

# Information Markets

KRISTOFFER AHLSTROM-VIJ

## What Are Information Markets?

Information markets (IMs), also known as prediction markets, are markets for trading contracts representing bets on future or otherwise unknown events. For example, we might set up an IM for trading contracts regarding the outcome of an election, such as the US presidential election, in the manner that has been done for some time on the Iowa Electronic Markets (IEM) (<http://tippie.uiowa.edu/iem/>), one of the oldest, and most prominent, online IMs. Assuming that either the Democrats or the Republicans will win, the market will feature contracts stating that the Democrats win and contracts stating that the Republicans win. If we are setting up a winner-takes-all market, the entire sum specified in the contract goes to the party representing the true proposition upon settlement. Say, for example, that the value specified in the contract is \$10. A trader who believes that the Democrats will win can buy the relevant contract for a fraction of \$10 from a trader who believes that the Republicans will win. The profit or loss made by any given trader is determined by what she paid or received for her contracts and the values those contracts settle at. For example, someone who bought a Democrat contract for \$1 will make a profit of \$9 (\$10 received at the contract's maturity minus the \$1 paid for the contract) if the Democrats win, and a loss of \$1 otherwise. The person at the other end of the bet will make a corresponding loss of \$9 (\$10 dollars paid at maturity minus the \$1 received when selling the contract) if the Democrats win, and a corresponding profit of \$1 otherwise.

By facilitating the relevant transactions, an IM elicits information possessed by the traders, and aggregates that information by way of a market value. That value can in turn be interpreted as the percentage chance assigned to the event at issue by the market. Interpreted thus, such market values tend to be highly accurate – which is not to say that they are infallible. For example, IMs failed to predict the outcome of Austria's

---

*A Companion to Applied Philosophy*, First Edition. Edited by Kasper Lippert-Rasmussen, Kimberley Brownlee, and David Coady.

© 2017 John Wiley & Sons, Inc. Published 2017 by John Wiley & Sons, Inc.

referendum on joining the European Union (Jacobsen et al. 2000), and the nomination of Judge Roberts to the US Supreme Court (Abramowicz 2007). In addition, IMs have been shown to be susceptible to the common betting bias of overestimating low probability events and underestimating high probability events (Leigh and Wolfers 2006). Still, the bulk of the evidence paints an encouraging picture as far as the epistemic promise of IMs is concerned. Indeed, in recent years, IMs have outperformed expert groups and opinion polls in a wide variety of domains (Hahn and Tetlock 2006), including politics (Berg and Rietz 2014; Berg, Nelson, and Rietz 2008; Forsythe et al. 1998), sports (Debnath et al. 2003; Deschamps and Gergaud 2007; Luckner, Schröder, and Slamka 2008), business (Chen and Plott 2002; Spann and Skiera 2003), medicine (Polgreen et al. 2007), and entertainment (Pennock et al. 2001).

Consequently, it should come as no surprise that IMs recently also have caught the attention of social epistemologists (Ahlstrom-Vij 2012; Bragues 2009), concerned with how we come to know things as a result of social interactions (Goldman 2011). In fact, as we shall see, IMs are a particularly fruitful object of research for *applied* social epistemology (Coady 2012), given their prospects for making a practical difference to several social institutions and practices. The relevant kind of epistemology is applied in two senses.

First, a social epistemology investigating IMs is applied on what Lippert-Rasmussen (Chapter 1, *The Nature of Applied Philosophy*) calls *the relevance conception* of applied philosophy, captured by Stevenson's (1970) suggestion that applied philosophy is "relevant to 'the important questions of everyday life'" (258). After all, epistemological investigations considering how the accuracy of IMs can be harnessed for purposes of answering questions relevant to the areas within which IMs have already had some success – that is, politics, sports, business, medicine, and entertainment – will help answer questions important to everyday life.

Second, a social epistemology investigating IMs will also be applied on what Lippert-Rasmussen calls the *empirical facts conception* of applied philosophy, by having an "essentially interdisciplinary nature," as Stevenson puts it (1970: 263). This will hold as long as we assume that it will be impossible (in some relevantly strong sense) to investigate IMs from a social epistemological point of view without factoring in the empirical details regarding how IMs work and why they are as successful as they are.

We have already discussed how IMs work. The plan for the remainder of the chapter is as follows. The next section will explain why IMs are as successful as they are. Then, a number of problems with IMs, as well as some potential solutions, will be surveyed. The final section will elaborate on (a) what we can learn from studying IMs about the nature, value, and methods of applied social epistemology, and (b) a number of social epistemological and political philosophical questions raised by actual as well as potential future uses of IMs.

## Information Markets and Expertise

Since Surowiecki's *The Wisdom of Crowds* (2004), it has become popular to describe the success of social epistemological phenomena with reference to "the wisdom of crowds." One of Surowiecki's examples of such wisdom is that of Francis Galton collecting a

large number of guesses by livestock exhibition goes about the weight of an ox, and finding the average to come extremely close to the actual weight of the ox. According to Surowiecki, this is exactly what we should expect to have happened, on account of a particular “mathematical truism” (see Chapter 6, *The Epistemology of Deliberative Democracy*; Chapter 27, *Deliberative Democracy*): “If you ask a large enough group of diverse, independent people to make a prediction or estimate a probability, and then average those estimates, the errors each of them makes in coming up with an answer will cancel themselves out” (Surowiecki 2004: 10).

The underlying thought is well known from so-called miracle of aggregation arguments for majority rule in democratic theory (see, e.g., Landemore 2013; Page and Shapiro 1993). Surowiecki (2004: 21) suggests that we can explain the success of IMs with reference to the same truism, but it is not clear that he is right about that. (Hanson 2013, too, calls into question the idea that the success of IMs is to be explained with reference to the “wisdom of crowds.”) IMs certainly require a type of *diversity*, simply on account of the no-trade theorem: unless traders on the market are willing to take on opposing bets, the market unravels (Wolfers and Zitzewitz 2006). It is not clear, however, that IMs require *independence*, in the sense Surowiecki is concerned with. In the case of Galton, independence – in the sense of individuals not simply making their judgment on the basis of what judgments others are making – is important because its absence means that the judgments in the group might systematically point in the same direction. That, in turn, will mean that any errors will *not* cancel each other out, and that the relevant mathematical truism will not apply. Consequently, if Surowiecki’s account of the wisdom of crowds explains the success of IMs, IMs will not be successful in cases of substantial dependency. As we shall see, IMs can be successful in cases of substantial dependency, which means that Surowiecki’s account does not explain the success of IMs.

To see why IMs can be successful in cases of substantial dependency, consider the fact that IMs can generate highly accurate outputs even with a strong herd effect on the market – so long as at least some informed traders on the market are willing to bet against the herd. Indeed, empirical evidence suggests that this is typically how successful IMs work: there is a large number of uninformed (and possibly systematically mistaken) traders on the market providing liquidity, and a small number of informed traders winning bets by taking the positions opposite to those of the uninformed. More specifically, Forsythe and colleagues (1999) found that IMs tend to be driven by a minority of *marginal traders*, which tend to trade higher-than-average sums, be active on the market on a higher-than-average number of days, show a lower-than-average degree of biases, and earn higher-than-average returns. Since they thereby have a disproportionately strong influence on the price signal, and the price signals of IMs moreover tend to constitute accurate verdicts, as I argue above, we have reason to believe that the accuracy of IMs are to a great extent due to the informed trades of marginal traders.

Does the success of IMs thereby presuppose a minority ripping off a majority? No. IMs *can* generate accurate predictions by having a minority make money off an uninformed majority, but accurate prediction does not presuppose such an arrangement. As we shall see in the next section, when it comes to the type of IM that would seem to raise the most acute worries about justice – that is, political prediction markets – play-money

markets are as reliable as real-money markets. Consequently, there is no need to rip anyone off in order to harness the predictive power of IMs.

Nevertheless, given the dependence of accurate prediction on uninformed traders, we might be tempted to follow Landemore in saying about the wisdom of the informed minority on IMs that, since the relevant wisdom only emerges on the market in the presence of an uninformed majority, “this wisdom is, ultimately, that of all the participants in the market, not just that of the few who have the correct information” (2013: 176). That, however, seems misleading at best and inaccurate at worst. Rather than saying that IMs are successful on account of the wisdom *of* crowds, it seems more accurate to say that IMs succeed by harnessing the wisdom *in* crowds, including in crowds where the majority is uninformed, and possibly systematically so. A different way to put the same point is by saying that IM prices tend to track expertise. In fact, there is more direct evidence for this hypothesis. Oliven and Rietz (2004) found that, on the 1992 presidential vote share market on the IEM, the average error rate of *market makers* – the subset of marginal traders that are particularly active in setting bid and ask prices – was close to one-sixth of that of the price takers – the traders who mostly accept others’ prices.

### Some Problems (and Solutions)

The account provided in the previous section of why IMs are successful receives further support from what we know about the conditions under which IMs fail. The main problem for IMs has turned out to be attracting a sufficiently large number of uninformed traders (Wolfers and Zitzewitz 2006). That is to be expected, if IMs tend to succeed on account of a minority of informed traders that provide accuracy by betting against a majority of uninformed traders, providing liquidity. As we have already seen, uninformed traders are needed on account of the no-trade theorem: in the absence of uninformed traders, the informed traders have no one to trade with. To prevent this from happening, the creators of the relevant markets can find ways to appeal to the entertainment value of participating in the market, or exploit our tendency for overconfidence, on account of which we typically overestimate what we know (Wolfers and Zitzewitz 2006; see Chapter 8, Epistemology for (Real) People). Alternatively, the creators can introduce an automated market maker, which continually responds to trade orders within pre-specified parameters for purposes of reducing the number of unfulfilled bets (Diemer and Poblete 2010). The latter option is particularly appealing in cases where we find it questionable to lure uninformed traders onto the market, and prefer to subsidize the market by in effect paying for the accurate predictions generated by informed traders.

If IMs were to become more widely relied upon, liquidity worries might be replaced by market manipulation worries. Consider IMs that predict elections, such as the IEM. If the reliability of such markets were to become common knowledge, powerful interests might start to pump money into the relevant markets for purposes of moving the market value, and thereby influencing people’s expectations and voting behavior. However, according to Rhode and Strumpf (2004), there is no historical evidence that IMs can be systematically manipulated beyond short periods of time. Political betting markets

on Wall Street in the late nineteenth and early twentieth centuries involved millions of dollars and many attempts at manipulation, but the prices consistently returned to their pre-manipulation levels within days. More recently, a series of random investments made by Rhode and Strumpf in IEM's 2000 presidential market led to large initial price changes, but the prices reverted to their initial levels in a few hours. Indeed, Hanson even suggests that "the more manipulation is expected, the more accurate are prices" (2013: 169), since manipulators create lucrative opportunities for informed traders. This, too, is consistent with the above account of the success of IMs, in terms of a minority of informed traders moving the market value by trading with an uninformed majority. More than that, Hanson's claim has been corroborated in theory (Hanson and Oprea 2009), experimentally (Hanson, Oprea, and Porter 2006; Oprea et al. 2007), as well as in the field (Berg and Rietz 2014; Camerer 1998; Wolfers and Zitzewitz 2004).

If liquidity increases, we might worry that the value of the trades made will be so high that some informed traders will not be able to afford to enter the market. This, too, is a worry about powerful interests having an undue influence: if participating in the relevant markets is so expensive that it effectively shuts out anyone not aligned with the relevant interests, there is no need to manipulate the market; the relevant interests will have a monopoly on setting the market price. If the lack of liquidity on the part of informed traders thereby shut out in turn tracks unjust distributions of wealth, this is not only an epistemic worry but also a worry about justice. The most common response to this worry is for the relevant IMs to operate with play rather than real money. While Servan-Schreiber and colleagues (2004) found no difference in accuracy between real-money and play-money IMs, Rosenbloom and Notz (2006) found slightly higher levels of accuracy for real-money markets. Consistent with the latter, Diemer and Poblete (2010) argued that, in direct comparison, real-money markets are slightly more accurate. However, it turns out that the context matters: McHugh and Jackson (2012) showed that, when it comes to the kind of IM where justice considerations would seem to matter the most, that is, political IMs, there is no accuracy difference between real- and play-money markets. (The same turned out to be the case for IMs dedicated to sport.) Consequently, any worries regarding financial barriers to entry to IMs seem surmountable.

Another worry about IMs is that they seem to leave room for insider trading. In fact, for all we know, "expertise" on an IM can often consist in having access to non-public information. Insider trading is prohibited on stock markets, and if IMs were to become subject to the same regulatory framework as such markets, similar prohibitions might come to apply to IMs. However, it is not clear that such a prohibition would serve the epistemic ends of IMs, given the manner in which the incorporation of non-public information on the market would seem to promote accuracy. For further discussion, see Abramowicz (2007: 208–211).

### Information Markets and Applied Social Epistemology: Three Applications

Given the epistemic virtues of IMs, and the fact that the most prominent obstacles to operating and relying on such markets seem surmountable, it should come as no surprise that there has been an increasing number of suggestions in the literature for how

they can be fruitfully applied. Three types of suggestions make for particularly fertile areas of study. As we shall see, each suggestion illustrates ways in which studying IMs tells us something important about the nature, value, and methods of applied social epistemology, and also raises a number of significant questions in epistemology and political philosophy.

### *Aggregating Information*

The most straightforward application of IMs serves to *aggregate information* spread among a large number of people. In recent years, IMs have been used for aggregative purposes in order to inform political campaigns (Berg, Nelson, and Rietz 2008), political parties (Berg and Rietz 2003), and policy makers (Wolfers and Zitzewitz 2009). One of the most ambitious attempts to inform policy by way of IMs – the Pentagon’s Policy Analysis Market (PAM), which would have used IMs to approximate economic growth, military activity, and political stability in a number of Middle Eastern countries – was canceled when described by some as an opportunity to profit from terrorism (Hanson 2006). This is unfortunate, not only because the description was inaccurate, but also because we can expect IMs to provide helpful, predictive tools in exactly these types of contexts, given what we have seen about the conditions under which IMs are successful: IMs provide a platform for a large number of people to reveal their judgments through bets, in a context where the market value moreover tracks the judgments of the informed. As such, IMs are *selectively* aggregative in a manner favoring the informed. In fact, the selectively aggregative power of IMs can be put to use in a more broadly political context that is not restricted to policy making and campaigning. More specifically, it can inform public opinion in a more successful manner than the kind of social deliberation traditionally associated with democratic citizenship (Ahlstrom-Vij 2012). To see why, consider the following:

When it comes to the information relevant for making informed political decisions, the majority tends to be uninformed. Indeed, according to Bartels (1996: 194), “[t]he political ignorance of the American voter is one of the best-documented features of contemporary politics.” Moreover, given what Friedman (2005: x) refers to as an “ocean of findings about political ignorance,” we cannot assume that Americans are unique in this respect. (For a general overview of the implications of public ignorance for the epistemology of democracy, see Ahlstrom-Vij forthcoming.) Given such public ignorance, we have reason to be skeptical about the deliberative democrat’s claim that social deliberation ought to be an integral part of democratic citizenship (e.g., Gutmann and Thompson 2004; Talisse 2005; see Chapter 6, *The Epistemology of Deliberative Democracy*, Chapter 11, *Freedom of Expression, Diversity, and Truth*, and Chapter 27, *Deliberative Democracy*). The reason is that we should expect social deliberation to either (a) suppress the information of informed minorities on account of informational and social pressure (Sunstein 2006), or (b) fall prey to *the common knowledge effect*, on which “[t]he influence of a particular item of information [on the judgment of a group] is directly and positively related to the number of group members who have knowledge of that item before the group discussion and judgment” (Gigone and Hastie 1993: 960). By contrast, if the above account is correct in taking IMs to succeed on account of enabling an informed minority to bet against an uninformed majority, we have reason

to believe that IMs tend to do a good job of harnessing accurate information exactly under the conditions where social deliberation does not: namely, in contexts of widespread ignorance (public or otherwise).

It is worth also noting that, since we are here concerned with deliberation, and as such with situations wherein there is a failure of independence, considerations regarding the wisdom of crowds fail to apply. The same goes for Hong and Page's (2004) diversity trumps ability theorem, according to which "a random collection of agents drawn from a large set of limited-ability agents typically outperforms a collection of the very best agents from that same set" (2004: 16386). Since Hong and Page's model does not take into account *communication* (16389), it has nothing to do with social deliberation. Moreover, even if we restrict ourselves to non-deliberative contexts, note Landemore's (2013) point that, since the theorem is decreasingly likely to hold as the relevant collection of agents grows, it "could seem slightly sobering from a democratic point of view favoring maximal inclusiveness" (164).

It might be objected that we simply need to work harder to make social deliberation work. After all, there is some evidence from Ackerman and Fishkin's (2004) work on "deliberative polling" suggesting that, under highly controlled and monitored settings, good information can be properly harnessed, spread, and utilized through social deliberation. However, Ackerman and Fishkin also acknowledge that mirroring the relevant settings in practice would require "a rethinking of the deliberative process from the ground up" (2004: 5). This raises questions about the feasibility of the relevant kind of reform – questions that become particularly pressing when considering that IMs seem to perform well without the need for substantial reform.

Of course, this point about feasibility would carry less weight were social deliberation valuable in ways that could not be fully captured in instrumental terms. For example, Gutmann and Thompson (2004) suggest that social deliberation also has an *expressive* value, because "[b]y deliberating with one another, decision-makers manifest mutual respect toward their fellow citizens" (22). At the same time, they also maintain that the expressive value of deliberative democracy depends on its instrumental value: "If deliberation tended to produce worse decisions than other processes in the long run, then it would not serve the expressive purpose" (22). Arguably, IMs constitute one such process, which is why what we have seen so far calls into question not just the instrumental but also the expressive value of social deliberation.

A related objection to the feasibility claim is that social deliberation, unlike IMs, can make the resulting policies politically legitimate. One of the most influential notions of political legitimacy is captured by Estlund's (2008) *qualified acceptability criterion*, a necessary condition on legitimacy on which "[n]o one has authority or legitimate coercive power over another without a justification that could be accepted by all qualified points of view" (33). However, as argued by Ahlstrom-Vij (2012), on any plausible reading of what makes points of view qualified, IMs carry more promise than does social deliberation when it comes to offering the kinds of justifications relevant to political legitimacy. Needless to say, given the large literature on political legitimacy, further work is required for purposes of determining whether this claim generalizes to other, plausible notions of political legitimacy. Still, given the centrality of conditions like the qualified acceptability criterion to contemporary political philosophy – indeed, Brennan (2011) refers to it as a "defining feature of liberal political philosophy"

(714) – the aforementioned result at the very least throws doubt upon the claim that social deliberation is alone in being able to confer legitimacy upon policies.

It should be stressed that it is not being claimed that IMs never fail or face no obstacles; the claim is simply that, under a prevalent type of condition wherein there is widespread ignorance, IMs do a significantly *better* job of harnessing the insights of the informed than does social deliberation. Consequently, Landemore’s claim that “[f]rom a comparative point of view, information markets are no worse than deliberation in giving better information greater weight” (2013: 182) is clearly an understatement.

Moreover, by reflecting on the role that IMs thereby can play in informing politics, policy, and public opinion, we also learn something important about the nature, method, and value of applied social epistemology that relates back to what was said at the outset of this chapter. After all, establishing that IMs can play such a role requires us to be informed about the relevant empirical research on IMs, and as such highlights the naturalistic nature of applied social epistemology – naturalism being the view that the methods of philosophy are continuous with those of the empirical sciences. This speaks to the empirical facts conception of applied philosophy. Moreover, due to our reliance on empirical research on the matter, the relevant kind of investigation also has to extend across disciplinary boundaries, and as such be interdisciplinary in method. Finally, given that IMs can be shown to help us think constructively about how to make better decisions, including in contexts crucial to people making informed political choices, the value of an applied epistemology engaging in the relevant kind of investigation should be manifest, rendering that investigation an instance of the relevance conception.

### *Voting on Values, Betting on Beliefs*

So far, we have seen that IMs can aggregate information in a highly effective manner. As such, IMs can also be used to help us make more informed decisions in contexts where we need to evaluate whether something is an effective means to a certain end. Hanson elaborates on this idea in relation to political policies in particular, and suggests that IMs provide a way to rethink government for the better:

To embed such markets in the core of our form of government, we could “vote on values, but bet on beliefs.” Today we vote on both to say what we want, and to say how to get what we want. We might instead continue to vote to say what we want, but use speculative markets to say how to get what we want. (Hanson 2013: 152)

More specifically, Hanson’s proposal involves retaining an elected legislature, the primary job of which would be to develop and maintain a measure of national welfare. On account of being elected, the legislature’s proposals for such a measure would be sensitive to our votes and thereby to what we want. Policy proposals would in turn be evaluated by way of IMs predicting the impact of the relevant policies on national welfare. A proposal becomes law when it is predicted to have not only a significantly positive impact on welfare but also a significantly *higher* such impact compared to available alternatives.

Hanson’s proposal invokes a kind of IM we have not encountered in the above: the *conditional* IM. The contracts used when evaluating a proposal would be formulated

as follows: *If policy P is implemented, national welfare would be affected in such-and-such a manner.* Bets regarding contracts conditional on policies that are not implemented are called off. As Hanson points out, “this called-off-bet approach gives us speculator estimates on the consequences of events that *never actually happen*,” without preventing the same mechanisms generating accurate predictions from operating prior to the relevant bets being called off: “until speculators know an event won’t happen, they can have incentives to accurately estimate its consequences” (2013: 159). Contracts pertaining to the policies actually implemented are then settled with reference to an after-the-fact measurement of welfare. (For more on conditional IMs and the unwinding of contracts, see Abramowicz 2007: 141–144.)

Whether or not Hanson’s proposal is ultimately feasible – Hanson himself suggests that his goal is not to “induce high confidence in this concept [of voting on values while betting on beliefs], but to raise reader confidence to a level justifying further exploration via more realistic prototypes” (2013: 162) – it serves to illustrate nicely the nature, method, and value of the kind of applied epistemology we have been concerned with here, for much the same reasons as those given above. After all, such an epistemology will be naturalistic, in being informed about the relevant empirical research on how IMs work and why they are epistemically successful. On that account, the relevant kind of epistemology will also be interdisciplinary in method. Finally, since it helps us think constructively about how to make better decisions about what policies to adopt in the pursuit of common and democratically determined ends, the relevant kind of investigation will also be valuable because relevant to the improvement of actual, socially important decisions.

### *Aggregating Preferences*

As we have seen, Hanson (2013) suggests that we vote on values while betting on beliefs. It is possible to go even further and bet on both, given that IMs also can be used to *aggregate preferences*. For example, Abramowicz (2007) suggests that IMs can be used for purposes of aggregating preferences regarding costs and benefits, by making for a kind of “predictive cost–benefit analysis” (173). A further example of how IMs can aggregate preferences is provided by political IMs predicting elections, on IEM and elsewhere (Berg and Rietz 2014; Berg, Nelson, and Rietz 2008). What those markets illustrate is that there are (at least) two ways to aggregate preferences: you can *ask* people what they want – for example, by holding an election or conducting a poll – or you can *predict* what they want by setting up an IM on which people can bet on what people’s preferences are. It might be objected that such IMs would have to be resolved with reference to an actual election, and as such never could *replace* elections (or polls). However, as argued by Abramowicz (2007), there is nothing in principle that prevents an IM from being *self-resolving*, that is, having its contracts be settled, not on the basis of some external event (e.g., an election outcome), but on account of some factor internal to the market. For example, a market could settle its contracts on the basis of the market value at some pre-specified time, unknown to the traders.

If reasons can be provided that self-resolving IMs are as reliable as externally resolved IMs – and as far as available evidence is concerned, that is a big “if” – we may raise questions about whether IMs could actually replace the voting mechanism.

How can we trust that a self-resolving IM will give us an accurate account of the popular will, if not resolved by an actual election? Presumably, on the same grounds that we would trust any other process: if we have strong enough evidence that process *P* tracks some phenomenon *Q*, we can trust *P*'s outputs to mirror *Q* without independent verification. This is typically how we treat electoral processes: we do not run an additional election to verify the outcome of the first one, nor do we verify the outcome through some independent, non-electoral process. Such trust is, of course, not unconditional. For example, if we find the electoral outcome to be radically different from what we expect on the grounds of polling data, we might start to question the electoral process. If further research were to show that self-resolving IMs are as reliable as externally resolved IMs, self-resolving IMs gauging the popular will might over time attain the same status as traditional elections, in that we might come to trust such markets to reveal the popular will, at least in the absence of independent evidence suggesting that something has gone awry.

Setting aside questions about feasibility, the mere possibility of self-resolving IMs replacing elections raises interesting political philosophical questions, such as: Are elections a necessary component for democracy? If Estlund (1990) is right that democracy is not merely rule *for* the people – that is, rule in accordance with people's preferences, conceptions of justice, or beliefs about what is in the common good – but also rule *by* the people, in a sense that makes voting necessary for democracy, such a replacement would have the relevant regime cease to be democratic. If voting is not necessary for democracy, however, IMs potentially provide a powerful tool in the service of having democratic systems be maximally responsive to the popular will. (Those who believe that votes do not express preferences, but rather constitute statements about what is just – as suggested by Rawls 1971 – or in the common interest – as argued by Estlund 1990 – can simply imagine that the traders are betting on what people take to be just or in the common interest.)

In these cases, too, we learn something about the nature, method, and value of an applied social epistemology concerned with IMs. Such epistemology will be naturalistic in nature, since informed about the relevant empirical research, and moreover interdisciplinary in method – all of which is in accordance with the relevance conception. Moreover, since the relevant kind of investigation helps us think constructively about how to get a more accurate picture of what people want, be it in a political context or in the context of cost–benefit analyses, the relevant kind of epistemology is clearly valuable in the sense highlighted by the relevance conception of applied epistemology. Additionally, reflecting on the possibility of self-resolving markets as an alternative to the voting mechanism is valuable on purely philosophical grounds, on account of the questions it raises about process trust and the conceptual relationship between voting and democracy.

## Conclusion

To sum up, IMs tend to generate highly accurate outputs on account of harnessing the insights of informed minorities. Moreover, available evidence suggests that the major challenges faced by such markets – relating to illiquidity, manipulation,

and barriers to entry – are surmountable. Finally, there are a number of promising applications of IMs that enlighten us on the nature, value, and methods of applied social epistemology, and also raise a number of significant questions in epistemology and political philosophy. While the field of IM research is still new, and the philosophical study of them newer still, the present overview should make clear that IMs warrant further investigation by applied philosophers, not least on account of how such markets provide a novel and promising tool for informing socially important decisions.

## References

- Abramowicz, M. 2007. *Predictocracy: Market Mechanisms for Public and Private Decision Making*. New Haven, CT: Yale University Press.
- Ackerman, B. and Fishkin, J. 2004. *Deliberation Day*. New Haven, CT: Yale University Press.
- Ahlstrom-Vij, K. 2012. "Why Deliberative Democracy is (Still) Untenable." *Public Affairs Quarterly* 26(3): 199–220.
- Ahlstrom-Vij, K. Forthcoming. "The Epistemic Benefits of Democracy: A Critical Assessment." In *The Routledge Handbook of Social Epistemology*, edited by M. Fricker, P. J. Graham, D. Henderson, N. Pedersen, and J. Wyatt. Routledge.
- Bartels, L. 1996. "Uninformed Voters: Information Effects in Presidential Elections." *American Journal of Political Science* 40(1): 194–230.
- Berg, J. and Rietz, T. 2003. "Prediction Markets as Decision Support Systems." *Information Systems Frontiers* 5(1): 79–93.
- Berg, J. and Rietz, T. 2014. "Market Design, Manipulation and Accuracy in Political Prediction Markets: Lessons from the Iowa Electronic Markets." *Political Science and Politics* 47(2): 293–296.
- Berg, J., Nelson, E., and Rietz, T. 2008. "Prediction Market Accuracy in the Long Run." *International Journal of Forecasting* 24: 285–300.
- Bragues, G. 2009. "Prediction Markets: The Practical and Normative Possibilities for the Social Production of Knowledge." *Episteme* 6: 91–106.
- Brennan, J. 2011. "The Right to a Competent Electorate." *Philosophical Quarterly* 61(245): 700–724.
- Camerer, C. 1998. "Can Asset Markets Be Manipulated?" *Journal of Political Economy* 106: 457–482.
- Chen, K.-Y. and Plott, C. 2002. "Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem." CalTech Social Science Working Paper No. 1131.
- Coady, D. 2012. *What to Believe Now: Applying Epistemology to Contemporary Issues*. Malden, MA: Wiley-Blackwell.
- Debnath, S., Pennock, D., Lawrence, S., and Giles, C.L. 2003. "Information Incorporation in Online In-game Sports Betting Markets." *Proceedings of the 4th Annual ACM Conference on Electronic Commerce (EC'03)*: 258–259.
- Deschamps, B. and Gergaud, O. 2007. "Efficiency in Betting Markets: Evidence From English Football." *The Journal of Prediction Markets* 1: 61–73.
- Diemer, S. and Poblete, J. 2010. "Real-Money vs. Play-Money Forecasting Accuracy in Online Prediction Markets – Empirical Insights from iPredict." *Journal of Prediction Markets* 4(3): 21–58.
- Estlund, D. 1990. "Democracy Without Preference." *The Philosophical Review* 99(3): 397–423.

- Estlund, D. 2008. *Democratic Authority: A Philosophical Framework*. Princeton, NJ: Princeton University Press.
- Forsythe, R., Frank, M., Krishnamurthy, V., and Ross, T. 1998. "Markets as Predictors of Election Outcomes: Campaign Events and Judgment Bias in the 1993 UBC Election Stock Market." *Canadian Public Policy* 24: 329–351.
- Forsythe, R., Rietz, T., and Ross, T. 1999. "Wishes, Expectations and Actions: A Survey on Price Formation in Election Stock Markets." *Journal of Economic Behavior and Organization* 39: 83–110.
- Friedman, J. 2005. "Popper, Weber, and Hayek: The Epistemology and Politics of Ignorance." *Critical Review* 17(1–2): i–lviii.
- Gigone, D. and Hastie, R. 1993. "The Common Knowledge Effect: Information Sharing and Group Judgment." *Journal of Personality and Social Psychology* 65(5): 959–974.
- Goldman, A. 2011. "A Guide to Social Epistemology." In *Social Epistemology: Essential Readings*, edited by A. Goldman and D. Whitcomb, 11–37. Oxford: Oxford University Press.
- Gutmann, A. and Thompson, D. 2004. *Why Deliberative Democracy?* Princeton, NJ: Princeton University Press.
- Hahn, R. and Tetlock, P. 2006. *Information Markets: A New Way of Making Decisions*. Washington, DC: AEI Press.
- Hanson, R. 2006. "Designing Real Terrorism Futures." *Public Choice* 128(1–2): 257–274.
- Hanson, R. 2013. "Shall We Vote on Values, But Bet on Beliefs?" *Journal of Political Philosophy* 21(2): 151–178.
- Hanson, R. and Oprea, R. 2009. "Manipulators Increase Information Market Accuracy." *Economica* 76(302): 304–314.
- Hanson, R., Oprea, R., and Porter, D. 2006. "Information Aggregation and Manipulation in an Experimental Market." *Journal of Economic Behavior and Organization* 60: 449–459.
- Hong, L. and Page, S. 2004. "Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers." *Proceedings of the National Academy of Sciences of the United States of America* 101(46): 16385–16389.
- Jacobsen, B., Potters, J., Schram, A., van Winden, F., and Wit, J. 2000. "(In)accuracy of a European Political Stock Market: The Influence of Common Value Structures." *European Economic Review* 44(2): 205–230.
- Landemore, H. 2013. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton, NJ and Oxford: Princeton University Press.
- Leigh, A. and Wolfers, J. 2006. "Competing Approaches to Forecasting Elections: Economic Models, Opinion Polling, and Prediction Markets." *Economic Record* 82(258): 325–337.
- Luckner, S., Schröder, J., and Slamka, C. 2008. "On the Forecast Accuracy of Sports Prediction Markets." In *Negotiation, Auctions & Market Engineering, Lecture Notes in Business Information Processing (LNBIIP)*, edited by H. Gimpel, N.R. Jennings, G. Kersten, A. Okenfels, and C. Weinhardt, 227–234. Dordrecht: Springer.
- McHugh, P. and Jackson, A. 2012. "Prediction Market Accuracy: The Impact of Size, Incentives, Context and Interpretation." *The Journal of Prediction Markets* 6(2): 22–46.
- Oliven, K. and Rietz, T. 2004. "Suckers Are Born but Markets are Made: Individual Rationality, Arbitrage, and Market Efficiency on an Electronic Futures Market." *Management Science* 50(3): 336–351.
- Oprea, R., Porter, D., Hibbert, C., Hanson, R., and Tila, D. 2007. "Can Manipulators Mislead Market Observers?" Chapman University, E.S.I. Working Papers 08-01. Available at <http://www.chapman.edu/ESI/wp/CanManipulatorsMisleadMarketObservers.pdf>
- Page, B. and Shapiro, R. 1993. "The Rational Public and Democracy." In *Reconsidering the Democratic Public*, edited by G. Marcus and R. Hanson, 33–64. University Park, PA: Pennsylvania State University Press.

- Pennock, D., Lawrence, S., Nielsen, F.A., and Giles, C.L. 2001. "Extracting Collective Probabilistic Forecasts from Web Games." *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 174–183.
- Polgreen, P., Nelson, E., Neumann, G., and Weinstein, R. 2007. "Use of Prediction Markets to Forecast Infectious Disease Activity." *Clinical Infectious Diseases* 44: 272–279.
- Rawls, J. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Rhode, P. and Strumpf, K. 2004. "Historical Presidential Betting Markets." *Journal of Economic Perspectives* 18(2): 127–141.
- Rosenbloom, E.S. and Notz, W. 2006. "Statistical Tests of Real-Money Versus Play-Money Prediction Markets." *Electronic Markets* 16(1): 63–69.
- Servan-Schreiber, E., Wolfers, J., Pennock, D., and Galebach, B. 2004. "Prediction Markets: Does Money Matter?" *Electronic Markets* 14(3): 243–251.
- Spann, M. and Skiera, B. 2003. "Internet-Based Virtual Stock Markets for Business Forecasting." *Management Science* 49: 1310–1326.
- Stevenson, L. 1970. "Applied Philosophy." *Metaphilosophy* 1(3): 258–267.
- Sunstein, C. 2006. *Infotopia: How Many Minds Produce Knowledge*. Oxford: Oxford University Press.
- Surowiecki, J. 2004. *The Wisdom of Crowds: Why the Many Are Smarter than the Few*. London: Abacus.
- Talisse, R. 2005. *Democracy After Liberalism: Pragmatism and Deliberative Politics*. New York: Routledge.
- Wolfers, J. and Zitzewitz, E. 2004. "Prediction Markets." *Journal of Economic Perspectives* 18(2): 107–126.
- Wolfers, J. and Zitzewitz, E. 2006. "Five Open Questions about Prediction Markets." In *Information Markets: A New Way of Making Decisions*, edited by R. Hahn and P. Tetlock, 13–36. Washington, DC: AEI Press.
- Wolfers, J. and Zitzewitz, E. 2009. "Using Markets to Inform Policy: The Case of the Iraq War." *Economica* 76: 225–250.

### Further Reading

- Abramowicz, M. 2007. *Predictocracy: Market Mechanisms for Public and Private Decision Making*. New Haven, CT: Yale University Press.  
The most comprehensive treatment available of the nature and possible applications of information markets.
- Berg, J. and Rietz, T. 2014. "Market Design, Manipulation and Accuracy in Political Prediction Markets: Lessons from the Iowa Electronic Markets." *Political Science and Politics* April: 293–296.  
A helpful overview of current issues with, and possibilities for, information markets in light of the Iowa Electronic Markets.
- Bragues, G. 2009. "Prediction Markets: The Practical and Normative Possibilities for the Social Production of Knowledge." *Episteme* 6: 91–106.  
A good introduction to information markets for a philosophical audience.
- Hahn, R. and Tetlock, P. 2006. *Information Markets: A New Way of Making Decisions*. Washington, DC: AEI Press.  
A collection of papers on information markets by several of the foremost people working on the theory and application of such markets.
- Hanson, R. 2006. "Designing Real Terrorism Futures." *Public Choice* 128(1–2): 257–274.  
A discussion about the predictive possibilities of information markets in matters of foreign affairs, in light of the cancellation of the Pentagon's Policy Analysis Market (PAM).

KRISTOFFER AHLSTROM-VIJ

Hanson, R. 2013. "Shall We Vote on Values, But Bet on Beliefs?" *Journal of Political Philosophy* 21(2): 151–178.

An interesting proposal for how to apply information markets for political philosophical purposes.

McHugh, P. and Jackson, A. 2012. "Prediction Market Accuracy: The Impact of Size, Incentives, Context and Interpretation." *The Journal of Prediction Markets* 6(2): 22–46.

A good discussion of a variety of factors potentially affecting the accuracy of information market output.

UNCORRECTED PROOFS