Moral Dimensions of Moral Hazards

WILL BRAYNEN
Stanford University

'Moral hazard' is an economic term which commonly refers to situations in which people have a tendency to increase their exposure to risk when the costs of their actions, should they get unlucky, befall someone else. Once insured, for example, a person might have little reason, financially speaking, to be careful if he will get fully reimbursed for his losses should things go wrong, especially if he does not risk an increase in his insurance premium fees. In this article, I argue that moral hazards are not morally neutral. To this end, I distinguish between concepts that (a) call for a moral value judgement but do not have a fixed moral value and (b) those that call for a moral value judgement and also have a fixed moral value. In short, this article examines questions that lie at the intersection of ethics and economics.

I. INTRODUCTION

Following Kenneth Arrow's seminal paper on moral hazards, the economist Mark Pauly promised that in his paper it will 'be shown that the problem of "moral hazard" in insurance has, in fact, little to do with morality'.[1] Surprisingly little has been written by philosophers on the topic of the ethics of moral hazards. Ronald Dworkin briefly mentions moral hazards, but not with an eye to evaluating whether markets in which moral hazards obtain are morally objectionable.[2] With an eye to *that* question, I have only been able to find one philosopher who has written on this important topic: Ben Hale.[3] Continuing where Pauly left off, Hale writes: 'Pauly suggests the thesis that I will be defending in this paper – that there is nothing inherently moral about the moral hazard'.[4] Hale argues this by arguing against three hypotheses of what might make a moral hazard inherently unethical: that it is a form of deception, cheating or stealing. Hale concludes that moral hazards are not deception, cheating or stealing. From this, however, Hale concludes that there is nothing inherently unethical about moral hazards and,

[1] Mark V. Pauly, 'The Economics of Moral Hazard: Comment', *The American Economic Review* 58 (1968), pp. 531–37, at 531. Also, Kenneth Arrow, 'Uncertainty and the Welfare Economics of Medical Care', *American Economic Review* 53 (1963), pp. 941–73.

[2] Ronald Dworkin, 'What Is Equality? Part II: Equality of Resources', *Philosophy and Public Affairs*10 (1981), pp. 283–345, at 325.

[3] Benjamin Hale, 'What's So Moral about the Moral Hazard?', *Public Affairs Quarterly* 23 (2009), pp. 1–25.

[4] Hale, 'What's So Moral about the Moral Hazard?', p. 3.

© Cambridge University Press 2013
Utilitas Vol. 26, No. 1, March 2014
doi:10.1017/S0953820813000204

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

moreover, that moral hazards are 'morally neutral'.[5] In this article, I argue that this is not the case. While Hale might be right that moral hazards do not have a moral dimension for the reasons of deception, cheating or stealing, I argue that it is nevertheless not the case that moral hazards are morally neutral.

Most obviously, moral hazards might be inherently morally unethical for reasons other than deception, cheating or stealing. So in order to show that moral hazards are not inherently morally unethical, it is not enough to show that they need not be instances of deception, cheating or stealing. Moreover, there might be something moral about the moral hazard if it turns out that moral hazards are in fact right or good (as a matter of conceptual truth or as a matter of empirical contingencies). I do not argue, however, that moral hazards are inherently morally unethical or inherently morally ethical. Rather, I argue that moral hazards are morally non-neutral and do in fact have a moral dimension, just as *power* is not a morally neutral concept and has an inherent moral dimension.

Hale is right to have (a) brought up this important topic and to have (b) shifted the burden of proof to those who wish to show that the term 'moral hazard' is not simply a misnomer. After all, at least at this point in our history, it is an economic term and not a term from moral philosophy. So why not suspect that it is a misappropriation of terms by one discipline from another? As I already mentioned, I think Hale is also probably right that, when it comes to moral hazards, the moral issue is not one of deception, cheating or stealing. Hale is also philosophically praiseworthy for having shifted the burden of proof to those who wish to show that there is something moral about the moral hazard and that it is not simply a misnomer. In this article, I meet Hale's challenge. But first, let me clarify the terminology.

'Moral hazard' is an economic term which commonly refers to situations in which people have a tendency to increase their exposure to risk when the costs of their actions, should they get unlucky, befall someone else. Within an economic context, we might similarly say that a moral hazard refers to a situation in which there is little or no economic incentive for an agent to guard against risk because the consequences of the agent's actions will befall someone else. Similar definitions have been adopted by others in this literature, definitions which I take to be sufficiently similar for our purposes. For example, Tom Baker writes: 'In the economics literature and in the law and policy debate that draws upon this literature, "moral hazard" refers to the tendency for insurance against loss to reduce incentives to prevent

[5] Hale, 'What's So Moral about the Moral Hazard?', p. 5.

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

or minimize the cost of loss'.[6] Hale adopts a similar definition: 'the term "moral hazard" describes the danger that, in the face of insurance, an agent will increase her exposure to risk'.[7] For the purposes of this article, I take these definitions to be sufficiently similar.

There are two points to notice here. The first is the bit about tendency and the second is the conditional bit about who will pay the costs should things break – break literally if the insured item is my phone or the computer on which this article is being typed, or perhaps break less literally in case of my health or flood insurance.

Let me first make a few remarks about the first bit: the bit about tendency. It is important to notice in any discussion of moral hazards that it is often natural and not necessarily a sign of foul play and bad intentions (or any intentions for that matter) to increase one's exposure once insured, as compared to one's level of exposure to risk when uninsured. Psychologically this can, in many contexts, not be all that different from being very careful with a fragile object in one's hands and being less careful with an object one perceives to be less fragile. Similarly, if I park my car in a neighbourhood I perceive to be dangerous, I might double-check to make sure I have locked my car and left it in a well-lit area and one that is less deserted. I am being more careful in reaction to the dangers and possible consequences as I perceive them. The point of insurance is to ameliorate the dangers and possible consequences – to make my environment safer. If so, then it is a perfectly natural reaction to be less careful in an environment I perceive to be safer. Cases of foul play or insurance fraud aside, it is probably perfectly natural for me to be less on my toes once insured than I would be otherwise. I might well be reacting to my environment just as a healthy cognitive agent should (morally speaking or medically speaking). Whether or not people always behave differently and in the direction of less caution when insured is an open question. The point for the present discussion is merely to assume that at least sometimes human behaviour might indeed change in this way once we find ourselves to be insured and to see what conceptually ethical implications might follow.

In other words, let us assume that moral hazards obtain, at least sometimes if not always. This is an assumption for the sake of argument since the question before us in this article is the morality of moral hazards. If there is no moral hazard, then the question of its morality does not arise. So, I presuppose that there is a moral hazard. This might be a controversial assumption, especially in the context of social

[6] Tom Baker, 'On the Genealogy of Moral Hazard', *Texas Law Review* 75 (1996), pp. 237–92, at 238–9.

[7] Hale, 'What's So Moral about the Moral Hazard?', p. 1.

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

insurance, given some of the empirical evidence. For example, in a research brief summarizing RAND's Health Insurance Experiment we find reported the following results: 'Finally, the experiment examined whether shouldering more of their own health care costs leads people to take better care of themselves. It did not. Risky behaviors were not affected – rates of smoking and obesity, for instance, did not change.'[8] But let us assume that moral hazards actually obtain: that people who are uninsured actually engage in less risky behaviours than people who are insured and that there is a causal connection of the right kind between being insured and engaging in risky behaviours. When such circumstances obtain, are the resulting moral hazards morally neutral? As I will argue, they are not.

Second, let me make a few remarks about the conditional bit: the bit about who pays the costs should things break – break figuratively or literally. If I unintentionally break my insured phone, then the insurance company pays. I might pay too because there might be a deductible or an increase in the premium. But whether or not I pay given how the insurance policy is set up, the insurance company pays or shares my costs. Let us focus on the insurance company's costs and formulate this as a conditional claim: if something breaks (literally or figuratively), then the insurance company pays.

Let us combine the tendency bit with the conditional bit. Furthermore, let us add the assumption that an insurance company insures lots of people like me; this assumption helps turn probabilities into frequencies and thus turn exposure to risk into actual costs as opposed to a mere danger of costs. Given these three ingredients – the tendency bit, the conditional bit and the large numbers bit – we get the following: if people were just as careful once insured (when perhaps they feel safer) as compared to when they are uninsured (when perhaps they feel less safe), then the insurance company's costs would be less. Since insurance costs can translate into higher or lower premiums for everyone, this means that moral hazards can translate into higher insurance premiums than the baseline, where the baseline is people behaving as they would if uninsured. In the case of social insurance (e.g. FEMA and flood insurance), this can mean higher or lower tax rates. This perhaps becomes a sore point if some people are actually more careful than others, or if some people need insurance whereas others do not need it because they have enough money to cover their own costs should things go wrong.

[8] RAND, 'The Health Insurance Experiment', The Rand Corporation, <http://www.rand.org/content/dam/rand/pubs/research_briefs/2006/RAND_RB9174.pdf> (2006), p. 4.

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

The important point in all this is that it is an open question whether you in fact have a moral duty or moral responsibility to be just as careful as you would if you were uninsured. For example, being able to relax might well be the very reason you purchased insurance. Moreover, being able to relax and increase your exposure to risk might well give you access to something to which you should be entitled, in which case the given moral hazard would be required as a matter of justice. Consider a quick example. This might well be the case if:

(a) in low-income families, parents enrol their children in school sports when and only when their children have medical insurance, and
(b) all children regardless of family income are entitled to participate in school sports (given, for example, its overall health benefits or given that sports are a central cultural value in our society).

Given all this, how could moral hazards have a moral dimension? Wouldn't it all depend on the particular circumstance, especially given that the tendency to increase one's exposure to risk once insured might be perfectly natural rather than a case of inherently irresponsible behaviour (e.g. drunk driving) or foul play (e.g. knowingly and intentionally committing insurance fraud)? As I will argue next, moral hazards do in fact have a moral dimension even if they are not inherently ethical or unethical. The most obvious argument here is that moral hazards affect the overall allocation of goods and bads. Namely, moral hazards affect the distribution of costs. If any allocation of goods and bads or, more specifically, a distribution of costs, requires moral justification, then so do moral hazards. The justification might depend on other facts – the fact that there is a moral hazard might well not be enough to conclude that a given distribution of costs is justified or unjustified, just or unjust – but such justification must be forthcoming.

Moreover, if negative externalities require moral justification, then a moral hazard requires moral justification when it is an instance of a negative externality. An 'externality' is an economic term which can best be described as transaction spillover. In the case of a 'negative externality', the spillover has to do with costs rather than benefits. The term 'negative externality' describes a situation when an agent does not have to pay the full cost associated with her behaviour because those costs, in full or in part, befall someone else. Transaction spillover is typically thought of as costs (or benefits) spilling over to parties that were not involved in the exchange. However, unexpected or unjustified costs could also spill over as a result of the transaction from one of the two parties involved in the transaction to the other. For example, if you hire a babysitter but as a result end up paying the babysitter's

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

speeding tickets, this might be a negative externality if the costs of the babysitter's actions unjustifiably spill over to you.

In the context of insurance or, more broadly, market transactions, we might specifically worry about a case in which two agents engage in a market exchange as a result of which one agent incurs the risk for another agent. Insurance is a particularly interesting case of this because the point of insurance is indeed to transfer risk (and also spread risk!), which means that there might be an important voluntary component that prima facie justifies why one agent should incur the costs associated with *another* agent's actions (e.g. because these costs are associated with this agent's actions – the action of taking on the liability for the costs of another) when, for example, the risk transfer is voluntary or perhaps is justified on some paternalistic grounds of what would be voluntary if the agent were only adequately informed. In what follows, I argue that moral hazards have an inherent moral dimension and that the concept of a moral hazard is not value-neutral. The gist of the argument is that a moral hazard can be a negative externality, negative externalities always require moral justification, and something that requires moral justification is not morally neutral; therefore, a moral hazard is not morally neutral when it is an instance of a negative externality. And if it turns out that moral hazards are always negative externalities, then moral hazards always require moral justification.

This article will unfold as follows. In section II, I draw some relevant distinctions with respect to value-neutral and morally neutral concepts. In section III, I argue that moral hazards are morally non-neutral: I discuss the aspect of moral hazards that has to do with responsibility as a moral concept and argue that moral hazards are not value neutral and do in fact have an inherent moral dimension. In section IV, I briefly suggest that markets in which moral hazards obtain might be in one respect unfair. Section V concludes.

II. VALUE-NEUTRAL CONCEPTS

We might call a concept 'value-neutral' or 'value-laden'. Domination and innovation, for example, are both obvious examples of value-laden concepts because both have three features: (a) they both call for a value judgement, (b) that value judgement must be a moral one, and (c) the moral judgement is fixed in the sense that it is either positive or negative depending on the concept itself and no other facts.[9]

[9] For the purposes of this article, I am intentionally keeping away from employing the famous distinction between thick and thin concepts. For more on the thin/thick distinction, see Bernard Williams, *Ethics and the Limits of Philosophy* (London, 1985),

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

For example, domination has a negative connotation while innovation has a positive connotation. Thus, domination and innovation are both obviously value-laden concepts. Some people use the word 'power' to mean *domination*. In that case, the concept of power would also possess all of the above three characteristics. If, however, the word 'power' is used more broadly to mean *authority* (as I will use it in the rest of this article), then power might be morally justified or unjustified, legitimate or illegitimate, have good consequences or bad, and thus be good or bad – depending on context, i.e. depending on other facts. If so, power would only have features (a) and (b) but would lack feature (c). Whether or not we should call this 'value-laden', I would not call it 'value-neutral'. Instead, I would call this kind of concept 'not value-neutral' and 'having a moral dimension' (but not a specific moral value, i.e. a value of good or bad in and of itself inherent to the concept). A chair and a hammer, on the other hand, are both concepts that might possess only feature (a) if we might wonder about their usefulness and utility but not in moral terms. This is so at least given how chairs and hammers are typically used: to sit on, to fix things and build things with, and perhaps as art objects. The word 'gun' might by some people be used in a way that only rings of feature (a) – they might conceptualize a gun as a tool which can have positive utility or negative utility but not a value that is necessarily moral (e.g. perhaps they usually use guns for hunting and don't think that animals should be in the sphere of proper moral concern). Other people might use the word 'gun' to denote a concept that possesses features (a) and (b) because, for example, they believe guns are used for hunting animals and believe that not only should animals be in the sphere of proper moral concern but that proper moral concern means that animals should not be killed recreationally or for food. Others might conceptualize a gun as mainly a tool used to kill and sometimes murder other people. Some people might think of the word 'gun' as possessing all three features because they think of guns as having the purposes of ending life, human or not, and they might think that ending life is always morally impermissible and/or states of affairs in which life comes to an end at the hand of another are morally bad. Some concepts, however, do not possess any of the above three features. Space, for example, as a purely physical concept appears to be devoid of features (a), (b) and (c), and I think this appearance is correct because space simply *is*. It is a truly descriptive concept for it does not call for any value judgements whatsoever. In addition to space, time might also be devoid of any value dimensions, at least as a physical concept, let alone moral dimensions. The same might be said of spacetime.

as well as Samuel Scheffler, 'Morality through Thick and Thin: A Critical Notice of Ethics and the Limits of Philosophy', *The Philosophical Review* 96 (1987), pp. 411–34.

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

Chairs, on the other hand, possess feature (a). Power, which might or might not be morally justified, along with guns perhaps, depending on the conception, possesses feature (b) in addition to (a). Domination and innovation typically both possess feature (c) in addition to (a) and (b), thus possessing all three features. Concepts that possess all three features are not value-neutral; instead, they are value-laden. Concepts that possess features (a) and (b), but not feature (c) have a moral dimension even if they do not have a fixed moral value. Any concept that possesses only feature (a), or not even that, does not have a moral dimension. The key point here is that a concept lacks a moral dimension if and only if that concept lacks features (b) and (c).

Because the context of the present article is about *moral value* and not just value of utility (e.g. the utility of chairs as measured in their sheer usefulness or aesthetic value), I will use the word 'value-neutral' to mean value-neutral with respect to moral value. Thus, I shall call a concept 'value-neutral' if and only if it does not possess features (b) or (c). In other words, if a concept possess feature (b), even if it does not possess feature (c), it will not be a value-neutral concept in the context of moral inquiry and will have a moral dimension (even if it will not have a fixed moral value). As I will argue in this article, the concept of a moral hazard possesses feature (b). Hale's contribution, as I understand it, is to have provided good evidence that moral hazards probably do not possess feature (c). As I will argue next, however, moral hazards possess feature (b), just as the concept of power does, even when 'power' is not used to mean domination.

III. RESPONSIBILITY

The argument I offer is one from responsibility.[10] If we are to take the idea of individual responsibility seriously, then any case where the costs of one's actions befall someone else and any transfer of risk require justification. By 'taking the idea of individual responsibility seriously', I mean the idea that each person is prima facie morally responsible for the consequences of her actions and bears the burden for associated costs. Consider, for example, the following line from Ronald Dworkin, which here I take somewhat out of context merely to illustrate the sentiment: 'We live our lives mostly like swimmers in separate demarcated lanes.'[11] Taken as an expression about the way things should be or about what is prima facie (or pro tanto)

[10] Perhaps the doctrine of individual responsibility has been overextended; that, however, is not the concern of the present article since in that case this doctrine will have been overextended elsewhere; the aim of this article is simply to shift the burden of proof to the other camp back from where Hale has placed it.

[11] Ronald Dworkin, *Justice for Hedgehogs* (Cambridge, Mass., 2011), p. 287.

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

morally justified, this way of thinking introduces a default baseline for a just distribution of costs deviating from which requires moral justification.[12] Here, I do not necessarily endorse this baseline, but merely make it explicit in connection with discussions of the morality of moral hazards. This baseline of course is not a very social one but, instead, a rather individualistic one.

Along similar lines, we might think that it is unfair to impose the costs of one's actions on another. My claim here is merely conditional: *if* we take the idea of individual responsibility seriously, then moral hazards require moral justification. This already introduces a moral dimension to moral hazards, at least in the United States where the normative doctrine of individual responsibility is very culturally and philosophically salient. In other words, we already have a modus ponens because, as a matter of descriptive fact, we do take the idea of individual responsibility seriously. Therefore, given the assumptions, moral hazards require moral justification.

This does not yet mean that moral hazards are unethical or unjust or that someone has been wronged or is even at risk of being wronged. Nor does it yet mean that moral hazards constitute a morally bad state of affairs. But it does already introduce a moral dimension and a moral concern. So, Hale is right to say that situations involving moral hazards are often not morally problematic and that for this reason, especially in the context of pure economic analysis and public policy, the term is somewhat misleading. On the other hand, moral hazards do in fact require moral justification, which of course is often forthcoming in the kinds of contexts that Hale considers because, for example, those contexts involve transfer of risk via voluntary exchange when the issue is not one of social insurance but rather a transfer of risk between private parties. For that reason, moral hazards do indeed have an inherent moral dimension.

Let us now consider the more specific but salient case of insurance. Voluntary transfer of risk might be one justification for a transfer of risk and thus for a moral hazard. Whether a transfer is voluntary or not will in part depend on what the two parties agreed to. Misinformed or uninformed choice, for example, can in many contexts undermine the

[12] By saying that 'X is prima facie morally justified', I mean that by default we are epistemically justified in believing that X is morally justified unless we hear otherwise, akin to a presumption of innocence. By saying that 'X is pro tanto morally justified', I mean that there is something unfortunate, morally speaking, in deviating from X even if, all things considered, deviating from X is the right (or the good) thing to do. See Shelly Kagan, *The Limits of Morality* (Oxford, 1989), pp. 17 and 47, as well as W. D. Ross, *The Right and the Good*, reprinted with an introduction by Philip Stratton-Lake (Oxford, 2002).

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

voluntariness of a choice or even invalidate choice altogether. While the inherent idea behind insurance might not necessarily be to mitigate uncertainty, without uncertainty actual insurance would not be called insurance but would simply be called compensation (even if it could be justified by appeals to hypothetical insurance). In other words, uncertainty is what makes insurance markets possible. But what is the relationship between informedness and uncertainty? Is not the very nature of uncertainty such that those acting under it are not fully informed? It is. However, the relevant parties must still be adequately informed with respect to some relevant facts. For example, the insurer (e.g. the insurance company) and the insuree (e.g. the person seeking to be insured) must be on the same page when it comes to what it is that is being insured, what the coverage is, and what each party's liabilities and costs will be. Otherwise, their contract might not be informed in the relevant sense.

The *level* of risk is something relevant to the informedness and the voluntariness of the exchange. If, for example, an agent buying medical insurance thought that she was purchasing enough coverage to cover $100,000 of medical bills but, too late to do anything about it, found out that she owes the hospital that exact amount but is only covered up to $25,000 and now must find a way to cover her remaining $75,000, this would wash away the relevant informedness and some of the voluntariness of the market exchange that took place when she purchased insurance. The same is true of the insurer. If you decide to run an insurance business and think you agreed to insure me for $25,000 (which, let us assume here for the sake of argument, is a considerable sum for you), but then find out that you are on the hook for four times that amount, your poor epistemic state might indeed undermine the voluntariness of our market exchange.

There is of course the further question of who ought to be responsible for your epistemic state in this case and whether, perhaps, you are blameworthy for it and if so, up to what amount you should be liable for your mistake: e.g. should there be a public fund for these kinds of epistemic accidents or should you and I split the difference if it was an honest mistake? By 'should' I do not here mean 'What does the law have to say?' but rather 'What should the law say?' or 'What is the right thing to do regardless of what the law should be?' The answer to that question, however, is not forthcoming in this article. And whatever the answer is, the point is that the degree of liability which one took on matters because it can impose higher costs on an agent than she had anticipated: someone buying insurance might end up not having reduced her risk to the degree that she had thought she did and someone offering insurance might end up having taken on more liability than he had anticipated.

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

Imagine a hypothetical illustration of the frustrations that might spring up as a result of moral hazards. Tom and Jerry decide to offer cell phone insurance in case your phone breaks. They reason as follows: we know that 4 in 100 phones break by accident without insurance and so we will offer rates with a thin profit margin with this number in mind as an alternative to all these scam insurance companies that overcharge their customers. Assume that the 4 in 100 statistic is correct. Moreover, they leave a buffer to be able to handle a 5 in 100 failure rate in case they are wrong in their estimate. We are inexperienced and could be off by 20 per cent in our estimates, they reason. However, once Tom and Jerry start offering their highly affordable insurance and sell lots of insurance policies, the rate of accidents among their customers goes up from 4 in 100 to 8 in 100. Assume this is not because of insurance fraud or ill intent but rather simply because people are less careful with their phones now that they have insurance. Assume also that this means that Tom and Jerry have to close shop as this breaks their bottom line.

Consider another thought experiment. You start a car insurance company by insuring 100 drivers who you know do not drink and drive. Assume that you know that drivers who do not drink and drive get into much fewer and much less costly accidents than those who do drink and drive and that you also know the relevant statistics. You build your business plan around these statistics and, to offer competitive rates, you make your profit margin somewhat slim while leaving what you believe enough of a reserve for a skewed distribution of accidents given that the number of drivers you are insuring is not that big statistically speaking. You responsibly do your maths and balance your budget. After you insure your drivers, however, half of them start drinking and driving and, as a result, they all predictably get into costly accidents (even though you had not predicted that they would start drinking and driving in the first place, e.g. that they would become that kind of driver). Whatever their motivations and reasons and whether or not they are even aware of their own motivations of reasons, you might feel this is not what you agreed to and that your clients started acting irresponsibly at your expense. After all, you thought you were insuring low-risk drivers (e.g. 'responsible' drivers), but it turns out you ended up being liable for high-risk drivers (e.g. 'irresponsible' drivers). Whatever the legal duties of each party in this case, the moral question is now obvious: did the others act irresponsibly and, if so, is this an unfair imposition of costs on you? For example, if you or both parties were not 'on the same page', then who is responsible for whose epistemic state? To disambiguate between two different senses of 'responsible', who is morally responsible and who should be liable for your and possibly the drivers' poor epistemic states with respect to the agreement

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

that you made and thus with respect to the unmet and even violated expectations of each? This is a moral question to be sure. And it arises precisely because the risks associated with the behaviour of one party befall another. In other words, is the negative externality in this case morally justified? If, by default, we assume that negative externalities require moral justification and this is a negative externality, then this moral hazard will require moral justification. Such justification might indeed be forthcoming and might even always obtain in certain types of cases and certain types of markets: e.g. markets with all voluntary and adequately informed market transactions. One way to frame the moral issue is: what does responsibility require in case of moral hazards? We could ask this question specifically in the context of insurance cases when the two parties have failed to control risk relevantly and have introduced new risks which stem from uninformed or misinformed market transactions. Who is responsible for whose costs in this case, who is responsible for whose epistemic state, and *is it irresponsible* to increase one's exposure to risk simply because one is now insured? This again is a question that outlines the moral dimension of moral hazards.

The case of social insurance (e.g. welfare safety nets or state health insurance) is not dissimilar if we focus on the question of what it is to treat individuals as responsible agents and what it is to be a responsible agent. Would it be irresponsible of me to increase my exposure to risk or to behave in more risk-prone ways simply because I am now insured? And, similarly but as a distinct question, would a state that over-insures its citizens fail to treat them as responsible adults? Would we, as members of a democratic state of this kind, be treating each other and ourselves as responsible adults if we were to over-insure each other? Again, these are moral questions that arise specifically in the context of moral hazards.

What does it mean to 'over-insure' in this context? To answer this question we would need a baseline. There are three candidates. The first is the amount of risk to which someone exposed herself before becoming insured; this denotes past behaviour. The second is the amount of risk to which someone would have exposed herself had she not become insured; this denotes counterfactual behaviour. The third is the level of risk exposure given moral hazards. If, or in cases when, this third candidate is not the right normative baseline, then moral hazards will require additional moral justification.

The question of responsibility for over-insurance is a moral question. Let 'over-insurance' be defined as the difference between actual liability due to risk and baseline liability. Whichever of the first two baselines is the normatively relevant baseline, over-insurance is the difference between the agent's actual risk exposure once insured and the baseline. If so, then it is an important moral question how the costs of

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

over-insurance should be distributed. For example, should they all be assigned to the insurer, the insuree, shared by the two and how exactly, or also shared by third parties (e.g. society as in the rest of us)? The answer to this question might in part depend on answers to other questions (e.g. was there deception and if not, what is a just assignment of responsibilities for epistemic states?) But what matters for the purposes of this article is that the moral question is inherent to the concept of a moral hazard.

IV. UNFAIRNESS

In this section, I argue that markets with moral hazards might be in one respect unfair. This conclusion obtains under very specific assumptions about circumstance. Such assumptions, however, might often be true of the real world given certain epistemic limitations that often obtain in real markets.

Assume that the insured indeed have a moral responsibility not to increase their exposure to risk once insured. Or, to be more precise, assume that the insured have a moral responsibility or moral duty to keep their exposure to risk at least as low as they would if they were uninsured. If so, as I argue in this section, we will face a distinct moral problem in a market in which moral hazards obtain. Given reasonable epistemic assumptions along with the assumption that some of the insured will be behave responsibly and some will not, the allocation of costs and benefits will be unfair to those of the insured who act responsibly.

Markets in which moral hazards obtain might come with a basic epistemic limitation: if different types of 'sellers' exist (which in this case are the persons seeking to be insured), then a 'buyer' (which in this case is the insurance company) might not be able to distinguish among them.[13] Or, equivalently, different types of 'sellers' cannot make themselves distinguishable to the 'buyers'.[14] Some individuals (the 'sellers') might increase their exposure to risk once insured or already have increased exposure to risk while others might not. In other words, moral hazards can come with an inherent epistemic limitation if it is not possible to tell whether risk obtains because of bad brute luck or

[13] This same epistemic limitation can be found in markets that economists call 'lemons markets'. For a seminal paper on 'lemons markets', see George Akerlof, 'The Market for "Lemons": Quality Uncertainty and the Market Mechanism', *The Quarterly Journal of Economics* 84 (1970), pp. 488–500.

[14] I refer, somewhat counter-intuitively, to those who are insured or are seeking insurance as 'sellers' and to insurance companies (or, similarly, the state should the state be providing insurance instead of private firms) as 'buyers' because this keeps the terminology parallel to what we might find in economic discussions of lemons markets.

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

because it is traceable to some element of the agent's behaviour for which the agent could perhaps be held responsible. If we assume that a market with moral hazards has both of these underlying causes of ensuing costs, then this kind of market would be unfair to some 'sellers' and thus inherently unethical. The costs of the 'irresponsible' behaviour of some insured would be transferred not necessarily to the insurance company, but to the more 'responsible' insured individuals.

The basic epistemic limitation I mentioned in the previous paragraph means that even if an insurance company can charge differential rates to properly incentivize 'responsible' behaviour, i.e. behaviour that does not increase exposure to risk when insured or which increases such exposure to a lesser degree rather than a greater degree, it cannot truly reward it. For example, a car insurance company might offer a 'safe driver discount'. The insurance company, however, cannot know whether any given customer who qualified did so because he behaved responsibly (e.g. because he was incentivized to drive safely) or because he merely got lucky despite driving recklessly. If we thought that, for deontic reasons, 'responsible' behaviour should be rewarded, then epistemic limitations make it impossible to reward those of the insured who behave 'responsibly' with sufficient degree of precision if we are to reward each individual if and only if he deserves the reward (e.g. is entitled to be rewarded for responsible behaviour).

This creates a distinct moral problem. If it is really true that at least some of the increased exposure to risk is under the insured agents' control and if some agents really behave more responsibly than others, then the outcome can be unfair to those who behave more responsibly. That's because, ex hypothesi, the insurance company cannot distinguish between those who act responsibly and those who do not and thus cannot reflect this difference in insurance premiums. This is a moral problem distinct from the one I have already mentioned and is one I mention here only as a side-note. It is beyond the scope of this article to defend this line of reasoning more thoroughly. If, however, this line of reasoning can be defended more thoroughly, then this would shift markets with moral hazards from (b) a concept that calls for a moral value judgement but does not have a fixed moral value independent of other facts to (c) a concept where the moral judgement is fixed and, moreover, one that is negative (see section II for this distinction).

In other words, it is elucidating to distinguish between two types of markets in which moral hazards can obtain. (This distinction is not meant to be jointly exhaustive as other types of 'moral hazards' markets are obviously possible.) The first is a market in which everyone's exposure to risk increases once insured, perhaps simply as a matter of some psychological fact about human nature and, moreover, perhaps not due to anything for which any insured agent can be properly held

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

responsible. The second is a market in which exposure to risk differs across the insured population and, moreover, where at least some significant part of this difference can be traced to elements of behaviour for which it would be appropriate to hold the insured parties responsible (or liable) but for which, given our basic epistemic limitation, we are unable to hold them responsible (or liable). The second type of market is an extension of the first. As I have argued in section III, in the first type of market we risk having costs unfairly imposed on insurance providers, costs which should instead befall the insured. This imposition of costs can be justified, but such justification must be forthcoming. Some story about voluntary or involuntary transfer of risk could provide such justification. As I have suggested in this current section, the second type of market is more morally problematic because it creates an unfair imposition of costs which would be more difficult to justify: it creates a situation in which the costs of less responsible insured agents are unfairly imposed onto more responsible insured agents.

V. CONCLUSION

Moral hazards inherently have a moral dimension, although Hale is probably right that this dimension is not one of deception, cheating, or stealing. None of what I have argued here implies that moral hazards are inherently immoral. Immoral, however, is not the same thing as morally neutral. Contrary to Hale's and Pauly's brave claims, moral hazards do in fact have an inherent moral dimension, just as power, for example, has an inherent dimension. Power can be justified or unjustified, legitimate or illegitimate. But whichever it is, it naturally invites questions of justification and legitimacy; power always requires justification. The same is true of moral hazards for two main reasons intimated in this article. First, presupposing an individualistic framework, liability for the actions of others always requires justification so as to avoid an unfair imposition of the costs of one's situation and choice onto others – the question of distributive justice. In the case of moral hazards, there are two liabilities that require justification: (a) liability for the baseline level of risk, which can perhaps be justified via voluntary market exchange, and (b) liability for the difference between baseline liability and actual liability. Both the baseline and the additional liability – liability above the baseline – require justification. Second, there is the question of what it is to be a responsible agent – the question of moral character. Third, there is the question of what it is to treat ourselves and each other as responsible agents – the question of relational justice and a relational ethic. These three questions outline not only one moral dimension of moral hazards, but its three dimensions. In short, moral hazards are not

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

morally neutral, even though they could well be morally justified or be morally legitimate depending on the case. However, if the assumptions of section IV are not true, then Hale is absolutely right that moral hazards are not intrinsically unethical, just as power is not intrinsically immoral, and that, given that the concept is almost never used or mentioned in moral philosophy and was coined and is indeed used in the context of economics and public policy, it has a distracting name. It is not the case, however, that moral hazards are morally neutral; they are not. That being said, their moral dimensions are better left to contexts that are appropriate to the exploration of their moral dimensions, which a course in economics might not be. In other words, we should distinguish between value-neutrality and moral badness: moral hazards are not value-neutral but it is true that, barring the suggestive remarks of section IV, they might not be inherently morally bad.

To recap, this article unfolded as follows. In section II, I drew some relevant distinctions with respect to value-neutral and morally neutral concepts. In section III, I discussed the aspect of moral hazards that has to do with responsibility as a moral concept and argued that moral hazards are not value-neutral and do in fact have an inherent moral dimension. In section IV, I briefly suggested that markets with moral hazards might be inherently unfair and thus inherently unethical for that reason. In this current section, I conclude.

The issues I raised in this article make salient three clusters of important questions. First, do insured individuals have a moral responsibility or moral duty to keep their exposure to risk at least as low as they would were they uninsured? I think clearly not as there are counterexamples I personally find convincing (e.g. the example of school sports in section I). However, the question then becomes: *when* do insured individuals have a moral responsibility or duty to keep their exposure to risk at least as low as they would were they uninsured? Second, who is responsible for whose epistemic state? Or, similarly, what is a just allocation of information costs and what is the relationship between voluntariness and informedness (e.g. can a choice be voluntary despite being misinformed)? For example, what if I insured you but did not know that you would have a tendency to increase your exposure to risk? Third, what are the moral implications of moral hazards and related markets (e.g. lemons markets) under conditions of the kind described by the information asymmetries and epistemic limitations in the market described in section IV? What makes the moral question of moral hazards especially important and interesting is not merely that moral hazards lie in that uncomfortable space of the polarization between the American political right and left, especially in the context of social insurance when taxes are involved and thus the justification of voluntary market exchange is unavailable. Rather, the

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press

question of the morality of moral hazards points to the three question clusters I just mentioned, which require more attention than they have received in the philosophical literature so far.

In this article, I hope to have met Hale's challenge to provide an argument that the concept of a moral hazard is not value-neutral and that moral hazards are morally non-neutral. As I suggested in section I, at its most basic, this is because moral hazards affect the allocation of costs (and benefits) and such allocations always require moral justification. Perhaps some allocations of costs are prima facie justified (or perhaps even pro tanto justified). This, however, would require an argument which addresses the difficulties raised by this article.[15]

wbraynen@stanford.edu

[15] I would like to thank all the members of the postdoctoral workshop at Stanford University's McCoy Family Center for Ethics in Society for their comments on earlier drafts of this article, as well as Drew Hutchins.

https://doi.org/10.1017/S0953820813000204 Published online by Cambridge University Press