

The Varieties and Dynamics of Moral Repugnance: Prediction Markets and Betting on Matters of Life and Death

Dan Weijers & Vadim Keyser

Abstract

In this paper, prediction markets that encourage traders to bet on matters of life and death are used to explore the varieties and dynamics of moral repugnance. We define moral repugnance as morally charged feelings of revulsion that correspond (correctly, incorrectly, and indeterminately) to moral reasons and contexts. Rich variations of moral repugnance and their dynamic qualities are presented by investigating the contextual frames in which they arise. These contextual frames constitute interacting conditions composed of information about states of affairs, moral reasons, and feelings of revulsion. Through careful study of two medical prediction markets that encourage betting on death, we can observe the interaction between these causal conditions to see *how* the varieties of moral repugnance emerge. We also present three interesting results that arise from analyzing the dynamics of moral repugnance in response to prediction markets. First, a prediction market can elicit several conflicting types of moral repugnance at the same time. Second, moral indeterminacy can arise in two different ways when judging prediction markets. Finally, some prediction markets can generate a moral endogeneity problem, a disruptive feedback loop between a given prediction market and the morally relevant outcome it predicts.

Keywords: repugnance, moral repugnance, prediction markets, betting on death, moral indeterminacy, endogeneity problem, CrowdMed, Iowa Electronic Health Markets

1. Introduction

Betting on matters of life and death is widely considered to be morally repugnant, which is to say that the idea of such betting causes many people to feel morally charged revulsion. Consider celebrity death pools, in which people bet (and win real money) on lists of famous people that they expect to die in the next year. Breitbart and Ebner (2004) identify the death pool website www.stiffs.com¹ as the most extreme case of schadenfreude that they have seen, pointing out that the website is replete with jokes at the expense of the deceased. For example, the site announced the death of Albert Broccoli with the quip: “Just as well. It's awful to think of anyone spending the rest of his days as a vegetable.” Indeed, to say that the site is unashamedly disrespectful of the dead is an understatement. The site’s homepage informs us that: “We've got multiple games, email alerts when the famous ones hit the dirt, and tons of other sick fun for the whole family. Come on in and have a look around ...” (www.stiffs.com).

Clearly not everyone finds betting on matters of life and death repugnant, since some people bet on death pool websites. But small numbers of people do many strange and, in the eyes of the vast majority, immoral things. Important here, is that the vast majority of people find betting on matters of life and death repugnant and they judge it to be immoral. Furthermore, widespread repugnance about betting on matters of life and death occurs even when the betting will not influence those life

¹ If you want to know if you are famous enough to be (to quote from www.stiffs.com) “good to go”, then you can check the celebrity database here: www.stiffs.com/celebrity/directory/. Several academics make the list.

and death matters in any way. Of course, if the betting is thought to influence the life and death matters, then it is even more widely viewed as repugnant.²

We use prediction markets in this paper to showcase the varieties and dynamics of moral repugnance that can arise from betting on death. . Modern prediction markets are websites in which anonymously registered traders buy and sell shares in predictions about real-world outcomes (Weijers 2013a). Prediction markets usually pay out a set fee (e.g., \$10) to traders who hold shares in a prediction that turns out to be true. For example, a trader might purchase shares in the prediction that ‘there will be a power outage affecting at least 1 million people in the United States in 2013’ for \$2 per share because she thinks that the shares are under-priced—that such a power outage is more likely than 20% ($\$2/\$10 = 20\%$ chance). If the trader holds on to the shares (instead of selling them to another trader), then she stands to make \$8 per share if the prediction is true, or lose \$2 per share if the prediction turns out to be false.

In this discussion, prediction markets that encourage traders to bet on matters of life and death are used to explore a moral phenomenon—moral repugnance, which constitutes morally charged feelings of revulsion that correspond (correctly, incorrectly, and indeterminately) to moral reasons and contexts. We present rich variations of moral repugnance and their dynamic qualities by investigating the

² Some forms of death pools, such as a death pool in Taiwan about when local elderly people will die, plausibly influence the deaths of the people who are bet on because those who care for the elderly can participate and have considerable financial incentives to withhold life-saving care. See: “Betting on when late-stage cancer patients would die, doctors and family members wagered over 1 billion [TWD]” for a discussion and comments by readers expressing repugnance. Available from:

<http://www.chinasmack.com/2013/stories/underground-gambling-on-when-cancer-patients-will-die-exposed.html>.

contextual frames in which they arise. The prediction market-user relation determines sets of causally interacting conditions composed of information about states of affairs, moral reasons, and feelings of revulsion, collectively labelled ‘contextual frames’. These contextual frames determine the varieties of moral repugnance that emerge from a prediction market. Through careful study of a given prediction market we can observe the interaction between these conditions to see *how* the varieties of moral repugnance emerge.

In this paper, we provide a classificatory scheme of varieties of moral repugnance and analyse the dynamics of repugnance within prediction markets by *systematically* tracking shifts in contextual frames. We also present three interesting results that arise from analyzing the dynamics of moral repugnance in response to prediction markets. First, a prediction market can elicit several conflicting types of moral repugnance at the same time. Second, moral indeterminacy can arise in two different ways when judging prediction markets. Finally, some prediction markets can generate a moral endogeneity problem, a disruptive feedback loop between a given prediction market and the morally relevant outcome it predicts. When the moral endogeneity problem affects a prediction market, the information used to make a bet in that prediction market determines the predicted likelihood of the event, which, in turn, affects the morality of using the information to make the bet. Although we exclusively use prediction market examples, the following discussion of the varieties and dynamics of moral repugnance could be applied to many social and ethical issues.

The structure of the paper proceeds as follows. In Section 2, the varieties of repugnance are discussed, especially as they pertain to complex moral information and contextual frames. In Section 3, two kinds of prediction markets that encourage betting on matters of life and death are explained in terms of contextual frames.

These are prediction markets that predict epidemics and prediction markets that predict medical diagnoses. In Section 4, three important results are discussed: conflicting moral repugnance, indeterminate moral repugnance, and the moral endogeneity problem. Finally, Section 5 summarises the paper and discusses the implications.

2. The Varieties of repugnance and ‘contextual frames’

Leon Kass (1998, p. 687) described repugnance as: “[when] ... we intuit and feel, immediately and without argument, the violation of things that we rightfully hold dear.” Kass famously defended repugnance as a source of wisdom by arguing: “Revulsion is not an argument; and some of yesterday's repugnances are today calmly accepted—though, one must add, not always for the better. *In crucial cases, however, repugnance is the emotional expression of deep wisdom, beyond reason's power fully to articulate it.*” (Kass 1998, p. 687; emphasis added). Kass's concept refers to a context-relative violation of values. It is important to note that such a violation, in the form of general repugnance, need not be moral. For example, people may find the smell of a specific food repugnant. This is a context-relative judgment about food smells. There are often reasons associated with the repugnance, but in many cases, like food repugnance, the judgement does not occur in moral contexts or is not entangled with moral reasons. In this discussion we focus on moral repugnance, which is a specific type of repugnance in which the feelings of revulsion occur in relation to moral contexts.

Moral repugnance involves an interesting complexity: The context-relative judgment often corresponds (although sometimes not clearly) to moral reasons for making that judgment. These reasons need not be known or understood, and as such can be explicit or implicit reasons. Furthermore, we propose that moral repugnance

is a relational property of a moral phenomenon. That is, moral repugnance is produced through the causal interaction between information about states of affairs, moral emotions, and moral reasons. Each causal condition will be analysed and illustrated shortly. This makes moral repugnance context-sensitive. Specifically, depending on the types of information, emotions, and reasons within a given context, the repugnance felt by an individual will take on a different character.³ Moral repugnance and context-sensitivity (sometimes referred to as ‘contextual frames’) are explored in detail below.

Moral repugnance occurs when person X feels morally charged revulsion about thing T, and this charged feeling is sensitive to the context in which moral information is presented.

The morally charged aspect of moral repugnance comes from the fact that the feeling of revulsion is about a moral issue, such as what someone (including ourselves) has or hasn’t done, or may or may not do (Haidt 2001). So, moral repugnance is an intuitive feeling of revulsion that is cognitively associated with a phenomenon (the target of the moral repugnance), when that phenomenon is considered to lie in the moral domain by the person experiencing the revulsion. To adapt an example from Weijers and Richardson (2014), thinking about the hypothetical manager of Safety First Autos, who knowingly disregarded safety advice in order to raise short-term profits, is likely to elicit moral repugnance; in response to considering this case, we likely experience an intuitive feeling of revulsion and attribute that revulsion to the behaviour of the manager (which we consider to lie in the moral domain).

³ Here, we are not committed to our model being realist or antirealist. It might be that there is a corresponding truth that these context-relative judgments can latch on to. We are not concerned with this. Rather, we are concerned with the relation between information, moral emotion, and moral reasons, and how those factors interact to produce varieties of repugnance. As we will see in section 4, this account may have an aspect of perspectival realism, where different perspectives are indispensable in giving us certain information to make moral judgments.

Furthermore, this feeling emerges within the context of knowing certain information—e.g., specific information about the nature of the manager’s disregard for safety.

Despite being based on intuitive feelings, which are notoriously hard to analyse with introspection (Weijers 2013b), repugnance can be morally justified by *moral reasons*. For example, Kass (1998) argues that human cloning elicits repugnance because it is an affront to human nature and normal human reproduction. However, Kass (1998) also argues that, even when it cannot be justified, *widespread* repugnance by itself is a strong, *prima facie* winning, moral argument—i.e. when the majority of people find something repugnant then the burden of moral argument falls on the proponent of the thing in question. In such reason-lacking scenarios, the implicit assumption is that there is some hidden moral reason, corresponding to the repugnance. This reason need not be known at the time of the feeling of revulsion, so the argument goes, because it will likely become obvious at some point, especially if the purportedly repugnant phenomenon becomes commonplace. But it is not always the case that the assumed reasons associated with the feeling of revulsion properly correspond.

There are cases in which even widespread moral repugnance is unwise because it does not correspond to the right set of moral reasons, such as when the repugnance is based on widespread false moral or factual beliefs.⁴ Even when repugnance seems like a good indicator of moral reasons, we cannot know it without carefully analyzing the reasons, and the states of affairs those reasons relate to, directly. For example, a child might find it morally repugnant that a stranger is wading in ‘the sacred pool’, that is, until he asks his parents about it, and hears that

⁴ See Kimberly (2002) and Niemela (2011) for further opposition to the wisdom of repugnance.

the stranger is the messiah, finally returned. Therefore, anyone experiencing repugnance about something should be open to receiving new information about it in order to form a more explicit judgment based on a more complete set of moral reasons and background information.

Although some instances of moral repugnance about a phenomenon may be deep-rooted, the context-sensitive interaction between feelings of revulsion, information, and moral reasons are dynamic. By discovering certain facts, or new moral reasons, or the fact that certain moral reasons are false, we may change the feeling of moral revulsion. There is a feedback loop between emotions, information, and reasons such that small shifts in information can change the contextual-frame and the character of repugnance.

The example of the sacred pool wader from above helps to demonstrate that the contextual-frame is a dynamic entity. Recall that the boy felt moral repugnance about the act of the stranger, wading in the sacred pool. His moral repugnance is based on framing information about what the sacred pool is and the moral rules about how to interact with it. Perhaps it was considered a most vile act to touch the water of the sacred pool in any way, and so the boy believed the moral rule that it is immoral for anyone to wade in the sacred pool. Upon seeing the stranger, he immediately feels moral repugnance at the stranger's actions. But then, upon inquiring, and learning the non-moral information that the stranger is the messiah returned, the boy revises his moral reason, which, in turn, eliminates his feeling of repugnance by removing both the moral and the revulsion aspects. The boy's moral reasons changed from 'wading in the sacred pool is morally wrong' to 'the messiah can wade in the sacred pool'. And, that the person wading in the pool is the messiah returned replaces the boy's previous background information that the wader was a stranger. The simple summary is that this shift in context, directed by the change of

certain information *corrects* our initial belief-based judgment that caused our feeling of revulsion.

Since repugnance that is based on false beliefs seems clearly mistaken in relation to moral reasons, the following definition will be used for mistaken moral repugnance in this discussion:

Mistaken moral repugnance occurs when person X feels moral revulsion about thing T, but careful analysis of the relevant information finds that X's moral revulsion about T is based on a false belief about T such that, if X held sufficient true, and no false, beliefs about T, then X would not feel moral revulsion about T.

So, moral repugnance is revealed to be mistaken moral repugnance when further information about T, changes the way X feels about T in one of two main ways.⁵ First, the morally charged aspect might disappear, while an amoral revulsion remains, such as when a child with no knowledge of modern medicine witnesses a surgery and then, upon inquiring, has it explained to him that the surgeons are helping and not torturing the patient. The child no longer finds the surgery morally problematic, but the graphic visual display still makes her feel sick to her stomach. Or second, the whole feeling of revulsion might dissipate, such as in the sacred pool wader example

⁵ Similarly, there is a way in which further information about a context, as opposed to a thing (e.g. new technology), can also change the way that someone feels about that thing. For example, imagine that a man wearing a crass shirt is presenting information on the benefits of human cloning and some people in the audience (who are unfamiliar with the possibility of human cloning) feel revulsion. Although the sickened audience members cannot quite put their finger on the reason for why they find human cloning morally repugnant, they are certainly feeling revulsion, and human cloning is both a moral issue and the likely cause of their feeling of revulsion. However, it is possible that the audience members are not emotionally affected by the idea of human cloning at all—they are not moved by any moral reasons or lack of moral reasons in relation to human cloning whatsoever. Instead, they have been unwittingly repulsed by the presenter's crass shirt, and have been mistakenly attributing that unjustified repugnance to the potentially morally relevant idea of human cloning. However, when the presenter stops talking, and turns off his presentation, the audience members take more notice of his shirt, and realize that it is the cause of their feelings of revulsion. The audience members gain more amoral information about the context, which leads them to realize that their judgment of moral repugnance was mistaken; they were simply experiencing sartorial revulsion.

above. When the boy learnt more about the situation, there was a shift in the contextual frame, and he lost his moral concerns and his feeling of revulsion at the same time.

In both of the examples above, each child felt moral repugnance, and came to the judgment that they were observing an immoral act. If these children had not sought out more information about what they observed, then they might still feel moral repugnance, and think that the things they saw were immoral. These examples help to highlight the danger in treating even widespread moral repugnance as being an argument in a moral debate. The information that shows the moral repugnance to be mistaken might not be obvious to most people, or it might be attainable, but only after some targeted research. Indeed, cautioning against Kass's wisdom of repugnance, Roache and Clarke (2009) have pointed out that using unjustified repugnance as an argument in a moral debate can stymie the discussion by making arguments without enough substance to object to. For (at least) this reason, repugnance should be investigated by those experiencing it to see if they can justify it, to make it "legitimate". Here, "justification" and "legitimacy" will be context-sensitive to someone's moral reasons, which are derived from a moral framework.⁶ A moral framework is systematic set of moral beliefs that can be reasoned through—e.g., consequentialism, deontology, and virtue ethics. This is not to be mistaken with a form of relativism. As we will see when discussing prediction markets and repugnance, moral reasons can shift between different moral frameworks but this does not mean that moral reasons are reducible to claims about subjective states or that moral reasons are a shell of moral emotions.

⁶ The justification need not appeal to a correspondence theory of moral truth. Our focus is the process of moral reasoning rather than the final outcome of the moral reasons latching on to some moral truth. However, we do focus on the correspondence between moral reasons and information about states of affairs.

For the purposes of this paper, legitimate moral repugnance will be defined as follows:

Legitimate moral repugnance occurs when person X feels morally charged revulsion R about thing T; and careful analysis using X's moral framework finds the set of moral reasons that correspond to R, which thus deems T immoral.

Legitimate moral repugnance designates a *process* of moral reasoning in which the feeling of revulsion begins the analytical moral search for reasons that will justify something being moral or immoral. If people experiencing moral repugnance can (reasonably) justify it with a moral framework then the repugnance should be considered legitimate and be given more weight in moral debates (assuming that the justification is methodologically careful). Such justifications should be given more weight again if they are intersubjective—i.e. generally compelling to other people. Justifications can be considered generally compelling if they only rely on background beliefs that are widespread.⁷ For example, if nearly everyone finds death pools repugnant, and people's self-reflection on their moral frameworks reveals that they find death pools immoral because of, e.g., the way that they disrespect weighty matters of life and death (moral reasons), then there would be widespread legitimate moral repugnance about death pools. In this case, there is a correspondence between the repugnance and the moral reasons that we use to analyse certain information (deontological reasons in this case). Alternatively, we could weigh consequentialist reasons against consequentialist reasons within a moral algebra. For example, given that there are few, if any, benefits from death pools that could not be sourced from

⁷ Here, we say that this is an 'indicator' of translation between frameworks because it may be that everyone just has one moral framework. It may also be that there is merely the coincidence of moral results without translation. These effects are not detrimental to the concept of legitimate moral repugnance, so we accept the indicator as being just that.

elsewhere, the widespread legitimate repugnance about death pools seems to be enough to outweigh the benefits, and thereby deem death pools immoral and impermissible.

If there is widespread legitimate moral repugnance about a new technology, then proponents of the new technology can try to argue in two ways. They might try to persuade all of the people experiencing legitimate moral repugnance that their moral framework is wrong and that the correct moral framework would not deem the new technology immoral. Or, they could try to show that other moral factors within the given moral framework outweigh the repugnance—a strategy that is much more likely to be effective if the other moral factors are part of the moral framework of the people experiencing legitimate moral repugnance.

But, the opaque nature of our moral intuitions means that we will not always be able to justify our moral feelings, such as moral repugnance, accurately. Indeed, we are not always aware of how the contextual frame, as well as any shifts in that frame, can make our attempted justifications inaccurate. For example, when we judge a situation to be morally repugnant, we may have reasons that correspond to the feeling of revulsion—we may be able to justify our moral repugnance by referencing our personal moral framework. However, it is possible that the moral reasons we use to justify our moral repugnance are not in fact related to the thing we feel the repugnance about.

Jonathan Haidt has organized many experiments that run participants through this process.⁸ In Haidt, Koller and Dias (1993), the researchers describe a scenario to participants (in an interview setting) in which a taboo is violated; the recently deceased family pet is eaten. Many participants are emphatic that eating the dead pet is morally

⁸ See Bjorklund, Haidt, and Murphy's (2000) as well as Haidt, Koller and Dias (1993) for this process of eliminating reasons but failing to eliminate the feeling of repugnance.

wrong, citing various plausible consequentialist and deontological principles. But, in response to the cited moral principles, the researchers modify the scenario so those principles no longer apply (e.g., there is no harm for the animal, no risk of the eaters getting food poisoning, no lack of respect for the animal is experienced, etc.). “Is it still morally wrong?”, the researcher would ask, after the participant’s potential justifications for their moral repugnance are shown not to be relevant. It turns out that, for many participants, even though every justification they could conceive of proved not to be relevant, they could not shake their moral repugnance—their feeling that eating the family pet was morally wrong.

The above example makes no conclusions about the objective, or even culturally relative, moral status of eating the family pet, but it does demonstrate that moral repugnance can remain even if no subjective justification is available for it. If moral repugnance were a slave to moral reason, then we should expect the repugnance to dissipate, or at least lose its moral charge. However, at least in some instances, the feeling of morally charged revulsion sticks. This might indicate that there are vague, ambiguous, and/or indeterminate moral reasons behind the repugnance (moral indeterminacy will be addressed in more detail in Section 4). More simply, there may be currently inaccessible moral reasons that correspond to the moral repugnance, but it is also possible that there are not. So, when there is a discrepancy between the feeling of moral repugnance and the moral reasons that we search for and cannot find, it is hard to know whether wisdom resides in that repugnance. As a result, sometimes one is left trying to reconcile a strong moral feeling with what rationally appears to be an *amoral* state of affairs, i.e. one without corresponding moral reasons. For the purposes of this paper, this phenomenon will be referred to as dumbfounded repugnance.

Dumbfounded moral repugnance occurs when person X feels moral revulsion about thing T, but careful analysis using X's moral framework finds T to be amoral (or has nothing morally explicit to say about T).

Kass, and many others with faith in the wisdom of repugnance, probably believe that their repugnance about an act carries more weight in the argument than their inability to find a fitting reason to morally condone or condemn the act. But, this may be because they believe that the repugnance is an indicator for some deeply seated moral reasons. Indeed, without a strong argument for why an act is morally praiseworthy, or why the repugnance is misguided, anyone siding with the wisdom of repugnance seems, epistemically speaking, reasonably entitled to do so. This means that the burden of proof is on anti-repugnant moral reasoning. The reasoning here is similar to choosing a scientific method that produces false positives. With repugnance as the moral status quo, we may get cases where that repugnance is wrong (false positive), but we will most likely get more cases in which it corresponds to the proper set of moral reasons. The contrasting strategy is to be wary of moral repugnance as an indicator of moral reasons. This has the consequence of ignoring many cases in which repugnance has moral justification that is not accessible at this time.

One benefit of allowing dumbfounded moral repugnance to play a role in moral debate about new technologies is that it encourages proponents of the new technology to make a stronger and more explicit moral case for why the new technology should be permitted, thereby decreasing the chances of permitting a potentially dangerous new technology. Indeed, like an evolved precautionary principle, our ingrained propensity to be suspicious of (and occasionally undervalue) the unknown is likely to have been adaptive for precisely the same reason—it

encourages us to cautiously investigate new and potentially beneficial things. It is a call for more moral evidence, and until then it is to serve as prima facie evidence against permitting the new technology. The reasoning here is pragmatic. If we are concerned with preventing negative effects, then we will halt the technology until we see more evidence. However, as Kass rightly argues, the onus now falls on the proponents of the new technology to argue why the widespread moral repugnance is mistaken.

When people experiencing repugnance about a new technology are trying to justify it with their moral frameworks, and when they are considering arguments from proponents of the new technology, they might find that the new technology receives a mixed verdict from their moral framework. That is to say that they can see reasons why the new technology is immoral (why they might find it repugnant), but also reasons for why it would be morally good to permit it. Three kinds of mixed verdicts deserve closer attention, starting with indeterminate moral repugnance.

Indeterminate moral repugnance occurs when person X feels moral revulsion about thing T, but careful analysis using X's moral framework finds T to correspond to both immoral and morally good reasons, and, all thing considered, X's moral framework finds T to be morally indeterminate, or it cannot not pass a clear judgement on X.

Given that different kinds of moral reasons are not always easy for people to accurately weigh up, many people experiencing repugnance about a new technology may find themselves with indeterminate moral repugnance. For example, a cancer researcher who believes that the sanctity of life begins at conception might feel that there are incommensurable moral reasons both for and against permitting research on embryonic stem cells. For this researcher, we could say that their moral repugnance is justified (by the sanctity of life concern), even though they might not

have a clear overall moral judgment about the permissibility of embryonic stem cell research. At first it seems that, widespread indeterminate moral repugnance about a new technology should be treated like widespread dumbfounded moral repugnance—it should play the role of prima facie evidence against permitting the new technology and act as a call for further investigation on the part of the proponents of the new technology. However, as discussed in Section 4, because some scenarios give us incommensurable moral reasons, our only option is to either change the contextual frame to make it a case of non-indeterminate repugnance or to accept that no new information will make the moral reasons commensurable.

Another type of repugnance that receives a mixed verdict from a moral framework is mitigated legitimate moral repugnance.

Mitigated legitimate moral repugnance occurs when person X feels moral revulsion about thing T, and careful analysis using X's moral framework finds T to correspond to both immoral and morally good reasons, and, all thing considered, X's moral framework finds T to be immoral.

The cancer researcher from above might also decide that, although embryonic stem cell research could help relieve suffering and save lives, the sanctity of life and the moral rule not to kill are categorical moral rules that allow for no exception. As such, the researcher's overall moral judgement would be that embryonic stem cell research is immoral. In most respects, widespread mitigated legitimate moral repugnance will play the same role as widespread legitimate moral repugnance—it should be given weight in moral debates about new technologies (and more weight than is given to dumbfounded moral repugnance).

The final type of repugnance that receives a mixed verdict from a moral framework is outweighed legitimate moral repugnance.

Outweighed legitimate moral repugnance occurs when person X feels moral revulsion about thing T, and careful analysis using X's moral framework finds T to correspond to both immoral and morally good reasons, and, all things considered, X's moral framework finds T to be morally good.

Perhaps in a nearby possible world, the cancer researcher has a family history of several severe cancers, and has children with a high chance of contracting cancer during their lives. The researcher still believes that harvesting embryos is morally repugnant because it is killing a human life (so their repugnance is legitimate), but this time his moral framework allows for comparisons between the good of potentially curing cancer and the bad of killing. As a result, the researcher thinks that, all things considered, it should be morally permissible to conduct embryonic stem cell research. Notice here that the contextual shift of the moral framework determines the final moral outcome (moral judgment).

It should be noted that people's moral frameworks change over time, including as a consequence of moral argument or example. Any change in moral framework during the moral debate about a new technology should result in the new technology being reassessed using the varieties of moral repugnance defined above.

Although there are other varieties of repugnance,⁹ the above varieties provide a sufficient framework for assessing the dynamics of repugnance that emerge

⁹ *Conflicted moral repugnance* occurs when person X feels moral revulsion about thing T, but careful analysis using X's moral framework finds T to be (only) morally good.

Dominant moral repugnance occurs when person X feels moral revulsion about thing T, but careful analysis using X's moral framework finds T to be morally good, and X revises their moral framework (ensuring that T is deemed immoral) to accommodate their repugnance.

in the prediction market scenarios about betting on death, as well as most of the moral results generated from tracking how repugnance shifts in contextual frames.

In the next section, the dynamics of repugnance are analyzed within prediction market scenarios. In Section 4, moral results are discussed.

3. Prediction markets and contextual frames

3.1 Prediction markets on epidemics and infectious diseases.

The Iowa Electronic Health Markets (IEHMs; <http://iehm.uiowa.edu/iehm/main/>) is a website that allows anyone to set up their own health-related prediction market or to bet on the existing health-related predictions. Many of the markets are designed to promote betting on epidemics—on the spread and deadliness of viruses and other infectious diseases. For example, traders can bet on “What will be the level of 2009 H1N1 [(also known as Swine Flu)] influenza mortality rate in the U.S. by the end of July 31, 2009?”¹⁰ Since most of the markets run at the IEHMs are for “funny money”¹¹ (status points) instead of real money, it seems fair to say that the IEHM encourages traders to bet on matters of life and death for the fun of it. Indeed, traders will be playing a betting game in which the real-life occurrence of mass death caused by a virus outbreak could help them win.

No doubt many people find the IEHMs repugnant, and, if they tried to justify those feelings using information about IEHMs, along with their moral frameworks, they would probably suggest that the IEHMs are crossing a moral boundary by

Moral repugnance deficit occurs when person X feels no revulsion about thing T, but careful analysis using X’s moral framework finds T to be immoral.

¹⁰ See http://iehm.uiowa.edu/iehm/markets/published/swine_mortality_0908.

¹¹ See <http://iehm.uiowa.edu/iehm/content/faq/>.

failing to show appropriate respect for the lives and deaths of those who have been directly or indirectly affected by deadly diseases.

However, by shifting the contextual frame, using more detailed information about IEHMs, moral conclusions can shift to IEHMs being highly respectful of matters of life and death. The IEHMs website is a not-for-profit initiative with the goals of improving knowledge about prediction markets and especially making health-related predictions that help medical professionals better protect the health of people around the world (such as by providing information that enables the timely production and distribution of appropriate vaccines). Furthermore, initial reports (e.g., Polgreen, Nelson, & Neumann 2007; see also <http://iehm.uiowa.edu/iehm/main/>) suggest that prediction markets might be a very effective way to quickly identify the danger posed by new infectious diseases, and thereby enable a faster and more accurate response that could result in thousands of lives being saved. Since the IEHMs have the intention and likely effect of saving lives, it seems as though they are treating the matters of life and death, to which they are relevant, with a high level of respect.

The contrast here is between two contextual frames. One frame is that individuals might bet frivolously, with a lack of appropriate respect for those whose lives or deaths have been affected by infectious disease. The other frame is that the design of the IEHMs has the intention of providing accurate predictive information about epidemiological phenomena, and has proved both effective and efficient in providing such information. This market is designed with deontological as well as consequentialist considerations. The benefits it provides go far beyond the frivolous fun of betting for no good reason. Here, the two contextual frames give us two sets of information and, thereby, two sets of moral reasons.

It is important to address that these shifts in contextual frames are determined by certain information about not only the IEHMs but also about human psychology. For example, what is the likelihood of people with frivolous intentions trading on the IEHMs? Can we conclude that it is low considering the minimal pay-offs available compared to trading on the more popularist predictions available on real-money for-profit prediction markets? Such statistical information will contribute to information about the IEHMs' disrespect for matters of life and death, and thus will contribute to the contextual frame for the likely varieties of moral repugnance about the IEHMs.

3.2. Prediction markets on medical diagnoses

CrowdMed (www.crowdmed.com) is a new online medical diagnosis prediction market that concerns some doctors (Hall 2013). People with undiagnosed and highly worrying illnesses pay CrowdMed \$200 to reveal their personal medical histories to CrowdMed's M.D.s ("medical detectives"). Anyone can join CrowdMed's ranks of medical detectives in order to suggest diagnoses and bet on their own or other M.D.s' suggestions. Upon hearing a brief description of what CrowdMed does, many will find it repugnant—the idea of strangers with no medical expertise suggesting diagnoses of the most painful and deadly illnesses to vulnerable people is likely to elicit feelings of moral revulsion in many, if not most, people. Adding to the intensity of this feeling is that the suggestions produced by the medical detectives are just that, suggestions, and might be mistaken as medical advice.

Let's develop the specific contextual frames. If people attempted to reconcile these initial bits of information about CrowdMed with their moral frameworks, it is likely that they would come up with at least one of the two following justifications. First, they would likely claim that CrowdMed is repugnant because it recklessly

encourages untrained traders to suggest high-stakes medical diagnoses to vulnerable people in a way that puts these people at risk of further emotional and physical damage (e.g., by not pursuing further professional medical advice). This would clearly violate the oath to do no harm taken by many doctors around the world. Second, they would likely claim that CrowdMed is repugnant because whether or not someone has a deadly illness is a weighty life and death issue that should be treated with the utmost respect, not gambled on frivolously by curious strangers. Reasons can be grouped into consequentialist reasons, focusing on negative outcomes (e.g. increasing harm of patients), as well as deontological reasons, focusing on moral principles and virtues (e.g. manipulating someone's autonomous decisions). Examining information from a different perspective about how CrowdMed works, however, reveals a possible shift in contextual frame.

CrowdMed does rely on the verdicts of unscreened “medical detectives”, but apparently not in a reckless way or a way that takes advantage of the vulnerable. The end result of medical detectives' trading on CrowdMed is a short list of possible diagnoses (often of rare diseases) based on the medical detectives' trading behaviour in the market. CrowdMed is explicit that these potential diagnoses should then be taken to a medical professional to discuss them and possibly arrange for the relevant diagnostic tests. So, CrowdMed does not steer its “patients” away from professional medical care. This means that autonomy is respected. In fact, this can be seen as a way to give individuals more options. Indeed, it seems likely that people will only use CrowdMed *after* they come to realise that the medical professionals in their area cannot properly diagnose their problem. Furthermore, CrowdMed then immediately directs their patients back to professional medical experts. With this information, we can make a judgment about the intentions behind the design of CrowdMed. The

intended purpose of the design seems to be to respect individuals and provide possibilities.

Now for the consequential analysis. Although CrowdMed charge \$200 to its patients, this does not mean that they are preying on the vulnerable. According to their website, CrowdMed refunds the \$200 if none of the suggested diagnoses are correct, making patients' use of CrowdMed financially favorable. Furthermore, if one of the suggested diagnoses is correct, then CrowdMed has likely saved the patient thousands of dollars in further medical examinations (not to mention the health benefits, since CrowdMed claims to be highly accurate; www.crowdmed.com/faq).

Assuming CrowdMed is accurate, and considering the description of CrowdMed above, then any repugnance felt about CrowdMed because of a belief that it is reckless, or takes advantage of the vulnerable, seems mistaken. Indeed, the way in which CrowdMed's prediction market on medical diagnoses seems to empower the vulnerable, rather than take advantage of them, makes it appear to be the opposite of reckless. Given this information within this contextual frame, if people feel that CrowdMed is repugnant *solely* because it recklessly endangers vulnerable people, then their repugnance appears to be mistaken.

But even people who believe that CrowdMed has good intentions in its design and that CrowdMed's prediction market on medical diagnoses seems to be beneficial to most patients, might still find the idea of betting on painful and deadly diagnoses repugnant because it does not appropriately respect that this is potentially a matter of life and death for the patient. Here, we are yet again shifting contextual frames. We are moving from the frame that takes into account the intentions behind the design, and also from the frame that takes into account the beneficial consequences. The current frame takes into account information from the patient's individually-located perspective to juxtapose it with the fact that CrowdMed's operation is

powered by betting. How can we respect this as a matter of life or death if betting is involved? For example, CrowdMed's medical detectives could be betting for amusement or the glory of being the most accurate predictor—motives that are disrespectful considering that people's lives are at stake. So, it seems that the betting-on-matters-of-life-and-death aspect of CrowdMed is legitimately repugnant because it does not appropriately respect the lives and potential deaths involved.

It might change our contextual frame again if we learn the information that CrowdMed's medical detectives win money for correctly predicting diagnoses, but also that the money goes to a charity, which allocates it to a real patient (of the medical detectives' choice) who needs financial support to treat their illness. But this bit of information changes the feeling of repugnance only if our moral framework has certain exceptions. For example, if betting on human lives is intrinsically immoral, does it change the nature of the bet to know that good consequences will result for human lives (e.g., in the form of more donations for medical treatment)? This brings us to an interesting point: Maybe we can prevent shifts in contextual frames. Why can't we just amalgamate all of the information in the frames to get complete information? The answer is explored in the next section.

4. Results of applying contextual frames to prediction markets

There are several important results that arise from applying the dynamics of moral repugnance to prediction markets. First, a prediction market can elicit several conflicting types of moral repugnance at the same time. Second, moral indeterminacy can arise in two different ways when judging prediction markets. Finally, some prediction markets can generate a moral endogeneity problem, a disruptive feedback loop between a given prediction market and the morally relevant outcome it predicts.

4.1 Levels of information and multiple moral frameworks

Prediction markets offer a vast space to observe the dynamics of moral repugnance. We have seen how moral repugnance shifts from mistaken moral repugnance to legitimate, back to mistaken, depending on the contextual frame. The first question is, why can't we prevent the shift, stabilize the contextual frame? The simple factor that prompts the contextual shift is the range of perspectives, and this factor cannot be eliminated. As we saw in the previous section, both the IEHMs and CrowdMed appear to have respectful intention in their design. That is, they seem to be designed to respect autonomy (e.g., by providing predictive information and options), as well as generate results that will help individuals. However, from the perspective of an individual that is involved in the life threatening event or condition, the juxtaposition between someone betting on your personal situation and the way that you experience that situation does not necessarily translate to the so-called respectful intention in the design of the prediction market. This is not merely unjustified offense. Rather, it points to the minimal level of treatment many of us would expect from others when care about our lives is involved; we expect our living or dying not to be part of a betting game, even when that game has benefits. Perhaps, we expect to be treated like humans rather than pawns. There are numerous other perspectives available. The careless better, for example, or the careful better, and so on. Importantly, these perspectives are not compressible into one simple perspective. There is also no invariant perspective that is foundational to every perspective. As such, we don't get a view from nowhere about these prediction markets. Rather, the prediction market-user relation produces layers of perspective, which are mostly incommensurable. Each layer offers a different set of information, each of which can determine different moral reasons. The interesting result is that while these

perspectives exist simultaneously within the prediction market-user relation, the user may only be able to adopt one at a time because of the incommensurable information.

Even if we were to choose one perspective over the others, this would still not help us reach a single contextual frame. This is because background information is not the only multi-layered feature of a contextual frame. There are multiple moral frameworks users can adopt. A person who subscribes to a different moral theory to another person may reach a contradictory moral judgment based on exactly the same non-moral information. As discussed in Section 3, we can take the consequentialist approach, which focuses on the output of successful prediction. And, we can take the deontological approach, which focuses on concepts like autonomy, no matter what the consequences are. These perspectives are usually interpreted as offering incommensurable frameworks for analyzing moral phenomena. Sometimes the results overlap, but the *processes* of moral reasoning do not. We could choose one framework over the other, but then we face an empirical ethical problem: have we represented all of the relevant moral parameters in the scenario? In the prediction market scenarios in Section 3, the shift between moral frameworks occurred by ignoring certain information. But this is rarely an option when trying to generate a complete representation of a moral scenario. For example, we might want to take into account aspects of human choice as well as the consequences (for individuals and populations). So, like the information perspectives, the macro-level moral framework perspectives are also numerous. Furthermore, these different moral perspectives will likely exist simultaneously in the prediction market-user relation, meaning that they are all accessible as possible moral frameworks to evaluate moral parameters of the prediction market. Perhaps, because of the incommensurable moral principles in each moral framework, a user may only be able to adopt one at a

time, but individuals can change their moral frameworks, and multiple individuals can occupy the prediction market-user relation at any one time.

4.2 Indeterminacy

In Section 3, contextual shifts were discussed in relation to mistaken moral repugnance and legitimate moral repugnance. Outweighed legitimate moral repugnance and mitigated legitimate moral repugnance work in the same way as legitimate moral repugnance, with one small addition; there is a conflict in moral reasons. But the conflict is resolved by one's moral framework. For example, sometimes such conflicts are resolved because certain factors are given a larger value in the moral algebra—e.g., successfully predicting the outbreak of a disease in a population is more important than preventing selfishness and gain in betting on the outcome of the outbreak. Mistaken, legitimate, mitigated, and outweighed moral repugnance all share the fact that there is a determinate outcome in which something is decided to be moral or immoral. However, there is another type of repugnance that produces a puzzling scenario.

Indeterminate moral repugnance occurs when careful analysis of conflicting reasons yields moral indeterminacy. This is different from dumbfounded moral repugnance, in which careful analysis determines something to be amoral. For the purposes of this paper, the nature of indeterminacy has to do with a relation between information, moral reasons, and revulsion.¹² Specifically, indeterminate repugnance is a process in which some conflicting combination of feelings, information, and moral frameworks interact in order to produce an output that has no determinate moral value. An interesting feature of indeterminate moral repugnance is that when

¹² Metaphysical issues about indeterminacy will not be discussed here. See Schafer-Landau (1995) for a helpful discussion on alethic and comparative indeterminacy.

we go through the process of careful analysis, we gain more relevant information, but unlike the other varieties of moral repugnance discussed here, the resulting moral judgment becomes less determinate.

One way to manufacture indeterminate moral repugnance is simply to keep the information limited. This way, the judgment cannot be passed on. We can keep any amount of moral information unknown—e.g., the intentions of agents, the consequences for individuals and populations, and the success of the market in generating information—thus making a concrete moral judgment difficult to reach until that the missing information is known. This may seem extreme, but there are many unknowns when judging a prediction market. For example, what is the exact success rate of CrowdMed and IEHMs' predictions in a particular subdomain? What are the majority of these prediction market users really like? Are they careless gamblers, thoughtful do-gooders, or bored individuals? In such cases, indeterminate moral repugnance takes form because of unknowns. But, there are other cases in which there is adequate information and indeterminate repugnance still takes shape.

Indeterminate moral repugnance can also occur when whether or not to implement the prediction market takes a form similar to many traditional moral dilemmas. For example, imagine a prediction market in which traders are encouraged to bet on matters of life or death, but implementing the market would considerably bolster conservation efforts. In cases like this, we might have all of the information we need; but some people will still find the prediction market morally indeterminate because they find the moral reasons on either side incommensurable. Someone might value conservation very highly, while also finding betting on life and death morally repugnant, and believing conservation and appropriate respect for matters of life and death to be impossible to trade off against one another.

4.3 The moral endogeneity problem

Claims that prediction markets would not work are mainly based on the endogeneity problem faced by some kinds of markets (e.g., Stiglitz 2003, Richey 2005). Generally, ‘endogeneity’ refers to a causal loop between two variables. We apply this to endogeneity in the stock market to demonstrate the general form of the endogeneity problem. Then, we use this form to present a new moral endogeneity problem.

Market prices are often used to ascertain the actual value of something because it is thought that the actual value of the thing in question sets the market price (in relation to extant supply and demand). For instance, the share price of a firm in a stock market should be set by the real-world value of the firm (in relation to extant supply and demand). The endogeneity problem occurs when the market price affects the thing in question, and often its actual value (Birchler & Facchinetti 2007). This kind of endogeneity can be problematic because it warps the incentives for trading in the market. Indeed, endogeneity can sometimes warp the incentives to such an extent that traders will avoid the market entirely. But, this is true only when the betting of a trader has an effect on some prediction being accurate.

Imagine a trader betting money that their best friend will discover a treasure, and the more that trader knows about the accurate spot of the treasure, the less likely it becomes for that trader’s best friend to discover it. The reason why it becomes less likely is because bets are public entities, so any information provided can shift the final outcome, given that the final outcome can be affected by the betting activity. Richey (2005) argues that if well-informed traders were very confident that a particular prediction was accurate, then the endogeneity of the market might dissuade them from buying shares in the prediction because buying shares in the prediction increases the price, and, depending on the prediction in question,

increasing the price is likely to change someone's real world behaviour in a way that increase the chances that the prediction will *not* come true.

We can apply the endogeneity problem to the dynamics of repugnance. Prediction markets face a moral endogeneity problem, where the information used to make a bet can determine the likelihood of the event, which determines the morality of using information to make the bet. This only works for prediction markets where the likelihood of the event can be affected by the bet. For example, suppose that we bet on how likely it is for a person to get sick while in a high-risk flu area. Further, suppose that we choose these individuals by real names and make this information publically accessible so that the individuals can see that they are being betted on. Let's imagine that Sal Harrington, who doesn't have the greatest immune response during the wintertime, is one of these individuals. Tray Der knows this about Sal and bets that he will likely become hospitalized due to the flu. Sal becomes aware of the bet on his health at the same time the public does. Determining the likelihood of Sal getting sick just by betting on Sal can be simple: People attempt to get Sal sick.¹³ Or maybe Sal will develop the nocebo effect, where Sal's negative expectation produces negative physiological results in the form of stress hormones and a weakened immune response.¹⁴ Either way, the likelihood of the event is influenced by the bet. But even more interesting is the fact that the morality of using information to make the bet is influenced by the likelihood of the event, which is influenced by the information used to make the bet. So, because Tray Der knows something about Sal, which makes Sal more likely to get sick, it becomes more immoral for Tray to use this information. This illustrates the moral endogeneity problem produced by prediction

¹³ Imagine something like the most corrupt sports betting—where bets sometimes rely on mechanisms for injuring the players during preseason just a little in order to ensure that the team suffers just enough for the bet to be won.

¹⁴ See June et al. (1997).

markets. Such a problem creates a feedback loop between information, likelihood of events, and morality, which can only be broken if the link between information and likelihood is broken. Additionally, we may not know if a given prediction market has an endogeneity problem until the prediction market is up and running.

5. Concluding Remarks

We have explored the varieties and dynamics of moral repugnance by characterizing the relationship between feelings of moral revulsion, information about states of affairs, and moral reasons within moral frameworks. By presenting a new model and classificatory scheme for repugnance in terms of contextual frames, we illustrated the dynamic shifts in repugnance, and applied these shifts to prediction markets, including the real markets, CrowdMed and the Iowa Electronic Health Markets. Three interesting results were discussed about analyzing the dynamics of moral repugnance within prediction markets. First, a prediction market can elicit several conflicting types of moral repugnance at the same time. Second, moral indeterminacy can arise in two different ways when judging prediction markets. Finally, some prediction markets can generate a moral endogeneity problem, a disruptive feedback loop between a given prediction market and the morally relevant outcome it predicts.

References

Birchler, U. W., & M. Facchinetti (2007). Self-destroying prophecies? The endogeneity pitfall in using market signals as triggers for prompt corrective action. *BIS Research Task Force Paper, March*. Available from: <http://www.bis.org/bcbs/events/rtfo7birchler.pdf>, accessed on 3 July 2013.

- Bjorklund, F., Haidt, J., & Murphy, S. (2000). Moral dumbfounding: When intuition finds no reason. *Lund Psychological Reports*, 2(1), 29.
- Breitbart, A. & Ebner, M. C. (2004). *Hollywood Interrupted: Insanity Chic in Babylon: The Case against Celebrity*. New York: John Wiley & Sons.
- Hall, B. (2013). CrowdMed Wants To Crowdfund Your Medical Care To Strangers, Readwrite, 18 May 2013. Available from: <http://readwrite.com/2013/05/08/crowdmed-wants-to-crowdfund-your-medical-care#awesm=~obOzozBrExVkNZ>, accessed on 3 July 2013.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological review*, 108(4): 814.
- Kass, L. (1998). The wisdom of repugnance: Why we should ban the cloning of humans. *Valparaiso University Law Review*, 32(2): 679–705.
- Kimberly, M. (2002). Reevaluating repugnance: a critical analysis of Leon Kass' writings on genetic reproductive technologies. *Princeton Journal of Bioethics*, 5: 8-24.
- Niemela, J. (2011). What puts the 'yuck' in the yuck factor? *Bioethics*, 25(5): 267-79.
- Polgreen, P. M., Nelson, F. D., Neumann, G. R., & Weinstein, R. A. (2007). Use of prediction markets to forecast infectious disease activity. *Clinical Infectious Diseases*, 44(2), 272-279.
- Richey, M. (2005). Thoughts on the Theory and Practice of Speculative Markets qua Event Predictors. *Essays in Philosophy*, 6(1): Article 26, no page.
- Roache, R. & Clarke, S. (2009). Bioconservatism, bioliberalism, and the wisdom of reflecting on repugnance. *Monash Bioethics Review*, 28(1): 4.1-4.21.
- Stiglitz, J. (2003). Terrorism: There's no futures in it. *Los Angeles Times*, 31 July 2003: B13.

Turner, L. (2004). Is repugnance wise? Visceral responses to biotechnology. *Nature Biotechnology*, 22(3): 269-270.

Weijers, D. (2013a). Technology: Prediction markets. *Observatory for Responsible Research and Innovation in ICT*. Available from: <http://torrii.responsible-innovation.org.uk/resource-detail/1181>, accessed on 27 August 2013.

Weijers, D. (2013b). Intuitive Biases in Judgments about Thought Experiments: The Experience Machine Revisited. *Philosophical Writings*, 41(1): 17-31.

Weijers, D. & Richardson, J. (2014). A moral analysis of effective prediction markets on terrorism. *International Journal of Technoethics*, 5(1): 28-43.