

## **Prediction Markets as an Alternative to One More Spy**

Dan Weijers  
California State University, Sacramento

### *Introduction*

Real-world policy decisions involve trade-offs. Sometimes the trade-offs involve both the efficacy and morality of potential policies. In this chapter, the morality and likely efficacy of hiring one more spy to help anti-terrorist intelligence gathering efforts is compared to the morality and likely efficacy of implementing a prediction market on terrorism. Prediction markets on terrorism allow registered traders to buy and sell shares in predictions about terrorism-related real-world events. The comparison at the heart of this chapter is based on the assumption that it would cost about \$5 million to bankroll the prediction market project, or to establish another spy, including equipment and head office support, for 15 years. The comparison reveals that implementing a prediction market on terrorism is likely to be more efficacious and less morally problematic than hiring one more spy.

The chapter is organized as follows: The first section, introduces prediction markets on terrorism, and offers some reasons to think that they might be effective. The second section addresses objections to the effectiveness of predictions markets, including the concern that participants in the market will not have the requisite knowledge to enable the market to produce reliable predictions, and the worry that the market will produce bogus predictions due to manipulative trading. The third section considers two moral objections to prediction markets on terrorism, both related to the use of the markets as a form of insurance. The fourth section introduces two further moral objections to prediction markets on terrorism: that the markets will (but should not) reward informants, and that the markets involve the repugnant act of betting on death. The fifth section completes the comparison of implementing a prediction market on terrorism and funding one more spy, concluding that implementing a prediction market is the preferable policy.

### *Prediction Markets On Terrorism*

Prediction markets allow registered traders to buy and sell shares in predictions about real-world events (Weijers 2013, para. 1). For prediction markets on terrorism, those real world events are always related to terrorism, and will likely include predictions about terrorist attacks (Hanson 2006a, p. 264–265). For example, shares in the prediction: ‘There will be a fatal terrorist attack on the US Embassy in Egypt during January 2015’, might be for sale at around \$1 in a market that pays out \$10 for each share in a prediction that turns out to be true. When acting rationally, traders with the belief that there is a 20% chance of there being a fatal terrorist attack on the US Embassy in Egypt during January 2015 will purchase any share offers that are cheaper than \$2 (20% of the \$10 pay-out). At any given moment, the price of shares in a prediction can be understood as the market’s prediction of the likelihood of the predicted event actually happening.

While it may seem counterintuitive that the predictions of (mostly) non-experts could be aggregated into something useful, prediction markets have been demonstrated to do exactly that in many domains (Surowiecki 2004, p. 77–83). Moreover, Wolfers and Zitzewitz (2004, p. 1) point out that prediction markets have produced better-calibrated (i.e., more accurate) predictions than ‘moderately sophisticated benchmarks’ in several domains, including

This is the last author version. The official version is forthcoming in Jai Galliot & Warren Reed (eds.), *Ethics and the Future of Spying: Technology, Intelligence Collection and National Security*, Routledge.

domestic politics, sales of consumer goods and box-office takings. Green, Armstrong and Graefe (2007), have also demonstrated that prediction markets can perform comparably to or better than Delphi methods, which pool the views of relevant experts.

Weijers (2013, para. 5) explains the remarkable success of prediction markets like this: ‘Most of us have some genuine information about the topics in question as well as some randomly incorrect information. When all of this information is aggregated in the right way, the random incorrectness tends to cancel out (because it’s random) and the genuine information accumulates into the basis of an accurate prediction’. Furthermore, in prediction markets, the incentives are in the right place. Participants are rewarded for their accuracy, and participants with more knowledge can trade many times. In these ways, prediction markets are quite unlike polling, which doesn’t reward accurate information and gives everyone an equal ‘vote’ regardless of their level of knowledge. It is the success of prediction markets in a diverse range of domains, and the broad applicability of their underlying logic, that has led several authors to suggest using prediction markets to gain anti-terrorist intelligence (e.g., Hanson 2006a; Looney 2004; Surowiecki 2004, p. 77–83; Yeh 2006; Weijers & Richardson 2014a, 2014b).

The general justification for setting up a prediction market on terrorism is that it is likely to produce some knowledge relevant to counter-terrorism efforts that would not have been gathered through existing intelligence-gathering means, at least not in as timely a manner. This was one of the many goals of the prediction market, Policy Analysis Market (PAM), being developed by Richard Hanson and his colleagues before the project was dramatically halted due to government intervention (Hanson 2006a; 2006b; 2007). The aspect of PAM focussing on gathering anti-terrorism intelligence was the focus of concerted criticism from many people. Nobel laureate economist Joseph Stiglitz criticised the project on both moral and efficacy grounds, commenting that it represented ‘market fundamentalism descending to a new level of absurdity’ (Stiglitz 2003, para. 8). His criticisms are notable because he offers such a broad range of them, and because he is an expert on prediction markets. Since the full range of his criticisms has not yet been adequately addressed, the rest of this chapter deals with them all.

### *Prediction Markets On Terrorism Can’t Generate Useful Information*

#### Participants won’t have new knowledge

In his previous work on prediction markets—which he refers to as futures markets—with Sanford Grossman (Grossman & Stiglitz 1976; 1980), Stiglitz identified the importance of traders bringing relevant information to prediction markets: ‘We focused on competitive markets in which we assumed participants had some relevant information. For instance, a farmer knows something about his own crop, so if he participates in a futures market he will bring his knowledge to bear on that market. Voters who participate in a futures market also bring relevant information – whom they and their friends are voting for – and that is why futures markets may predict presidential elections reasonably well’ (Stiglitz 2003, para. 5). However, Stiglitz expressed doubt that a prediction market on terrorism would attract traders with this kind of relevant knowledge, since he found it hard to believe that ‘there is widespread information about terrorist activity not currently being either captured or appropriately analyzed by the “experts” in the FBI and the CIA’, and perhaps harder still to believe that traders in a prediction market on terrorism would have that information (Stiglitz 2003, para. 10).

Whether prediction markets on terrorism could attract traders with relevant information is an empirical question that seems unlikely to be completely answered without a real-world test. Nevertheless, there seems to be good reason to think that the right kind of prediction market on terrorism could indeed attract traders with relevant information. Consider an anonymous prediction market on terrorism that was actively promoted to a wide range of potential participants from around the world (except, of course, for allied intelligence agency operatives, since then the market would not be one among many distinct sources of information, and would possibly provide too much information to terrorists). Several groups of participants of different sizes and with different kinds of information and incentives to invest could be expected, as outlined in Table 1.

<b>Group</b>	<b>Size</b>	<b>Relevant Direct Private Information</b>	<b>Relevant Ambient Private Information</b>	<b>Location</b>	<b>Incentive</b>	<b>Power over event</b>
Ordinary civilians	large	none-low	low	mainly in allied countries	mainly interest/fun	none
Connected civilians	small	low-moderate	low-high	mainly in or nearby suspected 'terrorist host' countries	mainly financial or informative	none-low
Allied security force members	small	none-low	moderate	diverse	mainly financial or informative	none-low
Allied intelligence agents	tiny	moderate	high	diverse	mainly financial	low-moderate
Terrorists	tiny	very high	very high	mainly in 'terrorist host' countries	financial or mis-informational	moderate-high

Table 1: The characteristics of the potential user groups of a prediction market on terrorism.

Contrary to what several writers on this topic have suggested (e.g., Posner 2004, p.175–176; Stiglitz 2003), even some ordinary citizens with little chance of ever crossing paths with a terrorist or terrorist supporter might have useful private information about the likelihood of particular types of terrorist attack. For example, as Hanson (2006a, p. 265) points out, these completely unconnected ordinary citizens might still have useful private information about the relative vulnerability of certain potential terrorist targets. Furthermore, if these unconnected ordinary citizens worked as security guards, or as janitors, or in certain other roles, then they might have information about the vulnerability of potential targets that intelligence agencies do not.

Regardless of whether ordinary citizens will bring relevant knowledge to the market, the most likely traders to bring relevant information to the market are connected civilians. Connected civilians are people who come into contact with terrorists and their supporters, but do not support terrorism themselves. Posner (2004, p. 175–176) denies the possibility that this group

This is the last author version. The official version is forthcoming in Jai Galliot & Warren Reed (eds.), *Ethics and the Future of Spying: Technology, Intelligence Collection and National Security*, Routledge.

exists, but surely not everyone sharing a geographical location with terrorists approves of their methods. The contact might be direct (connected civilians could know someone and suspect or know that person to be a terrorist) or indirect (they might have overheard snippets of conversation because they live or work in places that terrorists and their supporters frequent). Connected civilians that live in known 'terrorist host' countries might hold both direct private information and a high level of ambient private information related to terrorist activities, particularly if they live in an area that at least partially supports the terrorist activity in question (since terrorist activity might be discussed more openly there). Ambient private information is the aggregate of the indirect information held by an individual, such as information from television news, opinion polls, and hearsay. Furthermore, different individual connected citizens are likely to hold different private information; one might hear about a timeframe, another about a location and so on.

Why wouldn't these connected citizens simply alert intelligence agencies instead of using a prediction market to get their message out? There are lots of potential reasons. Some connected citizens would have financial motives for participating in the market. They would have little reason to approach the authorities. A payment might be offered for further information, but the informer may be investigated, constantly followed, or even tortured. Even if the informant had anti-terroristic motives, getting in touch with intelligence agencies is risky for the same reasons; raising the heat could result in reprisals from terrorist supporters, and the act of informing could result in being tortured for further information. The informant might also worry about not being taken seriously, or being ignored. It might also be difficult or impossible for some informants to contact the right agencies. There might not be the right kind of embassy nearby, and they might not have the resources to find the number for and make the call to the Pentagon. On the whole, it seems much more likely that the majority of connected civilians who desire to relay anti-terrorist intelligence, and especially those in 'terrorist host' countries, would prefer to use an anonymous prediction market than to risk direct personal contact.

So, if a prediction market on terrorism had enough connected civilians, and it played its aggregating role effectively, a lot of relevant information that was previously unavailable to intelligence agencies might be elicited. Now, compare this new knowledge to the knowledge that might be created by one more spy. The CIA already employs 1,600 spies (Miller 2012, para. 3). Is it reasonable to think that a radically different method of gathering intelligence will produce less novel information than increasing the number of spies by 0.0006%?

### Manipulation dilemma

Even if a prediction market on terrorism could attract traders with relevant information, Stiglitz cites a further issue that might prevent it from being effective—the manipulation dilemma: 'If trading is anonymous, then it could be subject to manipulation, particularly if the market has few participants – providing a false sense of security or an equally dangerous false sense of alarm. If trading is not anonymous, then anyone with information about terrorism would be, understandably, reluctant to trade on it. In that case, the market would not serve its purpose' (Stiglitz 2003, para. 11).

It seems unarguable that if participation in the prediction market required the registration of a lot of personal details, then many would-be participants would be deterred. Some would doubtless worry about being tracked down and investigated, harassed, or tortured. Therefore, the 'manipulation by anonymous traders' problem seems to be the best horn of the dilemma to tackle. Hanson (2006a, p. 266–268) has discussed some ways to deal with this problem, but

This is the last author version. The official version is forthcoming in Jai Galliot & Warren Reed (eds.), *Ethics and the Future of Spying: Technology, Intelligence Collection and National Security*, Routledge.

the most useful solution would be a *mainly* anonymous model. The market would be set up to be anonymous and advertise itself as such. However, authorities could gain access to whatever general information is collected through the market. For example, the CIA would know whether suspicious trades are made from the same account, or from the same geographical area. Furthermore, given the right warrant, some financial institutions or intermediaries involved with the transactions might be able to provide more detailed information.

The main worry with anonymous traders manipulating the market, expressed by Posner (2004, p. 175), is that the genuine relevant information gathered by the market will be swamped by the purposefully deceptive trades made by terrorists. Whether manipulative trades will regularly, occasionally, or ever swamp the genuine information is hard to discern theoretically. Nevertheless, several mathematical models indicate that prediction markets can still perform just as well, or even better, when some of the participants are attempting to manipulate the price (Hanson 2006a, p. 266–68; Wolfers & Zitzewitz 2004, p. 16–17). The models show that, in order to survive manipulation, a prediction market needs a sufficient number of participants with some relevant private knowledge and the money and inclination to trade on that knowledge (Hanson 2006a, p. 266–68; Wolfers & Zitzewitz 2004, p. 16–17).

For example, imagine that a group of terrorist participants were selling stock in a particular prediction short so as to reduce the price of that prediction, and thereby reduce the possibility of interference with their planned attack. Cases like this can be referred to as ‘diverting cases’ because the terrorists are attempting to manipulate the market in a way that diverts attention away from the prediction that might warn intelligence agents about their planned attack. In diverting cases, the connected citizens might think that, based on their private information about a planned terrorist act, the price is a bargain because it does not reflect the probability of the event happening and so they would buy as much of the under-priced stock as they could. As a result, the attempted manipulation of the market would have failed and, rather than create a diversion, the terrorists would merely have wasted their money. Even if terrorists could keep the price down, the increased activity in the market, with the traders pushing the price down all coming from the same geographical area, would be a clear sign that terrorists might be attempting a diverting manipulation, and thereby provide the same signal to intelligence agencies that a high price would.

The reverse kind of case, which can be referred to as a ‘disrupting case’, is trickier, however. Terrorists might attempt to buy all the stock in a particular prediction so as to make it look like a terrorist attack is imminent. Possible motives for this include disrupting and scaring their enemies (hence the name), and possibly undermining confidence in the prediction market. If terrorists did buy all the stock of a particular prediction for these reasons, very few participants are likely to have the kind of direct private information required to give them reason to correct the market. It’s close to impossible for connected citizens to have information to the effect that *all* of the main terrorist cells around the world *will not* attack a certain place at a certain time. There is, however, a slim chance that some connected citizens have information about a terrorist plot to artificially inflate that prediction. Most likely, when connected citizens see the price of a stock in the prediction steadily rise, they might assume that other participants have some relevant private information that they do not, and therefore view the price change as genuine. On the other hand, it seems plausible that all genuine participants in the prediction market have the default assumption that any particular terrorist attack is very unlikely. Indeed, this default consideration, in combination with their ambient information, could lead them to see the price rise as attempted manipulation or inept trading by inexperienced participants. If they did see it this way, then the genuine traders might try to sell stock in the prediction until

This is the last author version. The official version is forthcoming in Jai Galliot & Warren Reed (eds.), *Ethics and the Future of Spying: Technology, Intelligence Collection and National Security*, Routledge.

the price returned to relatively normal levels. Given the broad range of factors that may be at work here, it would be presumptuous to claim to know how most, or even many, real-world prediction markets on terrorism would react to attempted disruptive manipulations. Indeed, only a formal real-world trial seems likely to resolve this worry to any significant degree of satisfaction.

Importantly, a prediction market's ability to prevent disrupting cases of terrorist misinformation, is not nearly as significant as its ability to prevent diverting cases. The consequences of successful disrupting manipulations are mainly inconveniences and slight economic costs. The consequences of diverting manipulations, however, could be the calamitous death and destruction that successful terrorist attacks can cause. And, as was just argued, a prediction market on terrorism is much less likely to be successfully manipulated in diverting cases.

So, while some doubts remain, an intelligently designed prediction market on terrorism – with a broad range of traders from around the world – appears to have a relatively good chance of being effective. 'Relatively' is important here, since the likely effectiveness of such a market should be compared to the likely effectiveness of one more spy. Since, as has been argued, there are no clear reasons why a prediction market on terrorism could not produce useful information like it does in so many other domains, and one more spy using the same old methods seems relatively unlikely to produce more than a little new anti-terrorist intelligence, it seems reasonable to conclude that prediction markets are likely to be more effective than one more spy.

### *Insurance*

However, more effective policy options are not always chosen, especially when they produce more moral issues than the alternatives. Stiglitz (2003) raises two moral issues based on the use of prediction markets on terrorism for insurance.

#### Perverse incentives

Stiglitz explains the moral hazard problem of insurance this way: '[B]y providing "insurance" for participants, futures markets can also create the long-noted "moral hazard" problem – the notion that insurance can alter incentives. Someone who has insured his house for 110% of its value has an incentive to set it afire' (Stiglitz 2003, para. 11). The implication is clear, traders who stand to profit if a terrorist attack occurs have an incentive to bring out such a terrorist attack. However, while the moral hazard problem does seem to pose real ethical problems in certain domains in which the benefits may outweigh the costs (e.g., life insurance, property insurance), prediction markets on terrorism is not one of them.

The vast majority of people in the world, and potential participants in a prediction market on terrorism, try to avoid harming others, and abhor the thought of killing an innocent person. Imagine that hundreds of normal people bought some stock in there being a terrorist attack because the price was low and they thought that an attack was more likely than the price indicated. It seems exceedingly unlikely that any people in that situation would suddenly get the idea that they should plan and carry out a terrorist attack just to reap the financial rewards. First, how much money would it take to convince a normal person to commit an act of indiscriminate killing? I suspect no amount would be enough. Second, even if someone could be swayed to kill innocents by a multi-million dollar payday, they might not have the

This is the last author version. The official version is forthcoming in Jai Galliot & Warren Reed (eds.), *Ethics and the Future of Spying: Technology, Intelligence Collection and National Security*, Routledge.

knowledge, connections, and financial resources to carry out the attack. Third, there is no point being rich if you are caught planning or carrying out an act of terrorism because money is not so useful if you are in prison for the rest of your life or dead, shot by security forces. And, being thwarted in this way is likely because by betting heavily on a predicted attack (which is required to substantially profit from it), the trader would alert the CIA to their intentions. Fourth, as Hanson (2006a, p. 269) has suggested, the markets should be set up with an upper limit so that it would be impossible to make more than several thousand dollars, in which case money can be more easily made elsewhere. Would someone who is not already a terrorist become one, and risk their life in doing so, just to kill innocent people for an amount of money that they could make by working another shift or investing in the regular stock market (Hanson 2006a, pp. 268–270)? Surely not; the benefits are swamped by the likely costs.

That just leaves traders who already are terrorists – people who think there is also some other benefit to carrying out a terrorist attack. It makes sense that an *arsonist* would set fire to her house if it was over-insured. Similarly, it makes sense for a *terrorist* to commit an act of terror if he could gain financially as well as furthering his other aims. But, terrorists are those who plan to and do carry out terrorist attacks anyway. So, no new acts of terror are encouraged by the existence of prediction markets on terrorism. Furthermore, if terrorists attempt to make money from the prediction market in this way, then they are doing exactly what will enable the prediction market to help prevent terrorism; they will be sending a clear message of their intentions, allowing the CIA to thwart and possibly capture them. Therefore, the moral hazard problem does not apply to prediction markets on terrorism; the notion that the markets would encourage more of the behaviour they are designed to protect people from is implausible.

#### Further inequities

Stiglitz also raises the insurance-related issue that a prediction market on terrorism would allow ‘those with the sophistication and money to “hedge” against the threat of terrorism, financially at least, leaving the rest of Americans fully exposed!’ (Stiglitz 2003, para. 12). Stiglitz sees this as a problem, stating, ‘the U.S. government should be concerned with the exposure of all Americans to terrorism’ (Stiglitz 2003, para. 12). This argument is flawed in two important ways.

First, note that one of the most harrowing features of terrorism is that it involves indiscriminate killing (Khatchadourian 1998, p. 24). As a consequence, terrorism exposes *everyone* to the threat of being indiscriminately killed! Rich and poor alike were killed in the 9/11 attacks. Rich and poor alike lost loved ones. Rich and poor alike experienced grief and terror for months and years afterward. No amount of insurance could have hedged against that suffering because money cannot fix those problems. Money cannot shield people from the grief caused by losing a loved one to an indiscriminate attack. So, attempting to hedge against terrorism seems misguided, something that will only achieve the narrowing of the wide gap between how much money and sense they have (not by increasing their ‘sense’).

Second, yes, the government should be concerned with the exposure of all Americans to terrorism. But rather than being a reason for banning prediction markets on terrorism, this seems a better reason for implementing them. The exposure of American’s to terrorism gives the government reason to investigate every potentially useful strategy to increase anti-terrorist intelligence gathering efforts. Indeed, this is the main reason for setting up a prediction market on terrorism – to protect everyone from the risk of terror, grief and death.

This is the last author version. The official version is forthcoming in Jai Galliot & Warren Reed (eds.), *Ethics and the Future of Spying: Technology, Intelligence Collection and National Security*, Routledge.

So, while insurance does raise moral concerns, Stiglitz has given no good reason to think that the attempted use of prediction markets on terrorism as insurance poses any moral concern at all. This conclusion places prediction markets on terrorism on an equal footing with one more spy – neither poses any moral issues related to insurance.

### *Other Moral Problems*

Stiglitz (2003) and others (e.g., Dorgan 2003) have also presented two further potential moral problems with prediction markets on terrorism – rewarding informants and betting on death.

#### Rewarding informants

Stiglitz has argued that traders with genuine information about terrorist attacks should be investigated, rather than given a pay-out for their correct predictions: ‘shouldn't these people be investigated rather than rewarded?’ (Stiglitz 2003, para. 10). However, it is not obvious that we could or should investigate such people.

First of all, without the prediction markets, it is unclear how spies would know who to investigate. Second, it seems likely that this kind of ‘investigate everyone’ thinking deters innocent would-be informers. Worries about being investigated are probably seriously restricting the amount of useful intelligence the CIA is currently gathering. And, as for rewarding the participants whose predictions are correct: why not? The participants who make money trading on the prediction market will have been lucky or they will have knowingly contributed useful information. The lucky ones are probably innocent, and their reward will probably be used to fund the pay-outs of other traders on future predictions (their luck will likely run out soon). Of the other participants who are financially rewarded, some will be innocent, others guilty and still others somewhere in between. Rewarding the innocent participants, who bring useful information, seems like the right thing to do. Rewarding the not-so-innocent traders, who bring useful information, is not ideal, but neither are the alternatives.

The important comparison here is between rewarding possibly less-than-saintly traders in the prediction market and spies rewarding clearly less-than-saintly informants by paying them to provide novel intelligence on terrorists’ plans. These alternatives are similar, but, the key difference between the two is the correct alignment of reward and useful information. Using the old-fashioned method of paying informants, we do not always know whether the informant is lying to us. But, we pay them anyway. Now consider prediction markets on terrorism; only those traders who provide useful information get rewarded by participating in prediction markets. Lying or not, if someone won money on a prediction, then they were right and they provided useful anti-terrorist intelligence. Therefore, while neither prediction markets on terrorism nor one more spy would get a clean pass on this issue, the prediction market seems preferable because it avoids the morally and informationally worst outcome of paying bad people for misleading information.

#### Betting on death

Perhaps the most discussed potential moral issue regarding prediction markets on terrorism is that the trades they require to be effective can be seen as a form of betting on death. In reference to PAM—the prediction market with anti-terrorism aspects designed by Hanson and colleagues—Stiglitz said this: ‘Under the proposal ... participants would [be] betting, in effect – and perhaps profiting – on such potential events as an attack by North Korea or an

This is the last author version. The official version is forthcoming in Jai Galliot & Warren Reed (eds.), *Ethics and the Future of Spying: Technology, Intelligence Collection and National Security*, Routledge.

assassination of Yasser Arafat' (Stiglitz 2003, para. 9). Democratic senator, Byron Dorgan, was more expressive: 'This betting parlor on the Internet, ... [this] Internet casino, ... is unbelievably stupid. ... It will be offensive to almost everyone. Can you imagine if another country set up a betting parlor so that people could go in... and bet on the assassination of an American political figure...? ... It is offensive.' (Dorgan 2003, para. 8–11). However, it is far from clear how offensive the idea really is, especially when compared to the relevant alternatives.

A prediction market on terrorism would be nothing like a casino. Profit motives are in play, but, as mentioned above, many traders will have more altruistic motives as well. Furthermore, making wagers on matters of life and death is not always immoral. Consider buying shares in a company that is involved in death, like a coffin-making company, a weapons manufacturer, or a life insurance company. Or simply consider buying life insurance. These transactions are all essentially wagers on matters of life and death, but they are generally considered to be morally permissible. A prediction market on terrorism would be no online 'death pool' – a website that encourages participants to make lists of famous people who might die during the next year, and then pays them real money if they correctly predict the most celebrity deaths. The point and likely effect of prediction markets on terrorism is to respect life by helping prevent terrorist attacks, not to be disrespectful of matters of life and death (Weijers & Richardson 2014b).

Surely some people will find the very idea of betting on terrorist attacks highly offensive. But all policy decisions are about trade-offs. Lives could doubtless be saved by lowering the speed limit on freeways, but that would inconvenience many. Implementing a prediction market on terrorism would effectively be trading off some offense, and a modest financial outlay, against the likely chance of helping to prevent terrorist attacks. Consider how offensive and costly the 9/11 terrorist attacks were. Even helping to prevent one terrorist attack like that would make it worth offending some people and spending a few million tax dollars. Given that a prediction market on terrorism could help prevent many attacks, this trade-off should be a relatively easy one to make.

### *The final analysis*

Of course, the most important trade-off for the purposes of this chapter is between funding a new spy and a prediction market on terrorism. Both options were assumed to have the same financial cost and, erring on the side of caution, both might be said to have at least a slim chance of generating novel anti-terrorism intelligence during any given period. Indeed, it was argued that, mainly because of their novelty as an approach to intelligence-gathering, a prediction market seems likely to be more useful as an intelligence-gathering tool than one more spy. However, if prediction markets on terrorism are offensive, then the right policy decision might still be to fund one more spy. But, that would not be the right decision, since the behaviour required of spies is morally worse than the offense that prediction markets on terrorism might cause.

It would be a mistake to think that spies are politely handed novel intelligence on terrorist activity from reputable sources. If a spy is going to produce any useful information, she likely has to extract it. To have any reasonable chance of being effective, spies must lie, cheat, steal, bribe, trespass, forge documents, invade privacy, threaten, coerce, blackmail, detain, incapacitate and possibly torture. Naturally, these activities are at least offensive, and in most cases clearly immoral. Whether torture is usually, or even *ever*, effective is still debated (e.g.,

This is the last author version. The official version is forthcoming in Jai Galliot & Warren Reed (eds.), *Ethics and the Future of Spying: Technology, Intelligence Collection and National Security*, Routledge.

Blakeley 2011, p. 546–550; Dershowitz 2002, chap. 4; Koh 2004, p. 653). But even if we assume that torture is effective and thereby justified in a very select group of cases, it is still morally problematic. Even when torturing someone in an attempt to prevent a terrorist attack, the act of torture may also harm the torturer. As Weijers and Richardson (2014b, p. 38) have argued, having to inflict cruel suffering on another human being might psychologically and morally scar the torturer in a way that gives even morally permissible torture a stark moral downside.

The moral depravity of torture and the other immoral actions required of effective spies are not easily overlooked, especially once they are out in the open. Even the President has admitted that CIA agents have continually resorted to torture in this post-9/11 period, *even though this runs contrary to our considered moral values* (Dilanian 2014, para. 2). Furthermore, torture and some of the other immoral acts to which effective spies must occasionally resort cannot easily be recast as inoffensive. However, the same is not true of prediction markets on terrorism. If the public first hears about prediction markets on terrorism as a kind of casino in which they can bet on when and where innocents will be slaughtered by terrorists, then of course they will find it offensive. Alternatively, if the public first hears about an anti-terrorism tool that financially incentivises people to share relevant information that they would have otherwise kept to themselves, they are likely to applaud the government for doing all they can to prevent the scourge of terrorism.

On balance then, implementing a prediction market on terrorism will produce some offense, with the amount varying based on the presentation of it to the public, but clearly less than any widespread advertisement of the various immoral activities that spies must engage in to extract useful information. Indeed, what sounds more offensive: trying to prevent a terrorist attack by encouraging people to bet on whether a terrorist attack will occur, or by sticking pins under their fingernails and subjecting them to long periods of waterboarding?

All things considered, implementing a prediction market on terrorism is likely to be more efficacious, given its novelty, and less morally problematic, given its moral issues are limited to offense, than hiring one more spy. Therefore, in pursuit of anti-terrorism intelligence, implementing a prediction market on terrorism is preferable to hiring one more spy.

### *References*

- Blakeley, R 2011, 'Dirty hands, clean conscience? The CIA inspector general's investigation of "enhanced interrogation techniques" in the war on terror and the torture debate', *Journal of Human Rights*, vol. 10, no. 4, pp. 544–561.
- Dershowitz, AM 2002. *Why terrorism works: understanding the threat, responding to the challenge*, Yale University Press, New Haven.
- Dilanian, K 2014, 'Obama admits US tortured detainees', Associated Press, viewed 3 December, <http://www.stuff.co.nz/world/americas/10340771/Obama-admits-US-tortured-detainees>.
- Doherty, B 2012, 'Intrade prediction market targeted by our commodity futures trading commission (for absolutely no good reason)', Hit & Run Blog, viewed 3 December, <http://reason.com/blog/2012/11/26/intrade-prediction-market-targeted-by-ou>.
- Dorgan, B 2003, 'Senators Ron Wyden and Byron Dorgan hold news conference on a terror financing scheme', viewed 28 October 2013, <http://hanson.gmu.edu/PAM/govt/senator-wyden-dorgan-pressconf-7-28-03.txt>.

- This is the last author version. The official version is forthcoming in Jai Galliot & Warren Reed (eds.), *Ethics and the Future of Spying: Technology, Intelligence Collection and National Security*, Routledge.
- Green, KC, Armstrong, JS, & Graefe, A 2007, 'Methods to elicit forecasts from groups: Delphi and prediction markets compared', *Foresight: The International Journal of Applied Forecasting*, vol. 8, no. 8, pp. 17–21.
- Grossman, SJ, & Stiglitz, JE 1976, 'Information and competitive price systems', *The American Economic Review*, vol. 66, no. 2, pp. 246–253.
- 1980, 'On the impossibility of informationally efficient markets', *The American Economic Review*, pp. 393–408.
- Hanson, R 2006a, 'Designing real terrorism futures', *Public Choice*, vol. 128, no. 1, pp. 257–74.
- (2006b). 'Decision markets for policy advice', in Alan S. Gerber and Eric M. Patashnik (eds.), *promoting the general welfare: new perspectives on government performance*, Brookings Institution Press, Washington DC.
- Hanson, R 2007, 'The policy analysis market (a thwarted experiment in the use of prediction markets for public policy)', *Innovations: Technology, Governance, Globalization*, vol. 2, no. 3, pp. 73–88.
- Katchadourian, H 1998, *The Morality of Terrorism*, Peter Lang, New York.
- Koh, HH 2004, 'A world without torture', *Columbia Journal of Transnational Law*, vol. 43, pp. 641–62.
- Looney, RE 2004, 'DARPA's policy analysis market for intelligence: outside the box or off the wall?' *International Journal of Intelligence and Counterintelligence*, vol. 17, no. 3, pp. 405–19.
- Miller, G 2012, 'DIA sending hundreds more spies overseas', *The Washington Post*, 1 December. viewed 12 March 2013, [http://articles.washingtonpost.com/2012-12-01/world/35585098\\_1\\_defense-clandestine-service-cia-spy-agency](http://articles.washingtonpost.com/2012-12-01/world/35585098_1_defense-clandestine-service-cia-spy-agency).
- Posner, RA 2004, *Catastrophe: risk and response*, Oxford University Press, New York.
- Stiglitz, J 2003, 'Terrorism: there's no futures in it', *Los Angeles Times*, 31 July, viewed 14 April 2015, <http://mason.gmu.edu/~rhanson/PAM/PRESS2/LATstiglitz-7-31-03.htm>.
- Surowiecki, J 2004, *The wisdom of crowds*. Doubleday, New York.
- Weijers, D 2013, 'Technology: prediction markets', *Observatory for Responsible Research and Innovation in ICT*, viewed 3 December 2014, <http://responsible-innovation.org.uk/torii/resource-detail/1181>.
- Weijers, D & Richardson, J 2014a, 'Is the Repugnance about betting on terrorist attacks misguided?' *Ethics and Information Technology*, vol. 16, no. 3, pp. 251–62.
- 2014b, 'A moral analysis of effective prediction markets on terrorism', *International Journal of Technoethics*, vol. 5, no. 1, pp. 28–43.
- Wolfers, J & Zitzewitz, E 2004, *Prediction markets* (No. w10504), National Bureau of Economic Research.
- Yeh, PF 2006, 'Using prediction markets to enhance US intelligence capabilities', *Studies in Intelligence*, vol. 50, no. 4, pp. 137–149.