, a practice outlawed by the European Court of Justice in 2012 (Rego, 2015). A second example is that of increasing premiums for cars owned by residents of poorer postcodes or neighbourhoods (O'Neil, 2017, p164). Assessing cases like these is of growing importance: as we automate decision-making and pricing, harmful mechanisms risk becoming entrenched unless we actively discourage them.

The acceptability of differential treatment is discussed in the philosophical literature, most notably by Kasper Lippert-Rasmussen (2014). However, I argue that much of the discussion is unsatisfactory, focussing unduly on whether instances of differential treatment fall within the definition of discrimination. Firstly, there is no clear, agreed definition of discrimination. Secondly, most philosophers accept that discrimination is not necessarily wrong, forcing us to use other criteria (eg its consequences) to determine its permissibility. Thirdly, differential pricing may be unacceptable despite not meeting the definition of discrimination.

My argument is that Lippert-Rasmussen is right to argue for consequentialism, that we should make decisions that lead to the best consequences. I also welcome his acceptance that consequentialism can incorporate notions of distribution and justice. However, I deviate from Lippert-Rasmussen by arguing that we should consider all forms of differential pricing, and not merely cases that meet the definition of discrimination. I argue that this will help us more easily evaluate controversial cases, including those that have beneficial as well as harmful consequences. I also argue that it provides a more pragmatic way of gaining consensus for effectively reducing harm.

There is a risk to this approach. By considering borderline cases of differential pricing, those with lesser harms, and even benefits, I could conceivably weaken readers' opposition to the most harmful forms of discrimination. This is not my intention, and I will do my best to steer readers away from this conclusion. In particular, I argue that focussing on the negative consequences of the worst instances of discrimination will also justify opposition to these. However, more generally, I would beg the reader to keep in mind throughout this dissertation that the existence of beneficial forms of differential pricing does not logically, and should not emotionally, weaken the case against the most harmful forms.

This dissertation proceeds as follows: In chapter 2 I provide an overview of how car insurance works. This includes discussion of what car insurance typically covers, and the factors which can be used for setting premiums. In addition I highlight some conceptual differences between car insurance and other forms of insurance. In chapter 3 I provide a justification for a consequentialist analysis of differential car insurance pricing. This begins with a consideration of the existing philosophical literature on differential treatment. I argue that much of this

literature focuses on a subset of differential treatment, discrimination. Unfortunately, this literature fails to provide a clear definition of discrimination, nor does it show that qualifying as discrimination is either necessary or sufficient for being unacceptable. However, I do believe the literature, particularly that of Lippert-Rasmussen, makes a convincing argument for assessing differential treatment based on consequences. I provide an additional pragmatic justification for consequentialist analysis: that it is likely to appeal to those of differing ethical perspectives.

In chapter 4, I identify and consider a number of consequences of differential pricing of car insurance. These include intended behavioural effects, unintended behavioural changes, measurement costs, effects on privacy, stigma, effects on economic inequality, and fairness. In chapter 5 I consider the implications of these consequences: the case for regulation, for social pressure and for compensatory redistribution. In chapter 6 I use this consequentialist framework to consider the two examples mentioned above: the use of sex and postcode as pricing factors. In chapter 7 I summarise my conclusions.

2 How car insurance works

2.1 Background

A car insurance policy is a contract where an insurance provider agrees to compensate for a specified set of costs. These costs are uncertain both in whether they will occur and how much they will be if they occur. In return, the policyholder pays a premium. Car insurance policies can cover damage to the car, damage to other cars or property, or harm to drivers or passengers. They can cover costs from driving accidents, theft, vandalism or natural causes.

In many jurisdictions, at least some level of car insurance is compulsory in order to drive a car, most commonly ‘third-party’ insurance covering harm to other people and their property. Additional coverage is optional, and typically covers damage to one’s own car, but may also cover additional risks (eg vandalism) or reduce excess payments payable by the driver. In this dissertation I consider both compulsory and optional car insurance.

In order to be profitable, a car insurance provider must collect enough premiums to cover the costs for all policyholders, administrative costs, and capital costs. An overly simplistic view of the market might suggest that the company estimate expected costs based on historical statistics, and then demanding all policyholders pay equal premiums to cover these. However, this is generally not viable, due to significant variation in expected costs between policyholders (Heath, 2014). Insurance providers and individuals may not know their precise expected costs, but will have some ability to estimate it. This means that those individuals being charged premiums significantly greater than their expected premiums have an incentive to go to a competitor or take less insurance, leaving the insurance provider with its more costly policies.

Car insurance providers must therefore attempt to estimate the expected costs for individual policyholders. This is challenging, as there are a number of factors, some difficult to evaluate, that have a significant impact on expected cost: for example, driving behaviour (frequency, distance, and skill), the costs of property that might be damaged and crime rates. Insurers rely on historical statistics of insurance costs, but these are limited, and generally fail to establish causal factors from merely coincident factors.

A final challenge of car insurance pricing results from the fact that expected costs are not independent of how whether people have insurance and how policies are priced. If a car is

insured against theft, it may be parked less securely, increasing the likelihood of theft. If a driver knows that a certain model of car will attract higher insurance premiums, they may choose a different model. Similarly, drivers may avoid claiming for losses if they know it will increase their future premiums. These incentive effects are challenging to quantify, and may be overstated, but create consequences that I consider in greater detail in chapter 4.

2.2 Factors influencing car insurance premiums

The previous section identified the need for insurance providers to estimate the expected cost for an individual driver, as well as the difficulty in doing this. In this section I mention some of the factors that are used in practice, or are discussed in literature as potentially useful. Before doing this, I wish to make three comments.

Firstly, throughout this dissertation I am focussing on the factors that are used directly. This is not intended to downplay the importance of indirect discrimination, in which the use of one factor (eg occupation) has a disproportionate impact on members of a particular group (eg women) (Lippert-Rasmussen, 2014, p55). I consider this phenomena when discussing the consequences of using particular factors. Secondly, insurance providers may believe that a factor is informative, but not use it, because of regulations or social pressure, or because of the cost of measuring them (Avraham et al, 2014, p221-223). Thirdly, it can be challenging to determine a complete list of factors used in pricing insurance. Insurance providers have incentives to not disclose their pricing mechanisms, to reduce competition, to prevent cheating, or to avoid negative publicity. There is always a possibility that insurance providers will price using information that a policyholder did not anticipate being used, for example the policyholder's choice of internet browser or email address.

I separate the factors that are often or potentially used into five categories: factors specific to the policy (what is being insured), to the car, to the location, to the driving history and to the owner/driver. This does not imply that the factors are independent; for example, the model of car being driven and the driving history are not independent of factors about the driver. Factors specific to the policy include what risks are being covered (eg whether damage to the car itself is covered), any limits, or any excesses. Factors specific to the car include the model of car, but also include the colour and any alterations made to the car. Factors specific to the location include where the car is registered and where it will be parked. Factors specific to the driving history include previous accidents or insurance claims, and may also include ongoing driving measurements. Finally, factors specific to the owner or driver may include the age, sex/gender, occupation, level of education, country of birth, religion, and computer / web browser / email address.

2.3 Differences between car insurance and health insurance

Much of the literature discussing car insurance considers it as part of an overall consideration of insurance. Examples of this include Avraham et al, 2014, Bieder, 1987, Gaulding, 1995, Rego, 2015. I believe that we can learn from experience and intuition in other insurance markets, especially health insurance which is frequently discussed, and I would hope that insights from this dissertation will have value beyond car insurance. However, I wish to tentatively propose two reasons to be cautious in assuming the same intuitions or conclusions necessarily apply.

Firstly, while health and car insurance costs both depend on luck and on choices, I believe that luck plays a much greater part in health than in car insurance. This is difficult to prove or quantify, and in both cases some people tend to emphasise the role of choice more than

others. However, given the low cost of damage incurred by people that rarely drive or do not even own a car, my belief does not seem entirely ungrounded.

Secondly, I and many people believe that healthcare should be provided for all, as indicated by widespread support for universal public healthcare in many societies. I do not have the same belief with respect to car ownership or driving. I recognise that this belief is not universal, and that in locations lacking public transport, car ownership may be seen as more essential. However, I would argue that, even where car ownership is seen as essential, few would advocate giving driving licenses to those that are unable to pass the necessary tests.

As a result of these two beliefs, I (and no doubt many others) have a much greater concern to make sure insurance costs do not prevent some people from missing out on healthcare. For some people this motivates the view that health insurance should not even exist and costs should be directly funded by the government, while others may still accept insurance so long as it allows access to all. However, in the case of car insurance, I do not believe there is the same level of concern.

One might think, on the other hand, that car insurance is more crucial than health insurance, as it determines your ability to compensate others for the damage you do to them and their property. This is the reason that many jurisdictions make third-party insurance compulsory. However, it should be noted that it is only ever compulsory for car owners/drivers; one can always avoid paying for car insurance by choosing not to drive, and one can also choose to avoid insuring against damage to one's own car. I would also question the suggestion that the decision to get health insurance is a purely personal decision, as many societies would not be willing to let a member go untreated even if uninsured.

In this chapter I have described what car insurance is, why insurance providers care about the expected costs for individual policyholders when setting premiums, and have given an introduction into the factors they might use to do this. I have noted, however, that the use of some factors is regulated by law or social pressure. In the remainder of this dissertation, I argue that regulation or social pressure is indeed often appropriate, and where it is, it is on the basis of the consequences of using these factors.

3 Justifying a consequence-based analysis of car insurance pricing

In the previous chapter, I noted that the use of some factors in setting car insurance premiums was regulated by law or social pressure. In this chapter, I argue that the most appropriate basis for judging which factors should be used is on evaluating the consequences of using the factors. This argument is significantly informed by the existing literature on discrimination, which I outline in 3.1, arguing for Lippert-Rasmussen's conclusion that consequences are what matter. In 3.2, I critically assess his attempts to distinguish between discrimination and non-discriminatory differential treatment, and also his attempt to construct a 'no-discrimination' baseline, arguing that these are problematic and ultimately unnecessary. In 3.3 I provide an additional argument of my own, that there are pragmatic reasons to prefer a consequence-based evaluation of differential treatment, in that it will appeal to those of different ethical perspectives.

3.1 The existing literature on discrimination

There is a large body of existing philosophical literature on discrimination and differential treatment, though that is unsurprising given the last half century's increased focus on eliminating racial and sex inequality. This literature largely agrees that some differential treatment is wrong or harmful, and that some is unproblematic or even essential, however there is a great deal of disagreement as to where the line lies and why (Altman, 2016). Central to this question is the concept of discrimination, however there is disagreement to what it means or which instances of differential treatment it includes.

I would suggest that many people, in everyday discourse, simply consider discrimination to be wrongful differential treatment. This may make it a useful term for labelling the most harmful cases, but it is less helpful when people disagree on whether a particular incident is wrong. An example is affirmative action, in which a group is favoured in order to overcome or compensate for a historic wrong. Existing laws on discrimination seldom attempt to define discrimination, and instead rely on examples. A number of philosophers and legal scholars have worked to more precisely define discrimination and to identify what makes it wrong. However, in their attempt to maintain the everyday notion that anything wrong must be included in the definition, they have created "a bewildering array of types of discrimination" (Altman, 2016): for example, direct, indirect, unconscious, institutional and statistical discrimination. These have little in common, other than an emphasis on comparatively negative treatment, and in most cases a basis in group membership (eg racial, sex).

On the question of what it is that makes discrimination wrong, philosophers have proposed a number of accounts, and although each sheds light, none is complete. For example, some accounts claim discrimination reflects inaccurate stereotypes. However, it is unlikely that merely making all stereotypes accurate would correct the problem, and Harford (2008) provides an illustration of how even rational discrimination can lead to undesirable outcomes. Another account, attributed to Hellman (2008), ties discrimination to a demeaning moral judgement or message. However, some cases of discrimination do not carry any such message, and there are instances where such a moral judgement is justified (for example, a jail sentence to a convicted criminal). The failure of these accounts has forced some, including Arneson and Lippert-Rasmussen, to resort to harm or consequence based accounts of what makes discrimination wrong (Altman, 2016). Under these accounts, discrimination is wrong because of its harmful consequences, and should it not harm anyone, it would not be wrong.

Lippert-Rasmussen (2014, ch 6) argues that we should assess discrimination using a consequentialist approach, that is, by determining whether discrimination leads to suboptimal consequences. There are many versions of consequentialism, spanning a range of choices for weighing and aggregating consequences. Utilitarianism may be the best known consequentialist model, however there are other models that avoid perceived weaknesses. For example, prioritarianism allows benefits to those that are worse off to be prioritised over benefits for those that are better off. Lippert-Rasmussen proposes a version of this, desert-prioritarianism, originally introduced by Richard Arneson (1999). This account, as well as prioritising those who are worse off, also prioritises those who are deemed more deserving. The desert-prioritarian account therefore overcomes an objection to many versions of consequentialism: that they do not accord with our intuitive desire for justice.

I support Lippert-Rasmussen's desert-prioritarian account of the wrong of discrimination for assessing car insurance pricing, however I note two challenges. Firstly, as with any version of consequentialism, it does not allow any act to be wrong, irrespective of consequences. In contrast, Hellman's account claims that a demeaning message is wrong, even in the absence of any actual harm. Furthermore, I would argue that human rights should not be breached on merely consequentialist grounds. However, I do believe desert-prioritarianism could account for these views by according maximal priority to such rights. Secondly, consequentialism requires more information to decide whether an act of discrimination is wrong, and adds

considerable subjectivity. Desert-prioritarianism adds even greater subjectivity, with its dependence on our intuitions of deservingness and justice. I recognise the challenge of this uncertainty, and discuss it further in chapters 5 and 6. However, I do nonetheless believe it would be inappropriate to ignore distribution and desert, especially where there is consensus that they matter.

3.2 Criticisms of Lippert-Rasmussen's account

In the previous section I introduced Lippert-Rasmussen's desert-prioritarian account of the wrongness of discrimination. While I do believe it overcomes weaknesses of other accounts, in this section I wish to raise two objections to it. Firstly, and most critically, I argue that its focus on the bounds of discrimination (shared with other accounts mentioned above) complicates rather than clarifies analysis. Secondly, Lippert-Rasmussen's struggle to solve what he terms "the baseline problem" (2014, p157) creates more complexity and subjectivity, and, I argue, is unnecessary for a consequentialist theory.

As mentioned in the previous section, philosophers have attempted to clarify the everyday notion that if differential treatment is wrong, it must constitute discrimination. Lippert-Rasmussen challenges it, highlighting nepotism (favouring one's family members) as an example of a wrong which should not be included in discrimination (family membership not being a sufficiently salient group) (2014, p23). He also acknowledges that different types of discrimination have different definitions. Nonetheless, his resultant definitions are complex; for example his criteria for when an action Φ constitutes direct group discrimination:

"X discriminates against Y in relation to Z by Φ -ing if, and only if,
 (i) there is a property, P, such that (X believes that) Y has P and (X believes that) Z does not have P,
 (ii) X treats Y worse than Z by Φ -ing,
 (iii) it is because (X believes that) Y has P and (X believes that) Z does not have P that X treats Y worse than Z by Φ -ing,
 (iv") P is the property of being member of a certain socially salient group (to which Z does not belong), and
 (v) Φ -ing is a relevant type of act etc., and there are many acts etc. of this type, and this fact makes people with P (or some subgroup of these people) worse off relative to others, or Φ -ing is a relevant type of act etc., and many acts etc. of this type would make people with P worse off relative to others, or X's Φ -ing is motivated by animosity towards individuals with P or by the belief that individuals who have P are inferior or ought not to intermingle with others." (Lippert-Rasmussen, 2014, p45-46)

I argue that focussing on the terms of such a complex definition is not productive. As Lippert-Rasmussen accepts, showing that an act meets these criteria does not establish it to be wrong. Conversely, showing that an act does not meet these criteria does not make it right (it may be another type of discrimination, or another kind of wrong). Furthermore, its clauses are open to accusations of arbitrariness, as proponents and opponents of certain acts work to exclude or include them from the definitions. I believe we should instead focus directly on whether any act of differential treatment is wrong per our preferred consequentialist view, such as Lippert-Rasmussen's desert-prioritarianism¹.

¹ While existing literature has highlighted the difficulties in determining discrimination, and has used other arguments against forms of differential pricing, especially irrationality, I have not come across any papers arguing that we should consider all differential treatment and not just discrimination. However, I am sure other philosophers have considered it, and I do not wish to overstate the originality of my argument.

My second critique is to question Lippert-Rasmussen's attempt to construct a baseline. Under his consequentialist analysis, we should deem an action wrong if it creates worse aggregated consequences when compared to the relevant baseline outcome. He casts doubt on a Straightforward account in which the baseline would be the world in which the act in question did not occur, but all other circumstances were unchanged. Under this account, a discriminatory act might be permissible if it reduced the opportunity for future discrimination. In contrast, his preferred baseline is the No-Discrimination account, in which the act in question did not occur, and no future discrimination occurred. My concern with this account is that by ignoring likely future discrimination, we might forego the chance to achieve genuinely better outcomes.

I wish to argue that if we are using a consequentialist approach, we do not need a single baseline account. Instead, whenever an action is to be assessed, we can use our preferred consequentialist account to compare whether the outcome is better (or likely to be better) than other outcomes that we could bring about. I accept that the ability to compare an action against multiple alternatives does prevent us ascribing it a single measure of badness. However, I argue that this is a necessary price to pay, and preferable to making an incorrect evaluation through using the wrong baseline.

3.3 Pragmatic arguments for a consequentialist evaluation of differential treatment

In the previous two sections, I have argued for evaluating instances of differential treatment based on their consequences, but not focusing on whether the instances satisfy particular definitions of discrimination. I have clarified that we should care about subjective as well as objective consequences, and that a single baseline not always available. In this section I wish to provide an additional, pragmatic argument for such a consequence-based evaluation: that its conclusions are likely to resonate with those of differing ethical frameworks. Most obviously, it is designed to appeal to consequentialists. I also discuss how it might appeal to those that believe obligations are limited to avoiding profiting from market failures, or those that believe their only obligation is to comply with laws.

A consequentialist evaluation of differential treatment will obviously appeal to other consequentialists, or those wishing to promote consequentialist objectives (for example a government policymaker). As highlighted in section 3.1, there are a number of versions of consequentialism. Different versions aggregate consequences differently, and form different conclusions as to which outcome is best. Despite this, what the consequences of differential pricing are, will matter to any of them. I also argue that considering the full range of consequences is likely to lead to better decisions than ignoring less objective ones, such as moral judgement and stigma.

An alternative ethical framework that a business or individual may adopt is the Market Failures Approach (Heath, 2014, p36). Under this framework, businesses are free to maximise profits, so long as they do not profit from market failures like negative externalities. Negative externalities are harms incurred by parties other than those involved in the transaction. Many of the consequences of differential treatment I introduce in the next section may indeed be classified as negative externalities. As such, a consequentialist analysis, by highlighting negative externalities, is likely to be of use to businesses seeking to follow the Market Failures Approach, or more generally, wishing to avoid causing negative externalities.

Finally, it must be recognised that some businesses believe they are free to maximise profits, subject only to compliance with the law. This stance is argued by Milton Friedman (although he included the obligation to comply with ethical custom) (2002, p133). However, I believe it would be wrong to assume that an appreciation of consequences would have little importance for such businesses. Many people accept a role of government to prevent harm to other

people, such as that proposed by John Stuart Mill in “On Liberty” (2011, p10), so a business should be aware that harms inflicted on others may lead to future regulations. Even where regulation is unlikely, society is likely to notice positive and negative consequences of a business’s activities. Society is likely to use its resources to punish companies that inflict negative externalities, or reward those that provide positive externalities. Recognising consequences and proactively managing them is therefore likely to be good business, even for one that is purely seeking to maximise profits.

This argument also justifies the role for citizens in evaluating the consequences of differential treatment by companies. Where companies are inflicting harm on members of society, governments can be lobbied to regulate (or better regulate) the activity. Where regulation is not possible, direct lobbying or negative publicity may lead to better outcomes. Finally, awareness of positive consequences can allow citizens to encourage or reward or behaviour that leads to these. These opportunities are discussed further in chapter 5.

I do wish to note two challenges to a consequence-based analysis of differential treatment. Firstly, one might argue that some kinds of treatment are so demeaning or so unjust, that a prohibition is justified outright. Requiring opponents of discrimination to identify negative consequences might make it more difficult to stop discrimination. I agree with this concern, however would argue that a consequence-based analysis can allow categories of treatment where prohibition was assumed by default, and justification required when the benefits justify the action. Secondly, I worry that consequentialism can encourage an emphasis on objectivity, for example weighting financial costs and benefits above emotional costs and benefits. I strongly encourage consideration of all costs and benefits, and do my best to demonstrate that in the next chapter and throughout this dissertation.

I have therefore argued that an analysis of the consequences of differential treatment makes sense for those from a number of perspectives; governments, citizens, and businesses working from a range of ethical perspectives. In the next chapter I explore the range of consequences that might arise from differential pricing of car insurance.

4 Identification and consideration of the consequences of car insurance pricing

In chapter 2, I explained what car insurance is, and why insurance providers set premiums that cover the expected costs of individual policyholders. I described the challenge in estimating these costs, both because of generally insufficient statistics, and because expected costs are themselves affected by insurance pricing. I outlined a range of factors that are used to price insurance: factors specific to the policy (what is being insured), to the car, to the location, to the driving history and to the owner/driver. In chapter 3, I justified the use of consequentialist analysis to evaluate the appropriateness of differential pricing. In this chapter, I begin that analysis, identifying and considering consequences: intentional behavioural effects, the potential for unintended behavioural changes, measurement costs, effects on privacy, stigma, the effect on economic inequality and on perceived fairness. I note that this categorisation of consequences is my own, although it draws on analysis of discrimination by a number of writers, and I am particularly indebted to the list of considerations in Avraham et al (2014, p204-220), and the work of O’Neil (2017).

4.1 Intentional behavioural effects

The way car insurance is priced will invariably affect behaviour: the decision to insure and the risks taken. In this section I consider these effects. I refer to these as intentional behavioural effects as they are likely to be acknowledged (and perhaps even welcomed) by the insurance provider. I first assume that we insure against all costs, and consider how insurance pricing can lead to more efficient decisions: whether to own a car, which car to own, where to park, how to drive, and how to reduce crime and accidents. It is important to highlight that these incentives may also have less positive consequences, which I consider in subsequent sections. I then consider how pricing affects the decision of whether to insure. Finally, I draw a connection between these effects and the question of control, which has been discussed in the discrimination literature.

Car insurance pricing may influence the decision to own or drive a car. It only makes economic sense to own a car if the value of ownership exceeds the costs of ownership, including insurance (under our initial assumption that we insure against all costs). Pricing insurance premiums significantly above expected costs may lead to drivers being unable to capture the positive net benefits of ownership. On the other hand, pricing them too low may lead to people buying cars when they have excessive expected costs. Car insurance pricing may also affect one's choice of car to drive. Some makes and models have higher expected costs, and to the extent that insurance pricing reflects these, this may sensibly incentivise drivers from purchasing cars with lower costs.

Effective car insurance pricing can influence how one drives or parks, in a number of ways. The most direct example of this is the use of 'black box' devices that record driving behaviour, allowing premiums to be reduced for drivers that drive safely. In other cases, discounts may be offered for drivers that make no claims, incentivising drivers to drive more carefully or park more securely. Additionally, these incentives may motivate drivers to take additional driving lessons to improve their skills in order to lower their expected costs (and premiums).

A final way in which car insurance pricing can effectively encourage optimal behaviour is by highlighting ways that communities can reduce car ownership costs. For example, it may make sense for residents to increase the security of parking, and offering premium reductions might incentivise this. Similarly, drivers may be more supportive of lower speed limits if they know they will benefit from lower insurance premiums.

So far in this section I have assumed that we would insure all risks, however in some cases, it may be economically rational to not insure against all risks. Firstly, there is the case of moral hazard, where being insured increases the the expected cost. Where this cannot be prevented, it is likely to reduce the use of insurance, especially for smaller risks. Secondly, the insurance provider may set premiums significantly higher than the potential policyholder estimates their costs to be. In this case, competitive pressures may reduce the gap, but this may take some time and is not guaranteed in imperfect markets. Finally, if drivers are risk-neutral or even risk-seeking, they may not benefit from insurance.

When we discuss the positive benefits that incentives can create, we implicitly assume the existence of individual control, a factor that is discussed at length in the discrimination literature. For example, some writers argue that it is unfair to charge more based on factors outside an individual's control (Avraham et al, 2014, p214). In practice, it is often difficult to identify how much control an individual has over some factor, and individuals may vary in their ability to respond to these incentives.

4.2 Susceptibility to unintended consequences

As well as the economically efficient consequences discussed in the previous section, any pricing system may lead to cheating or other unintended consequences. In this section, I provide examples, and discuss the challenges of avoiding such outcomes.

When registering for car insurance, an owner is likely to be asked for information to be used in determining the appropriate premium, for example their address or occupation. One can significantly reduce the premium by providing a different address or occupation, either by lying or on more justifiable grounds (eg referring to your occupation as “finance” rather than “banker”). Insurance providers may seek to address this by ensuring factors used are verifiable, however this can prevent use of factors that have real predictive power over expected costs. Alternatively, insurance providers may seek to disguise the effect of such factors, but this can prevent positive consequences of price signals discussed in the previous section.

In the previous section, I discussed how lower premiums associated with one model of car can appropriately incentivise a driver to choose that car. However, this can lead to results that do not lower overall costs. For example, if a certain model merely happens to be driven by mainly safer drivers, it does not improve outcomes for another driver to switch to that model. To address this, statistical analysis should attempt to distinguish between causal and non-causal factors.

Finally, with new and more complex mechanisms for pricing, such as black box devices, there is increased risk of drivers identifying behaviour that leads to lower premiums, without in fact lowering expected costs. For example, if smooth driving is rewarded, a driver may drive more at night to ‘dilute’ the effect of everyday city driving. Alternatively, if swerving or hard braking is penalised, a driver may feel pressure to not brake for a cat on the road. It is important to be aware of the possibility of such unintended consequences.

4.3 Measurement cost

Some take the view that the ideal car insurance pricing mechanism would use all possible resources to estimate the expected cost associated with each driver (Bossert and Fleurbaey, 2001, p114). However, this view ignores the very real costs of those resources. These costs can include verifying the information given by policyholders, regulatory compliance (eg privacy obligations), sophisticated data acquisition and analysis, and driver monitoring devices. In addition, more complex pricing mechanisms can weaken customers’ trust, if they believe they are being treated unfairly. While technology offers ways to calculate better estimates more cheaply and more transparently, it is likely that insurance providers will always need to weigh the costs and benefit of obtaining better estimates.

4.4 Effect on privacy

A number of the mechanisms used to price car insurance premiums have potentially harmful privacy consequences. Black box monitoring devices could conceivably record all your driving and parking movements. If this data were stolen, it could enable burglary or robbery. Even without the data being stolen, it could be used by your insurance provider or the government in ways that you might reasonably not want.

I wish to note that not everyone values privacy equally; some people would be happy to allow their insurance provider to track their movements in return for lower insurance premiums while others would not. I believe that it is important to especially consider the consequences for those that have reasons to value privacy. This could be done by limiting the ways in which

personal data could be misused, as well as ensuring people are able to opt out of providing personal data without being forced to pay significantly higher premiums.

4.5 Effect on stigma

One of the biggest arguments against discrimination is that it often represents or encourages negative moral judgement or stigma; indeed, this is the primary criteria of discrimination according to Hellman's account introduced in section 3.1. While our consequence-based evaluation of differential treatment considers a much wider range of harms, it is important to not ignore the risk of stigma, despite it being somewhat less objective than other harms.

I believe that stigma is less present in car insurance pricing than in other forms of differential treatment. People are generally unaware of what premiums other people are paying. They also often do not know how personal factors are used to derive prices, and many of the factors that lead to higher prices do not relate to any moral judgement like them being a bad driver.

However, in some cases stigma can still arise. For example, the idea that young male drivers are fundamentally riskier drivers could be considered a stigma. It is also important to recognise that stigma can affect a group, even without using group membership directly as a factor. This can arise if a factor that is used applies disproportionately to one group. Furthermore, there is always a risk that insurance pricing mechanisms could evolve in ways that increase stigma. This might arise if insurance providers are currently avoiding using certain factors because society recognises they cause stigma. Finally, increasing transparency of how factors are correlated with risk, while leading to better outcomes in some respects, may in fact increase the potential for stigma.

4.6 Effect on economic inequality

Altering how car insurance is priced affects economic inequality. Innovations that lead to richer customers paying less, or poorer customers paying more, will increase inequality. This might arise from charging more for people that live in poorer neighbourhoods, or charging more for people that are accessing an insurance provider's website from a public (eg library) computer. Even overall increases in premiums affect inequality, though its overall effect of transferring wealth from car owners to shareholders, employees and tax revenue may be more difficult to evaluate.

Many consequentialists see increases in economic inequality as undesirable. Even utilitarians may believe that less well-off people may derive greater utility from a given amount of money than those that are better-off. I do believe that economic inequality should be taken into account when evaluating insurance pricing mechanisms, and note that desert-prioritarianism does so. This should not merely be a matter of discouraging innovations that increase inequality, but also an argument for encouraging innovations that decrease inequality as well as achieving other positive consequences.

When evaluating the economic inequality effects of car insurance pricing, it is important to recognise that redistribution (taxation, benefits, and public services) can also be used to reduce income inequality. Whether a new pricing mechanism can be fully compensated for by increasing redistribution will depend on a number of factors. Increases to tax and benefits may be politically infeasible or increase stigma. Redistribution may lower overall efficiency (though possibly by less than an inefficient insurance pricing system). Finally, the economic gains and losses from a change in insurance pricing may have complex effects, hurting only some poorer people; such effects may be challenging to compensate for using redistribution.

4.7 Questions of fairness

When I talk about my dissertation with friends and colleagues, almost everyone mentions fairness as an important factor when evaluating insurance pricing. In this section I first consider whether fairness should be considered, concluding that it should. However questions of fairness tend to be challenging to consequentialist models, with their focus on outcomes rather than how they come about. However, I do believe it is important to consider fairness concerns, and in this section I discuss two ways to incorporate them into a consequence-based analysis of insurance pricing. I also consider a fairness-based challenge to certain kinds of differential pricing that appears in the literature.

The literature on discrimination has considered whether interpersonal comparison should be taken into account, or whether discrimination can be accounted for as some individual harm. Scheinman (2014) considers individualistic ways of explaining the harm of discrimination, for example, failing to receive dignity or to meet essential needs. In the case of car insurance, unfair treatment may be wrong because it prevents people charged higher premiums from meeting their needs. Alternatively, it may be wrong because it stigmatises one or more groups. In such cases, the harm would not be mitigated by 'leveling down', for example, charging similarly higher premiums to everyone else. However, Scheinman ultimately concludes that individualistic accounts cannot explain our intuitive desire for interpersonal fairness. I support this conclusion, and suggest that it can be done in two ways.

Firstly, while consequentialist models tend not to take into account questions of fairness, they can if desired. Most notably, Lippert-Rasmussen's desert-prioritarian account does. Under this account, a state in the world where a gain goes to a more deserving individual, is preferable to one where it goes to someone less deserving. This accords with popular intuition. However, it would never choose to take an unfair benefit away from anyone; in some cases this is likely to be intuitive, and in other cases not. It also underplays the difficulty in agreeing the fair outcome, especially in complex situations without an agreed neutral baseline.

A second approach to incorporating fairness is through its relationship with incentives, particularly those considered in section 4.1. If incentives are to be effective, people will have to trust that if they act, they will get their reward; this can be considered a form of fairness. Therefore we should expect to see greater desire for fairness when it supports intentional incentives. Desire for fairness may also be observed beyond those cases. For example, as an overarching social norm, people will respect and enforce fairness conditions so long as others do (Bicchieri, 2006, p101). In this case, undermining people's sense of fairness might have wider-reaching negative consequences, though it would be challenging to estimate these.

Rego has proposed an account of discrimination in which insurance providers should not charge individuals more, despite their higher expected costs, if the higher costs result from historical injustice (2015). For example, if people from a certain group were prevented by law from driving until recently, and consequently caused more accidents, premiums should not reflect those costs. I believe this idea reflects the desire for fairness, in that it seems unfair for people to be penalised for something forced upon them. I would challenge the idea that it is the job of the insurance provider to compensate a group of individuals for historic injustice, and also that this desire for fairness necessarily outweighs all benefits that may result from differential pricing. However, I do accept that the account has value in recognising that current differences in expected costs may reflect past injustices and thus not be fair, and also that certain pricing mechanisms may exacerbate the past injustice.

5 Responding to the consequences of differential car insurance pricing

In chapter 3, I justified using a consequentialist analysis of car insurance pricing, in particular emphasising its ability to guide and justify various governmental, individual and societal responses to possible consequences. In this chapter, I delve deeper into the responses that might be appropriate. I must emphasise that I am hesitant to argue for specific responses, as which are appropriate depends on our differing beliefs of how likely different consequences are, and our differing views on how consequences should be aggregated.

This chapter is structured as follows: Firstly I discuss the use of governmental regulation to require, or to prohibit, particular methods of differential pricing. Secondly, I discuss the use of social pressure as a way of influencing how insurance providers calculate premiums. Thirdly, I consider the possibility of responding through redistribution via fiscal policies, that is, changes to taxes, benefits and government policies.

5.1 The case for governmental regulation

Car insurance is a contract between an insurance provider and an individual car owner or driver. Within this dissertation I assume it to be entered into freely by both parties, because the individual can choose to insure with another company, or perhaps to not insure. As such, both parties are likely to be better off for the transaction, and one might question why the government should interfere. I offer two answers to that question.

Firstly, the strongest case for government regulation is when insurance activity harms others. Indeed, according to Mill's "On Liberty" (2011, p10), this is the only justification for government interference in private activity. For example, if car insurance pricing mechanisms undermine privacy (see section 4.4) or create stigma towards vulnerable groups (see section 4.5), regulation may be justified to prevent these harms. However, even in this case, it is not always clear how the government should best regulate; it may be possible to ensure the victims are compensated, or it may be better to prevent the harm in the first place.

A further challenge with harm-based justifications of regulation is determining whether negative consequences constitute harm or not. This is made more challenging if we do not have a clear baseline, as discussed in 3.2. For example, in section 4.1, I noted that car insurance mechanisms that rewarded drivers that attended driving courses might improve road safety for everyone. However, I am not sure that the absence of such a policy constitutes harm, as such, to other drivers. As a result, in some cases we may struggle to justify the use of government regulation to prevent negative consequences.

The second answer to why government should regulate is on consequentialist grounds: that it produces better outcomes, under some measure. This approach fits more naturally with our consequentialist analysis of car insurance pricing. It also has the advantage of allowing us to maximise positive consequences as well avoiding negative consequences. As was discussed in chapter 3, we can aim for a consequentialist theory that takes into distribution and fairness, for example Lippert-Rasmussen's desert-prioritarian approach.

A major challenge to regulation is people's tendency to ignore "how little they really know about what they imagine they can design" (Hayek, 1988, p76). There are uncertainties in the consequences of the insurance practices that we are looking to regulate, and also in how people are likely to respond to regulations. When I consider two actual examples of car insurance pricing in the next chapter, I emphasise our uncertainties. I do not wish to underplay the significance of uncertainty, and believe it should be taken into account, and should lead to cautiousness in regulation. However, I do believe that we should use regulation where it is likely to lead to significantly better outcomes.

5.2 The case for social pressure

In the previous section, I noted challenges in use of government regulation, in particular opposition to regulation except in cases of harm, and uncertainties that make it difficult to justify consequentialism based regulation. In this section, I wish to argue that where current insurance pricing mechanisms are suboptimal, it may be rational for people to use their resources, individually and collectively, to apply social pressure to create change.

In the case of car insurance, individuals have a number of ways to apply social pressure. Communicating positive and negative consequences of premium pricing can alter incentives, making it rational for insurance providers to improve their pricing mechanisms. It can also lead to other consumers taking the behaviour into account when choosing an insurer, creating more direct pressure. An advantage of using social pressure is that it is generally less coercive than government regulation. This is especially valuable where there are considerable uncertainties, and coercion may be less justified. However, using social pressure can also be a way of achieving government regulation, especially if there proves to be popular support for it.

I note two potential downsides to social pressure as a way of influencing car insurance pricing. Firstly, social pressure can reflect inaccurate information, and can turn into misplaced outrage. Alternatively, other people can remain apathetic, particularly if there are short term individual costs to achieve the larger, collective benefit. Secondly, there is the risk that social pressure may only be effective on larger, established insurance providers. This can give an unfair advantage to newer startups that are less concerned about public opinion, and can even lead to a worse situation as socially responsible providers are driven out of the market. I do not believe either of these succeed as reasons to never use social pressure, but they should be considered when deciding between social pressure and governmental regulation.

5.3 The case for compensatory redistribution

The final response I wish to discuss, to the consequences of differential car insurance pricing, is that of compensatory redistribution, for example, through the tax and benefits system. This can reduce economic inequality without directly impacting car insurance providers. Such an approach would have the advantage of maintaining the intentional incentives in the insurance pricing system, while reducing the inequality. However, this proposal is not without risk, and could result in greater stigma, and a sense of unfairness by all. I also note the impossibility in adjusting taxes and public services to respond to every change in pricing. As a result, this suggestion may be better suited at a higher level, to justify enhanced levels of redistribution by noting the increases in cost of living for those that are worse off, which are not limited to car insurance (The Economist, 2015).

This approach is particularly relevant for the situation where an insurance provider charges more for car insurance to those living in poorer neighbourhoods, on account of the increased risk of crime. This situation is discussed in further detail in the next chapter.

6 Considering two car insurance pricing controversies

The previous two chapters have been at a reasonably generic level, considering the type of consequences that might arise from car insurance pricing, and possible ways we might respond to these. In this chapter I apply our consequentialist framework to consider two specific examples of differential pricing that have been discussed in the literature and were mentioned in the introduction: sex based pricing and postcode based pricing. This involves

identifying the consequences, assessing the different ways in which we could respond, and evaluating which leads to the best expected outcome. This evaluation is subjective, and those with different beliefs about likely consequences and how to evaluate these may reasonably come to different conclusions, but I argue that discussion along these lines would support building consensus.

6.1 Sex-based pricing of car insurance

Historically there has been a strong tendency for car insurance providers to take into account sex or gender when pricing policies. I believe this was done based on historical statistical evidence, which showed that women incurred on average lower losses than men. Other reasons that sex is used as a factor include its relative objectivity, ease to determine, rarity of cheating, and substantial historical data (all relative to other factors that could be used). Some have also argued that the historical differences are due to fundamental biological differences between men and women, however this is very much in doubt, due to the significant influence of social factors (Saini, 2017, ch 3). It should also be noted that some women incur higher costs than some men; this is most obvious when considered after the event (*ex post*), but undoubtedly also in expectation (*ex ante*). For example, a man that almost never drives will expect fewer accidents than a woman that regularly drives long distances.

A number of jurisdictions have subsequently prohibited the use of sex, for example, almost half of US states (Avraham et al, 2014, p245) and the European Union (Rego, 2015, p119). In at least the case of the European Union, the decision to outlaw use of the factor was not specific to car insurance, but across all insurance markets, including some where women would incur higher average costs (eg life annuities). I believe that this prohibition is based on general suspicion of sex-based discrimination, and that car insurance providers were offered the chance to argue that sex-based discrimination was necessary for their business, and chose not to. Despite the ban, women still pay less on average, and in fact the gap may have widened (Collinson, 2017). While insurance providers could be ignoring the ban, this is relatively easy to detect, and it is most likely that they have increased the significance given to the car model and driver occupation, both of which are correlated with sex.

Some people might claim that these other factors are merely proxies for sex, however I would argue instead that sex was a proxy for these other factors. The causal connection between the car you have and your expected costs is likely to be stronger than any causal connection between sex and expected costs, and likewise between driver occupation and expected costs. As a result, these factors seem intuitively fairer (at least to me) than using sex to set premiums. I would also argue that these factors work better as incentives, as they are things a driver can to a large extent control. However, as noted in section 4.2, to the extent that car model or occupation does not cause costs but merely selects people with different costs, these incentives are likely to be unhelpful. If this is the case, it may prove even fairer to use blackbox devices, such that actual driving behaviour can be taken into account, rather than that predicted based on an occupation or car model.

It is difficult to assess the likely effect of the prohibition on economic inequality and stigma. Women earn less on average, so one might think even lower prices for women would reduce inequality. However, if the occupations that incur higher premiums are more likely to be held by poorer people, or the cars that cause more accidents are cheaper, this could exacerbate economic inequality. Regarding stigma, it may well be that a stereotype of (younger) men being dangerous drivers has harmful effects, and weakening that through the prohibition is positive. It is less clear what the impact on women would be, and even being judged as a safe driver is not purely positive. Finally, I would note that the use of factors other than sex is likely to be more expensive to capture and model. However, it may be that car insurance providers were already taking them into account, in which case the change would be less costly.

Taking all these consequences into account, I am inclined to see the prohibition as an improvement and justified, especially in terms of improving incentives to take steps to lower one's expected costs. However, the effects on economic inequality and stigma are sufficiently unclear that I would be open to being convinced that they were causing harm, and that other steps were appropriate to counteract these.

6.2 Postcode based pricing of car insurance

A driver's postcode is a significant determinant of insurance premiums. To a large extent this is statistically based. Traffic variation makes certain areas more accident prone, and accidents more costly. Variations in crime make parking in certain suburbs riskier. Geographic factors make some areas at greater risk of weather related damage. These factors may well lead to those that live in poorer neighbourhoods incurring on average higher costs than those in richer neighbourhoods. However, we cannot ignore the likelihood that insurance providers are charging people in poorer neighbourhoods even more than the statistics would imply (O'Neil, 2017, p164). This may not be intentional, for example, economic data unrelated to car costs (for example, burglaries or loan default rates) may be used within the pricing model. Lack of competition may lead to some people being forced to pay these higher prices in order to drive.

Increases in premium based on statistically irrelevant data increases economic inequality, and may exacerbate stigma (depending on how aware people are of what premiums people pay and how these are calculated), and lacks any offsetting positive consequences. It does not incentivise people to lower their costs, and may prevent people from driving when they could valuably do it. I struggle to perceive such premiums as fair or good. In this case, I would see charging more accurate premiums as fairer. It may be challenging to regulate with enough specificity to prevent such overcharging, but I would at least argue for social pressure to be applied to discourage overcharging.

Where the premiums are based on statistically relevant data, however, I do perceive some benefit in terms of incentives. It may encourage people to move to neighbourhoods with less risk of losses, or to choose not to drive. It may also increase incentives to lobby for policing resources to reduce car related crime, in order to reduce premiums. I recognise that poorer people are still being forced to pay more for their insurance, and I do not consider this to be fair, in an overarching sense. However, if I was asked to make it fair, I would not choose to prevent insurance providers from taking into account postcode. I would instead prefer a transfer of wealth from all richer people to all poorer people, rather than one from the insurance provider or other drivers to poorer drivers. As such I would prefer to use the overarching redistribution system to correct for the existing inequality, as described in section 5.3 above. I do recognise that this conclusion is based on my belief that increased redistribution is achievable; otherwise I would be more open to social pressure or regulation to reduce such increases in economic inequality.

7 Conclusion

In this dissertation, I have argued for a consequentialist analysis of differential pricing of car insurance. In chapter 3, I have built on the work of Lippert-Rasmussen, who argues that other accounts fail to explain what is wrong with discrimination and when it is acceptable, and that the best account is a consequentialist one. I agree with Lippert-Rasmussen that we must be careful in selecting a consequentialist account; distribution and fairness matter, as do less objective consequences like stigma. However, I deviate from the approach in much of the

philosophical literature by applying this consequence-based analysis to all differential pricing, rather than limiting myself to that labelled discrimination. I do this for largely pragmatic reasons; I am concerned that present battles over the definition of discrimination are less helpful at preventing the harms of differential pricing.

In chapter 4 I have described a range of consequences of differential pricing of car insurance: intended behavioural effects, the potential for cheating or other unintended behavioural changes, measurement costs, effects on privacy, stigma, and the effect on economic inequality and fairness. I strongly believe in the appropriateness of responding to these consequences, and in chapter 5 I have considered the benefits and challenges of using government regulation, social pressure and compensatory redistribution. In chapter 6 I applied the approach to two specific examples, demonstrating how such a consequence-based analysis adds clarity.

Throughout this dissertation I have acknowledged the challenges of consequence-based analysis in terms of uncertainty as to the consequences, and differences in opinion on how to weigh these. However, I nevertheless believe that discussing the consequences is a fruitful way of seeking consensus, and a common basis for working together to tackle the negative consequences of car insurance pricing. However, even beyond car insurance pricing, I believe this approach would be a productive way of addressing the consequences of other commercial activity, for example those of artificial intelligence or large privately owned technology platforms.

References

- Alexander, Larry, 1992. "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies", *University of Pennsylvania Law Review*, Vol. 141, No. 1 (Nov., 1992), pp. 149-219.
- Alexander, Larry, 2016. "Is Wrongful Discrimination Really Wrong?" (May 15, 2016). San Diego Legal Studies Paper No 17-257.
- Altman, Andrew, 2016. "Discrimination", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2016/entries/discrimination/>>.
- Arneson, Richard, 1999. "Egalitarianism and Responsibility," *Journal of Ethics* 3.3, 225–247.
- Arneson, Richard, 2006. "What is wrongful discrimination", *San Diego Law Review* 43, 775–808.
- Avraham, Ronen et al, 2014. "Understanding Insurance Antidiscrimination Laws", *Southern California Law Review*, Vol 87(2) p195-274, 2014.
- Baumle, Amanda K and Fossett, Mark, 2005. Statistical Discrimination in Employment: Its Practice, Conceptualization, and Implications for Public Policy. *The American Behavioral Scientist*. 49: 1250-1274.
- Beider, Perry, 1987. "Sex Discrimination in Insurance", in *Journal of Applied Philosophy*, 4(1)
- Bicchieri, Cristina, 2006. "The grammar of society : the nature and dynamics of social norms", Cambridge University Press.
- Bossert, Walter and Fleurbaey, Marc 2001. "Equitable Insurance Premium Schemes," *Social Choice and Welfare* V, 19/1: 113–125.
- Collinson, Patrick, 2017, "How an EU gender equality ruling widened inequality", in *Guardian*, 14 January 2017, <https://www.theguardian.com/money/blog/2017/jan/14/eu-gender-ruling-car-insurance-inequality-worse> accessed 19 August 2018.
- The Economist, 2015. "It's expensive to be poor", 3 September 2015.
- Eidelson, Benjamin, 2014. "Treating people as individuals". In Hellman, D and Moreau, S (eds) *The Philosophical Foundations of Discrimination Law*. OUP: Oxford. p354-395.
- Eidelson, Benjamin, 2015. "Discrimination and Disrespect", Oxford University Press.
- Friedman, Milton, 2002. "Capitalism and Freedom", University of Chicago Press, Chicago.
- Gaulding, Jane, 1995. "Race, Sex, and Genetic Discrimination in Insurance: What's Fair?", in *Cornell Law Review*, Vol 80, September 1995.
- Harford, Tim, 2008. "The dangers of rational racism". In "The Logic of Life", London: Abacus.

- Hayek, Friedrich, 1988. "The Fatal Conceit", University of Chicago Press, Chicago.
- Heath, Joseph, 2014. "Morality, Competition, and the Firm: The Market Failures Approach to Business Ethics", Oxford University Press.
- Hellman, Deborah. 2008. "Why is Discrimination Wrong?", Cambridge, MA: Harvard University Press.
- Horta, Oscar, 2010. "Discrimination in terms of moral exclusion". *Theoria*, 76: 314-322.
- Lippert-Rasmussen, Kasper, 2006. "Private discrimination: A prioritarian, desert-accommodating account". *San Diego Law Review*, 43: 817-856.
- Lippert-Rasmussen, Kasper, 2007. "Nothing personal: On statistical discrimination". *The Journal of Political Philosophy*, 15, 385-403.
- Lippert-Rasmussen, Kasper, 2011. "We are all different: Statistical discrimination and the right to be treated as an individual". *The Journal of Ethics*, 15, 47-59.
- Lippert-Rasmussen, Kasper, 2014, "Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination", Oxford University Press.
- Maitzen, Stephen, 1991. "The Ethics of Statistical Discrimination", *Social Theory and Practice*, 17: 23-45.
- Mill, John Stuart, 2011. "On Liberty", Andrews UK Limited.
- Moreau, Sophia. 2010. "What is Discrimination?" *Philosophy and Public Affairs* 38: 143–179.
- O'Neil, Cathy, 2017. "Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy", Penguin Books, UK.
- Palmer, Daniel E, 2007. "Insurance, risk assessment and fairness: An ethical analysis. Insurance ethics for a more ethical world", *Research in Ethical Issues in Organisations*, 7: 113-126.
- Rego, Margarida, 2015. "Statistics as a basis for discrimination in the insurance business", *Law Probability and Risk* (2015) 14 (2): 119.
- Saini, Angela, 2017. "Inferior: The True Power of Women and the Science that Shows It", 4th Estate, London.
- Scheinman, Hanoch, 2014. "Two Faces of Discrimination", in In Hellman, D and Moreau, S (eds) *The Philosophical Foundations of Discrimination Law*. OUP: Oxford. p28-50.

Schauer, Frederick, 2003. "Profiles, probabilities and stereotypes", Harvard University Press.

Scholes, Vanessa, 2014. "Stereotypes and discrimination". Paper presented at 2014 Conference of New Zealand Association of Philosophy.

Scholes, Vanessa, 2016. "Risky Business - The Ethics of Judging Individuals based on group statistics", Contemporary Issues in Applied and Professional Ethics.

Scholes, Vanessa, 2017. "You are not worth the risk: The ethics of statistical discrimination in organisational selection of applicants", Thesis.

Segall, Shlomi, 2012 "What's so Bad about Discrimination?", Utilitas.