

Shall We Vote on Values, But Bet on Beliefs?

By Robin Hanson, December 2012. To appear in *Journal of Political Philosophy*.

Introduction

Policy disputes arise at all scales of governance, in clubs, non-profits, firms, nations, and alliances of nations. Both the means and ends of policy are disputed. While many, perhaps most, policy disputes arise from conflicting ends, important disputes also arise from differing beliefs on how to achieve shared ends. In fact, according to many experts in economics and development, governments often choose policies that are “inefficient” in the sense that most everyone could expect to gain from other feasible policies. Many other kinds of experts also see existing policies as often clearly inferior to known alternatives.

If inferior policies would not have been adopted had most everyone known they are inferior, and if someone somewhere knew or could have learned that they are inferior, then we can blame inferior policies on a failure of our “info” institutions. By “info” here I just mean clues and analysis that should change our beliefs. Our info institutions are those within which we induce, express, and evaluate the acquiring and sharing of info. They include public relations teams, organized interest groups, news media, conversation forums, think tanks, universities, journals, elite committees, and state agencies. Inferior policies happen because our info institutions fail to induce people to acquire and share relevant info with properly-motivated decision makers.

Yes, disputes that appear to be about means may actually about ends, economists and other experts may greatly overestimate how often there exist policy changes that could benefit most everyone, and contrary to appearances current info institutions may now acquire and share pretty much all useful info. This paper, however, will focus on the opposite possibility, that large policy gains may be possible from better info institutions.

Where might we find better info institutions? According to most experts in economics and finance, speculative markets are exemplary info institutions. That is, active speculative markets do very well at inducing people to acquire info, share it via trades, and collect that info into consensus prices that persuade wider audiences. This great success suggests that we consider augmenting our political info institutions with speculative market institutions. That is, perhaps we should encourage people to create, trade in, and heed policy-relevant speculative markets, instead of discouraging such markets as we do today via anti-gambling laws.

While many intellectual strategies are available for exploring this possibility, this paper will use the colorful strategy of taking the idea to an extreme. In particular, I will consider modifying the basic form of democratic national government to take full advantage of the info abilities of speculative markets. While such an extreme approach may not be required to gain policy advantages from speculative markets, it might better help us to explore their potential, and limits.

To embed such markets in the core of our form of government, we could “vote on values, but bet on beliefs.” Today we vote 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. That is, elected representatives might define some national goals, and a formal measure of attainment of those goals, like Gross National Product or the Human Development Index, and oversee its measurement after the fact. (I call this measure “national welfare,” but don’t presume economic concepts dominate.) An agenda process would propose specific new

policies to consider as alternatives to status quo policies. For each new proposal, market speculators would say whether they expect adopting that policy to raise national welfare relative to the status quo.

The basic rule of government would be this: when speculative markets clearly estimate that a proposed policy would increase national welfare, that policy becomes law. As we will discuss in more detail, speculative market prices can estimate national welfare conditional on a proposed policy being adopted, and conditional on that policy not being adopted, via called-off trades in assets that pay in proportion to measured national welfare. The difference between these two prices gives an estimated policy effect.

In this paper I offer neither a formal model nor a statistical practice analysis of this proposal. I instead take an engineering approach, seeking relevant design issues and a basic design that plausibly addresses these issues well enough to be worth testing at higher levels of prototype realism. This paper will thus review democracy's info problems and speculation's info successes, outline how speculation could aid democracy, consider twenty-five objections, and present a somewhat detailed design intended to address most of these objections.

Info Failures of Democracy

How well-informed is current democratic policy? Half a century ago, empirical research on individual U.S. voters seemed to confirm the worst fears of democratic skeptics.

Few citizens paid much attention to politics ... Even on important issues such as government help with jobs, aid to education, or the stationing of American troops abroad, large proportions of the public did not know what the government was currently doing, where the opposing parties stood, or even what they themselves wanted to government to do. ... Less than 20% ... had real and stable attitudes on ... electric power and housing (Page & Shapiro, 1992).

Such ignorance continues today. For example, only 29 percent of U.S. adults can name their House representative, and only 24 percent can identify the first amendment of the U.S. constitution (Delli-Carpini & Keeter, 1997).

Yes, formal analysis often suggests that voter info problems need not be severe (Wittman, 1995). For example, it should suffice to have policies chosen by a single central leader, and to have voters evaluate that leader retrospectively based on whether local outcomes visible in their personal lives seem better than expected. In fact, however, real votes are often based on beliefs about outcomes experienced by millions of distant others (Sears & Funk, 1990), and often rely heavily on prospective policy evaluation. Furthermore, since power is divided, voters must discern which politicians are responsible for which outcomes.

It is also true that voters have many information aids. Voters collect info as a side effect of living, large elections average individual errors, and political entrepreneurs can advertise. If voters can identify well-informed sources with incentives to honestly share info, voters can defer to such sources.

Unfortunately, voters usually have little chance to be pivotal in their voting or lobbying, and so have only weak incentives to use such sources well. Yes, rational but ignorant voters could recognize their info problem, and reduce it by being more uncertain about their opinions, more

suspicion of sources, and more open to ways to mitigate their ignorance, such as politician financial incentives, source track records, or juries.

In fact, however, voters seem not just ignorant, but overconfident in their political views and sources, and disinterested in mitigating mechanisms. Yes, some say voters have stable coherent opinions (Page & Shapiro, 1992), and learn what they need via cheap “info shortcuts” (Lupia, 1998). But more plausibly, voters are “rationally irrational,” proudly seeing “one man one vote” as validating their opinions, and indulging in such pride because the personal costs of doing so are small (Caplan 2007).

A proud unwillingness to defer to expert opinion seems confirmed in consistent differences between public and expert opinions. Within all demographic groups studied, those who know more tend to be more economically conservative and socially liberal (Althaus, 2003). Relative to public opinion, expert toxicologists tend to estimate lower risks and higher benefits from using chemicals (Kraus, Malmfors, & Slovic, 1995). Expert economists worry less about foreign aid and down-sizing, and more about reduced productivity growth; these tendencies remain in the well educated, after controlling for age, income, and ideology (Caplan, 2002).

While better info sources can help rational but ignorant voters, arrogant and ignorant voters may ignore such sources, forcing politicians and policy makers to also ignore them. And in fact, public policy often seems closer to public opinion than to what relevant experts advise (Dixit, 1997; Stiglitz, 1998); governments often use expert advisors more to legitimate predetermined policies than to gain info to determine policy (Barker & Peters, 1993).

Wealth differences between nations have greatly increased over the last two centuries (Maddison, 1995), and many see this as resulting from some nations adopting what are clearly better economic policies (Olson, 1996; Sachs & Warner, 1995; Ascher, 1999). This suggests a huge info institution failure in the nations that refused to believe in and adopt the clearly better policies. Some explain this as due to irrational and proud voters (Caplan 2003).

While many factors contribute to bad policy, it seems hard to imagine nations adopting them nearly as often if most everyone knew they were bad. Thus an important contribution to bad policy is a failure to aggregate relevant info and put it in the hands of decision makers with sufficient incentives to notice and use that info.

Info Successes of Speculative Markets

While government policy may often suffer info failures, speculative markets show striking info successes.

Most markets for stocks, bonds, currency, and commodities futures are called speculative markets because they let people speculate on future prices by buying or selling today in the hope of reversing their trades later for a profit. Such opportunities to buy low, sell high occur when identical durable items are traded frequently in a market with low transaction costs. Given such opportunities, everyone is essentially invited to be paid to correct the current market price, by pushing that current price closer to the future price. Such invitations are accepted by those sure enough of their beliefs to put their money where their mouth is, and wise enough not to have lost their money in previous bets.

Betting markets are speculative markets trading assets designed let people speculate on particular matters of fact, such as which horse will win a race. Final bet asset values are defined in terms of

later official judgments about the facts in question. By construction, such assets are durable, identical, and can be created in unlimited supply.

Betting and other speculative markets have been around for centuries, and for decades academics have studied their info properties. The main robust and consistent finding is it is usually quite hard, though not impossible, to find info not yet incorporated in speculative market prices (Lo, 1997; Hausch, Lo, & Ziemba, 1994). In laboratory experiments, speculative markets usually aggregate info well, even with four ignorant traders trading \$4 over four minutes (Sunder, 1995; Forsythe & Lundholm, 1990).

There are, however, noteworthy exceptions. For example, betting markets typically overestimate low probability, or “long shot,” events, though with low transaction costs specialists usually eliminate this bias. Also, with complex trader info, “info traps” can arise where traders lack info revelation incentives (Noeth, Camerer, Plott, & Webber, 1999). Such traps can often be avoided by trading in more kinds of related assets.

Financial markets also seem to have excess long-term aggregate price variation, such as financial market bubbles (Shiller, 2000), though bubbles can be hard to distinguish from rational info responses (Jorion & Goetzmann, 2000). Risk and delay discourage specialists from correcting such errors, and irrationally-risky traders can actually get better returns, if not utility, and so are not selected out (De Long, Shleifer, Summers, & Waldmann, 1990). Even if speculative markets are distorted by irrational bubbles, however, no known info institution does better. For example, what other info institution consistently and clearly predicted that we were over-investing during the dotcom or housing bubbles? Yes individual academics or reporters so predicted, but so did individual stock investors.

The key issue is not absolute accuracy, but accuracy *relative* to other institutions, on the same topics, given similar resources. We have some data on this. In addition to lab studies, a few studies directly compare real speculative markets with other real info institutions. For example, racetrack market odds improve on the prediction of racetrack experts (Figlewski, 1979), orange juice commodity futures improve on government weather forecasts (Roll, 1984), stocks fingered the guilty firm in the Challenger crash long before the official NASA panel (Maloney & Mulherin 2003), Oscar markets beat columnist forecasts (Pennock, Giles, & Nielsen, 2001), gas demand markets beat gas demand experts (Spencer, 2004), betting markets beat Hewlett Packard official printer sale forecasts (Chen & Plott, 2002), and betting markets beat Eli Lilly official drug trial forecasts (Servan-Schreiber, 2012).

The largest known tests compare bet prices to major opinion polls on U.S. presidential elections. For example, in 709 out of 964 comparisons, bet prices were closer to the final result (Berg, Nelson & Rietz, 2008). This gain comes in part from prices being disproportionately influenced by active traders, who suffer less from cognitive biases (Forsythe, Nelson, Neumann, & Wright, 1992; Forsythe, Rietz, & Ross, 1999).

Some have noted that statistical models built from poll data predict past elections better than did betting prices, and could have been used to win bets (Erikson & Wlezien, 2008). However, if we compare apples to apples, statistical models built from bet prices predict better than statistical models built from polls (Rothschild, 2009). More important, one group’s statistical model on one topic is not a general *institution*, to which we could return time and again for answers on many disputed policy topics.

My claim is *not* that betting prices are always more accurate than other sources, but that they are a *robustly* accurate *public institution* estimating policy-relevant topics. When supported by similar resources and compared on the same topic, they are often substantially better, and only rarely substantially worse, than other info institutions with publicly visible estimates. Speculators provide a valuable service even when they just evaluate other public info institutions, and echo the most accurate of them.

My claim also has little to do with any “wisdom of crowds” (Surowiecki 2004). Speculators tend to rely more on crowds when crowds know more, and on experts when experts know more. Yes, the best track bettors have no higher I.Q. (Ceci & Liker, 1986), but speculative markets if anything still over-emphasize experts, both public and private (Figlewski, 1979; Metzger, 1985; Lichtenstein, Kaufmann, & Bhagat, 1999).

Instead, the main reasons for superior speculative market accuracy seem to be incentive and selection effects: stronger accuracy incentives tend to reduce cognitive biases (Kruglanski & Freund, 1983; Eply et. al. 2004), those who think they know more tend to trade more, and specialists are paid to eliminate any biases they can find. These attempted explanations, however, are not the main reason to believe in speculative market accuracy. The main reason is the robust and consistent empirical track record.

We want policy-related info institutions to resist manipulation, i.e., to resist attempts to influence policy via distorted participation. Speculative markets do well here because they deal well with “noise trading,” i.e., trading for reasons other than info about common asset values. When other traders can’t predict noise trading exactly, they compensate for its expected average by an opposite average trade, and compensate for its expected variation by trading more, and by working harder to find relevant info. Theory says that if trader risk-aversion is mild, and if more effort gives more info, then increased noise trading *increases* price accuracy. And in fact, the most accurate real speculative markets tend to be those with the *most* noise trading.

What do noise traders have to do with manipulators? Manipulators, who trade hoping to distort prices, *are* noise traders, since they trade for reasons other than asset value info. Thus adding manipulators to speculative markets doesn’t reduce average price accuracy. This has been verified in theory (Hanson & Oprea, 2009), in laboratory experiments (Hanson, Oprea, & Porter, 2006; Oprea, Porter, Hibbert, Hanson, & Tila, 2007), and in the field (Camerer, 1998; Wolfers & Zitzewitz, 2004).

Measuring Welfare

Since speculative markets excel at a task where democracies struggle, we might try to improve democracy by having it rely more on speculative markets. However, to let bettors tell us which policies will make a nation better off, we need a way to settle bets on how well off a nation has been.

Many say we know too little about individual preferences to be able to evaluate in any detail how an individual’s quality of life fluctuates from hour to hour (Buchanan, 1988). As we aggregate across people, time, and contexts, however, these uncertainties often cancel, allowing more accurate judgments about how well off people are on average. It seems reasonably clear, for example, that the average Ethiopian today is less satisfied with her life than the average Swede. This can hold even if we could not advise a random Ethiopian on what to eat for lunch.

The development of national income accounts, i.e., the collection of statistics such as Gross Domestic Product (GDP), has been justifiably called one of the greatest economic innovations of the twentieth century. While such measurements have error, and leave out much, they seem a sufficient basis for many policy recommendations. For example, most researchers in empirical growth economics seem willing to recommend policies that causally induce sustainably higher GDP. And frequent travelers find it hard to escape the impression that, comparing nations with large differences in GDP, most who live in high GDP nations are better off than most people who live in low GDP nations.

Furthermore, researchers spending small budgets have already explored many promising extensions to current GDP measures, such as lifespan, home production, leisure, and environmental assets (Boskin, 2000; Nordhaus, 2000). Researchers also have standard ways to weigh future welfare, risk aversion, and inequality. Direct measures of happiness have been developed (Oswald, 1997), as have entire alternative welfare measures such as the Human Development Index.

With more funding and attention from researchers and elected representatives, we could devise even more accurate and robust measures of national welfare. Whatever your opinion of current measures like GDP, you must admit that over the last century we have dramatically improved our ability to measure social outcomes, and we are continuing to make dramatic gains.

Decision Market Mechanics

Even if we can see afterwards how well off a nation has been, how can we see beforehand which policies speculators expect to make that nation better off? The trick is to use markets whose prices can be interpreted as decision-conditional estimates (Hanson, 1999).

Consider bets on the event D , i.e., the Democratic Party wins the 2016 U.S. presidential election. A “bank” (i.e., financial firm) could without risk accept \$1 in payment¹ for the pair of contingent assets, “Pays \$1 if D ” and “Pays \$1 if not D .” This transaction carries no financial risk because exactly one of this pair will be worth \$1 in the end. The expected dollar value of the asset “Pays \$1 if D ” is $\$p(D)$, i.e., the probability of the event D . So if someone is willing to buy this asset for \$0.60, we can interpret this roughly as their saying the chance Democrats will win is at least 60%. And a market price of \$0.60 can be interpreted as a consensus among traders that $p(D) \approx 60\%$.² After averaging in their minds over plausible scenarios, traders would have judged that Democrats win in about 60% of such scenarios.

Now imagine a market like the one above, where people can trade an asset like “Pays \$1 if D ” for some fraction of \$1, but where such trades are called off unless there is also the event H , i.e., Hillary Clinton is the nominee of the Democratic Party in the 2016 election. When a trade is called off, the exchange does not happen, and each side instead retains their original assets. Traders thinking about the prices they are willing to accept here should again average over plausible scenarios, but now only considering scenarios consistent with the event H , that Clinton is the nominee.

A market price of \$0.70 here can be interpreted as a consensus that about 70% of these scenarios have the Democrats winning. This would be a consensus that $p(D|H) \approx 70\%$, i.e., that the

¹ “\$1” here is really the riskless bond “Pays \$1 after the 2016 election.”

² This neglects possible differences in the value of money between the sets of states D and not D .

conditional probability of Democrats winning, given that Hillary Clinton is the Democratic nominee, is 70%.

Note that this called-off-bet approach gives us speculator estimates on the consequences of events that *never actually happen*; until speculators know an event won't happen, they can have incentives to accurately estimate its consequences.³ Note also that speculators do *not* bet that a decision will later be shown to be best. We thus never need to prove later that a particular decision actually had better outcomes than would have the options not chosen. We will never know this, and never need to know. We only ever need to know the consequences of choices that were actually made.

Speculative markets can estimate not only probabilities, but also expected values. Let x denote the Democratic party's fraction of the total vote in the presidential election. If a bank accepts \$1 in payment for the pair "Pays x " and "Pays $1-x$ ", then the fraction of \$1 required in trade for the asset "Pays x " can be interpreted as a market estimate of the expected value $E[x]$. Similarly, the fraction of the asset "Pays \$1 if H" traded for the asset "Pays x if H" can be interpreted as a market estimate of the conditional expectation $E[x|H]$. Since 1996, the Iowa Electronic Markets have actually had markets with prices describing each party's expected vote fraction conditional on various possible party nominees (Berg & Rietz, 2003).

Similarly, if national welfare W is a measure normalized to be between zero and one, then we can create speculative markets to estimate whether a new policy N is an improvement over status quo policy Q . Markets trading "Pays W if N " for some fraction of "Pays \$1 if N " give estimates of $E[W|N]$. Just as people considering what fraction of \$1 to pay for "Pays W " will try to average W values across all plausible scenarios, people considering what fraction of "Pays \$1 if N " to pay for "Pays W if N " will average only over plausible scenarios consistent with the event N .

Standard decision theory says to take the action associated with the highest action-conditional expected utility. When the decision market estimate of $E[W|N]$ is clearly greater than its estimate of $E[W|Q]$, speculators are saying they expect more national welfare, given this new policy. In this case, standard decision theory recommends that, if you were in a position of power to make the choice, you should choose this new policy N if what you wanted was national welfare, and if the decision market prices matched your beliefs.

To illustrate this general mechanism, let us consider two concrete applications, before we consider more general proposals.

Corporate Governance Example

A public corporation is in effect a small democratic government. Its charter is its constitution, its stock owners are its citizens, and those citizens' welfare, at least as it relates to firm actions, is well summarized by the firm's stock price. The charter authorizes citizens to vote in proportion to their holdings for a board of directors, who choose and compensate a CEO, who runs the company. While CEO compensation usually gives CEOs strong incentives to maximize stock value, directors usually have weak incentives, and are often chosen by CEOs. As a result, boards of directors tend to wait too long before firing bad CEOs.

³ Conditional estimates also come from trading "Pays \$1 if D and H " for some fraction of "Pays \$1 if D ."

Imagine WidgetCo amends its corporate charter to make its CEO's tenure depend on certain market prices. WidgetCo hires a finance firm to maintain two new stock markets, trading WidgetCo stock conditional on dumping, and on keeping, its CEO in the current quarter.

In an ordinary stock market people trade cash for stock, and the ratio of cash to stock is the market price; bid and ask prices are prices where people offer to buy and sell stock. In a dump (or keep) market, in contrast, people would still agree to trade cash for stocks, but these agreements to trade would be called off (i.e., made null and void) if the CEO were not dumped (or not kept) by the end of the current quarter. The WidgetCo charter might say that if the dump price were clearly higher than the keep price over the last week of the quarter (e.g., dump bid price above keep ask price 90% of the time), the board must dump him at quarter's end.

If Disgruntled Jones, a junior WidgetCo employee, could trade anonymously, he could avoid retribution for making buy offers in the dump market. The CEO might accept sell dump offers to try to make himself look good, but the WidgetCo charter might say his trades and offers must identify him, to help observers correct for this distortion. If enough traders agreed with Disgruntled Jones to dump the CEO, and if WidgetCo prospered because of this, traders like Jones would be directly rewarded for their helpful advice.

Monetary Policy Example

Once firms succeeded with governance decision markets, federal agencies might try them. For example, most central banks periodically change short-term interest rates, reserve requirements, currency conversion rates, and other monetary policy parameters. They do this to achieve goals of smoothing out business cycles, limiting inflation, and promoting high average growth rates.

While monetary specialists sometimes disagree about relative weights for these goals and outcomes, their main disagreement is on how to achieve them. For example, people argue for independent expert judgment, for political accountability, and for fixed pre-announced policy rules (Taylor, 1999). This makes monetary policy a good candidate for decision markets.

To get market advice, the central bank could declare an explicit function describing good outcomes. This outcome function might, for example, be a weighted sum over each of the next hundred quarters of that quarter's GDP growth rate, unemployment rate, inflation rate, and the absolute value of each parameter's quarterly change.

The outcome would be some number, and each outcome asset could later be exchanged for that number of a base asset (perhaps inflation-indexed bonds). These outcome assets could not pay less than zero nor more than some large maximum value. Anyone could trade this maximum amount of base assets for one outcome asset, plus one asset that paid the maximum minus the outcome value.

To determine whether interest rates should change at week's end, there could be three markets, up, down and same. Up markets would trade cash for the outcome asset, trades called off unless interest rates were raised that week. Down and same market trades would be similar, called off unless interest rates were lowered or left the same respectively. If the up or the down market had a consistently higher outcome price over most of that week relative to the same market, interest rates would be raised or lowered accordingly.

In this system, info about the economy that is now given privately to monetary officials would instead be made public.

The Engineering of Institutions

This general idea of using speculative markets more widely to aggregate info has been discussed before, both in academic journals (Hirshleifer, 1971; Leamer, 1986) and in science fiction (Brunner, 1975; Stiegler, 1999), as has the more specific application of aggregating policy-relevant info (Zeckhauser & Viscusi, 1990; Hanson, 1995, 1999, 2006; Berg & Rietz, 2003).

A scientist (or at least a caricature of one) insists on saying “I do not know” about a theory until it has robust empirical support, or has clear theoretical support from a well-supported theory. A scientist thus bases policy recommendations only on relatively direct data, or on well-supported theory, and so stays quiet about radical new forms of governance, which cannot possibly have direct empirical support, and which are too complex for direct theoretical application.

An engineer, on the other hand, is happy to work on a concept with a five percent chance of success, if success pays thirty times the cost of trying. An engineer first uses analogies and theory-informed intuitions to consider a wide range of design issues, and then seeks modular design concepts that seem, relative to the status quo, likely to reduce more problems than they worsen. Promising design concepts progress to increasingly realistic proof-of-concept prototypes, from simulations, to wind tunnel models, to field tests. While scientists have little use for prototypes, prototypes make engineers’ worlds go round.

While academic social studies are now dominated by a scientific style, this paper takes an engineering style to considering a possible new form of governance. The purpose of this paper is not to induce high confidence in this concept, but to raise reader confidence to a level justifying further exploration via more realistic prototypes.

Simple Approaches

Let us now begin to consider this paper’s core proposal, to use speculators’ info successes to mitigate democracy’s info failures via decision markets that estimate national welfare given government policies.

The simplest way to do this is via advisory decision markets; let anyone pay to create them, and let anyone consult them for info. This would likely improve on the current situation, where such markets are mostly illegal. But as discussed above, voters often have insufficient incentives to notice and use better info sources. In addition, without a subsidy, too few may volunteer to pay to create such better info sources.

Also, advisory decision market prices can be seriously distorted when decision makers might know things that market speculators do not. In such cases, the fact that a certain decision is made can indicate hidden info held by decision makers. Market estimates of outcomes conditional on a decision then become estimates of outcomes given this hidden info, instead of estimates of the effect of the decision on outcomes. For example, even if a firm was better off firing its CEO, the fact that the board did not fire him might indicate that they knew something very positive about him. Thus speculators might set prices to estimate a higher expected value of the firm’s stock conditional on keeping the CEO, relative to firing him, even if they believed that firing would increase stock values more.

To avoid such decision selection biases, it is sufficient either to have informed insiders trade in the markets, so that prices reflect what decision makers know, or to put market estimates directly

in control of decisions, rather than being advisory. In either case it is important to specify a short time duration during which the proposal can be chosen. Otherwise speculators will have to wonder what they might know in the future than they do not know now.

Another simple way to make use of decision markets that estimate national welfare given government policies is to have decision markets veto legislative votes. In addition to their usual tasks, legislatures would also have to pass bills to create and manage an official measure of national welfare. On ordinary topics, a bill passed by the legislature would not become law if the market estimate of national welfare given the bill becoming law were clearly lower than the market estimate of national welfare given the status quo.

This approach might also improve on the current situation, by preventing large bad changes. However, a vetoed bill could probably be approved anyway, if broken into many smaller changes, none of which was clearly enough bad to trigger the veto. Also, legislators who saw the veto mainly as a formal obstacle to their plans might get into the habit of constantly redefining national welfare as needed to ensure market approval of each new favored bill, undermining the entire approach.

Both of the simple approaches considered so far, purely advisory markets and markets that veto legislative bills, are admirably incremental and modular; their minimal changes to the status quo make it easier to analyze their consequences. But in both cases there is a serious risk that the decision markets will be ignored. So in order to better explore the possibilities for decision markets to actually improve our policy info institutions, let us consider a reference proposal that is less incremental, but more likely to make a difference if adopted.

A Reference Proposal

Let us now consider the more extreme option of modifying our basic form of government to ensure sufficient incentives to notice and use decision market advice. The following reference proposal is intended to be as simple as possible while still ensuring substantial use of decision market advice, and plausibly addressing most of the concerns raised in the objections section of this paper. While it could apply at a city, county, state, nation, or national alliance level, for concreteness the following discussion focuses on nations.

Rule-based governance systems need ways to propose new rules, to approve or reject proposed rules, and to ensure that approved rules are applied and meet constitutional constraints. Let us now consider changing how we propose and approve rules, but not how we check or implement them.

That is, let us retain courts to adjudicate disputes on law application and constitutionality. Let us also retain legislatures, but greatly limit their scope. Instead of being allowed to pass bills on most any subject, legislatures can now pass bills only on their internal procedures, on the process by which legislators are chosen, and on the creation and management of an official measure of *national welfare*, i.e., a formal after-the-fact measure of how well we have attained national goals.

Welfare-changing bills are implemented after a standard notice period (e.g., a week). Courts should invalidate any welfare definition change that seems too directly an attempt to favor a specific policy whose main rationale comes from other measurable outcomes. For example, no fair putting miles of roads in the welfare definition, to favor a road-building proposal, if the reason for roads is to reduce travel times, which is already measured.

The agencies charged with measuring welfare would ideally have enough redundancy and monitoring to deter most corruption in measurement, and be independent enough from other social institutions so that changes elsewhere will not substantially and systematically change the welfare values agencies declare.

Some *base* asset, such as a bond or a stock index fund, is defined. It must be widely available, hold a relatively stable value, and offer a competitive rate of return. A *welfare* asset would also be created which, relative to this base asset, pays later in proportion to national welfare.

At any one time there would be a set of status quo laws. Such laws could be nested, with more fundamental policies authorizing less fundamental policies, similar to the way that a law today authorizing the creation of a government agency implicitly authorizes whatever regulatory rules that agency creates. At any one time there is also an official process by which authorized agencies estimate welfare. An agenda process, discussed below, offers new rules to consider as proposed laws.

Each valid proposal says, directly or indirectly, what people must do to comply with it, what happens if they do not comply, what previous laws it invalidates, and how and by whom ambiguities are to be clarified. Courts need only clarify any remaining ambiguities, and say if laws are valid. Valid laws do not violate constitutional guarantees, and may not substantially and systematically influence how legislators are chosen, or how welfare is defined or measured. Each valid proposal specifies a short decision period (e.g., an hour) during which it will be decided.

For each valid proposal, two pairs of markets let people make conditional trades of welfare assets for base assets. For each pair, in one market trades are conditional on adopting the proposal, while in the other trades are conditional on the opposite. The difference between the prices in a pair estimates how much the proposal changes national welfare. The two pairs differ only in how welfare is defined; in one pair asset values are tied to national welfare as it is currently defined, while in the other pair asset values are tied to national welfare as it will be defined a fixed period (e.g., a year) later.

The main rule of government is: a valid proposal is adopted if decision-period prices estimate that it improves currently-defined welfare, and if it is not vetoed by decision-period prices estimating that it hurts future-defined welfare. That is, an official agenda-process-approved proposal becomes law if, over the proposal's decision period, its current-welfare adoption-conditional price is clearly higher than its current-welfare opposite-conditional price, while its future-welfare adoption-conditional price is not clearly lower than its future-welfare opposite-conditional price.

A *clear* price difference has a precise definition (e.g., at least twice the bid-ask spread for 90% of the trading volume), and averages over time, trades, and valid marketplaces. Valid marketplaces trade valid assets, with trade fees or subsidies within limits, and with prices based on offers open to all valid traders (e.g., anyone) who spend resources within limits (e.g., \$2,000 each per year).

An *agenda process* picks proposed rules for decision markets and courts to officially consider. A good agenda process limits the number of proposals that evaluators must seriously consider, and gives them sufficient time for such evaluation. A good process gives incentives to propose rules that are likely to be approved, and to not propose rules that are quite unlikely to be approved. A good process also has many competing channels whereby advocates might get a proposal onto the agenda; there should be no single place that can reliably stop a proposal.

Many simple approaches can roughly meet these requirements. For example, during each year each of up to 730 legislative representatives could have an assigned half-day decision-period, to which they could assign any proposal if they gave a week's notice.

To this economist's eyes, however, the following approach seems more effective. Let there be a future calendar of non-overlapping decision period slots (e.g., one a day), and hold an auction to decide which proposal will be decided in that slot. The auction is held a waiting period (e.g., a week) before the slot. The sponsor who pays the auction price is paid a reward (e.g., \$100 million, or 1% of the bill's estimated national welfare increase) if that proposal is approved by markets and courts to become law.

Welfare Possibilities

While national welfare could be anything the legislatures choose, it helps to see that reasonable options exist to choose from. For example, a reasonable initial definition of national welfare could augment current measures of national consumption or product (i.e., G.D.P.) with simple measures of health, leisure, happiness, and the environment.

To express risk-aversion and sensitivity to inequality between identified subgroups, one could break the population into subgroups by geography, ethnicity, age, gender, and income, and then define monthly welfare to be a weighted average of the square root of each group's augmented consumption. Total national welfare could then be a discounted average (over the indefinite future) of monthly welfare.

By varying subgroup weights one would induce various types of affirmative action or discrimination policies. Consequences for foreigners and animals could be included, and nations might agree by treaty to explicitly give welfare weight to other nations' welfare. Some especially difficult choices surround how to treat changes in national population due to births, immigration, and emigration.

The agencies responsible for measuring welfare must eventually publish welfare values for each time period according to all past official welfare definitions that gave weight to this period. A recursive approach avoids an explosive burden here: welfare today could be a weighted average of a specific monthly welfare over the next twelve months, plus some weight times however welfare will be defined a year from now, applied to all months from then on.

Note that we can give speculators incentives to estimate the welfare over an arbitrary number of future time periods, without there ever being a point in time where we know for sure what is the value of total welfare.

Respect for rights and freedoms might be encouraged if the government collected statistics on the number and types of violations of basic rights and freedoms, and if the welfare definition gave a negative weight to such outcomes.

Objections

Many ideas that seem promising at first glance seem less so after one considers more details. Let us therefore, in engineering fashion, try to identify and briefly discuss twenty-five design issues, each expressed as an objection to this overall approach.

1. Maybe Poor Info Is Not Democracy's Major Failing

What if democracy's biggest problem is coordination or commitment failure, or preventing the government from being overthrown? If so, this new form of government might still be attractive if it did not make this bigger problem much worse, and if it were compatible with whatever reforms could reduce this bigger problem.

2. We Should Appeal To Higher Motives Than Money

If we could reliably pay off bets in higher forms of value, we might want to do so. We want to induce people to tell us what they know, and as a reward we need to offer them something that we can reliably produce, distribute in controlled amounts, and know that most people value. Only money now satisfies these requirements.

3. A Market-Based Approach Is Biased In Favor of Markets

Even though the mechanism of betting on beliefs is market-based, it need not favor markets. It could in principle approve anything from an extreme socialism to an extreme minarchy, depending on what legislators say they want, and on what speculators think will achieve those goals. We do need there to be some community of speculators who have assets they are free to bet, and we need some extra benefits to accrue to those who win their bets. These benefits need not be large, however, nor need the community of speculators be large. There is an old Marxist saying, a capitalist will sell you the rope to hang him with. Market speculators who were rich and selfish could easily approve policies that greatly hurt the rich.

4. People Might Reveal Less And Lie More

Perhaps many people now often talk about their politically relevant info because they feel it is their civic duty. If related betting markets became available, and if civic duty were fragile, then such people might be tempted to keep quiet or lie about their info, to benefit from related bets.

In speculative markets, the incentive to keep quiet about one's info is strongest before one trades on it. As one makes more trades, one approaches budget bounds that limit further trades. One is then increasingly tempted to publicly reveal one's info, to persuade other speculators and move the price favorably. One could then reverse prior trades for a profit, generating cash to allow this whole process to be repeated on other info. Thus info should be revealed via *both* bets and talk.

The option to bet on politics increases incentives to lie, but today lies are limited much more by a limited willingness to believe lies than by a limited willingness to tell lies. Those seeking to avoid lies have long been well advised to carefully consider the track records and incentives of their sources. With decision markets, the public will have a new powerful way to avoid lies: ignore the talk and just look at market prices.

5. Expressive Voting Could Still Cause Problems

Voters often seem to vote expressively, i.e., to care less about the policies their votes may influence than about how their votes will make them look and feel (Brennan & Lomasky, 1993). Voters may, for example, want to show they are tough or that they care, and they may want to have people they like represent them. These may be worthwhile goals, but policy choices can suffer as a result. By limiting democracy primarily to values, we would presumably focus voters more on expressing their values, rather than their beliefs. It is not clear that this would make expressive voting more (or less) of a problem overall.

6. The Rich Would Get More Influence

Betting markets are not opinion polls where the rich get more votes. The rich have more potential votes, but they lose those votes if they use them without having proportionally more info. And the rich do not tend to throw their money away easily; those who do, do not stay rich long. Also, those rich in money but poor in info will have incentives to seek out and team with those rich in info but poor in money.

Note that there is also the option to limit how much money each person can use to trade on policy proposals, though this might well reduce the accuracy of market estimates.

Today, civil servants have far more info influence on policy than do most citizens. This inequality is mostly seen as good, as it allows more total info via specialization. Unequal info influence in markets can also be good, for the same reason.

7. One Profits Little By Supporting Unlikely Proposals

Imagine you liked a proposal, but few agreed with you. Because your opponents have consistently held down the market estimate of national welfare given your proposal, your proposal has never been tried, and so you have never been rewarded, and they have never been punished, for your being right.

All political systems have similar problems. In this system, the insights that tell you that your untried proposal is promising are probably also useful for predicting the consequences of actually adopted proposals. If you prove yourself by winning bets on those topics, you will gain money and credibility to better argue for your untried proposal.

8. Markets Might Be Too Thin To Make Good Estimates

If decision markets were very thin or noisy, the relevant price difference would usually not be clear enough, and so the status quo might not change when it should. A market known to influence important policies, however, simply cannot be very thin (Hanson & Oprea, 2009).

After all, if a very thin market were to influence policy, interested parties could pay very little to move the price and favorably influence policy. If interested parties on both sides had similar budgets, their combined trades would thicken the market. Alternatively, if one side could spend much more, then speculators who knew this could make easy profits by trading against that better funded side, again thickening the market. Either way, the market would not be thin.

9. One Rich Fool Could Do Great Damage

Even Bill Gates, the richest person alive today, has only a tiny fraction of the wealth compared to all market speculators. So if most speculators were confident that he was wrong on a market estimate, even Bill Gates could not substantially change that estimate. If Gates pushed the price one way, those speculators would push it back. More likely, however, speculators would allow Gates to move the price some because they reasonably suspect that he had access to relevant info. Given his success, this would not be unreasonable.

10. People Could Buy Policy Via Trades

If market prices determine policy, and if trades predictably change prices, then interested parties could make trades that predictably influence policy. For example, someone with an unexpectedly strong desire to raise the market estimate of a particular stadium's profitability could succeed in doing so.

However, market speculators would have already corrected for an average expected desire to bias the price, and have increased their willingness to collect info and trade to match expected variation in such hidden desires. So every case where a manipulator raises the price too high must be matched by an equally likely case where the price is made too low, even if manipulators wanted on average to raise it. As discussed above, the more manipulation is expected, the *more* accurate are prices.

This is also why one needn't worry about a nation's enemies hurting it via adding misleading trades. On average, adding noisy trades *increases* price accuracy.

11. If Such Markets Are So Useful, Why Don't They Exist Now?

Such markets are mostly illegal. The history of financial regulation is basically: everything legal today was once banned as gambling (or usury) until an exception was granted for a newly legitimized higher purpose. A thirteenth century decree by Pope Gregory IX prohibited maritime insurance as usury. The 1570 Code of the Low Countries outlawed life insurance as gambling (Brenner & Brenner, 1990). In response to speculation in the South Sea Bubble, in 1720 Britain basically banned the formation of joint-stock companies (Kindleberger, 1984). And futures markets were banned as gambling in the late nineteenth century U.S. (Brenner & Brenner, 1990). We might thus hope to someday legitimate and legalize markets whose main function is to aggregate info on important policies (Bell, 1997).

12. People Could Do Harm to Win Bets

People are sometimes murdered for their life insurance money. One might fear speculators similarly selling a firm's stock short before directly damaging its capital or product, in order to profit from a stock price fall. But this almost never happens, apparently because it is very hard to combine enough financial capital with a willingness and ability to sabotage.

It is far harder to substantially damage national welfare. Furthermore, damaging actions should only favor market evaluations of a particular proposal if speculators expect that such damaging actions are more likely if the proposal is not implemented. Thus in order to bias a decision via such sabotage, one has to commit to substantially harming national welfare if the decision goes one way, but not if the decision goes another way, and either have deep financial pockets to bet on this, or credibly convince deep-pocket speculators that you had so committed. This seems an unlikely scenario.

13. Welfare Measurement Might Be Corrupted

The straightforward way to profit from an ability to change a measurement would be to artificially raise or lower that measurement, after one had bet that this change would happen. In order to bias a decision, however, one has to also commit to substantially changing an important contribution to measured national welfare if a decision went one way, but not if a decision went another way, and either have very deep financial pockets or convince deep-pocket speculators that you had so committed, without revealing your corruption to enforcers. This seems an unlikely scenario.

14. A Policy Might Influence How Welfare Is Measured

If speculators expected a policy to raise estimates of national welfare by influencing how welfare is measured, rather than by influencing real welfare, they still would approve that policy. It is therefore important that welfare measurement be a relatively independent process. Markets should therefore not be allowed to approve proposals that substantially change how welfare is measured, or the political process that leads to changes in the welfare definition.

15. Policies Might Be Encoded In The Welfare Definition

Elected representatives could in principle directly get whatever decisions they wanted by encoding them in welfare definition. For example, if they wanted a certain road built, they might put a term in the welfare definition that takes on a large positive value if the road is built as specified, and zero otherwise. Speculators would then have to agree that building the road would raise expected national welfare.

To avoid this, courts might be empowered to enforce a rule against such rigged welfare definitions, just as courts now enforce the rule that tax laws cannot too directly single out a particular person for higher taxes. More generally we might hope the public would consider such welfare terms as scandalous. At worst, this new form of government would just revert to our something like familiar democracy.

16. Most Decisions Can Not Detectably Impact Welfare

Speculative markets do not create infinitely precise prices, since many random contingencies influence exactly who trades what when. To avoid letting such contingencies overly influence policy, we can require markets to clearly favor a proposal over the status quo, via consistently higher prices over a substantial period. But then how could small changes ever be approved?

Many small changes added together can have large impact. For example, instead of proposing to build a particular stadium, one could propose a general policy saying how all stadium decisions will be made. Instead of proposing one road, one can suggest a general road policy.

A useful general approach is to let laws recurse, via derivative decision market regimes with more focused welfare measures. For example, a general stadium policy might be approved at a basic level which says to approve any proposed stadium if markets estimate that it would increase some measure of regional welfare, stadium profitability, or an ex post cost-benefit calculation (Abramowicz, 2004, 2008). A stadium that would not noticeably affect national welfare may noticeably affect regional welfare.

17. The Military Needs Secrecy, But Markets Are Public

This form of government allows public decisions, but does not require them. The general rule of government could be used to approve a regime in which military decisions are made in private. That private regime could be private markets, where only certain people could see prices and trade, or it could be any other institution that speculators expect to produce good military policy. We would still allow ordinary speculators to override general military policy in any specific case where they estimate that doing so would substantially raise national welfare. In weighing the gain of such overrides, speculators would take into account the fact that such overrides would be public, and therefore known to potential enemies.

18. How Can We Set All Of These Parameters Well?

There are many detail questions about how exactly to implement this proposal. How can we ensure that such parameters would be set well?

Legal recursion could robustly address these detail questions. The most basic rule of government could be very conservative, only approving changes that get very clear support from very trusted participants. Then within that framework, people could propose as policies particular less conservative regimes. So, for example, if the basic rule required that price differences must last a year to be clear, under that rule speculators could approve a policy which declared approval for all policies chosen by a certain new regime where price differences need only last one week. Once a year had passed, so the basic regime could approve this new general policy, from then on only a week would be needed to approve new policies. And later if a week seemed too short, a new bill at the basic level might cancel the one week regime and replace it with a two week regime. When such questions have a clearly best way, different fundamental conservative rules might result in nearly the same in-practice rules.

19. This Ignores The Cost of Considering Proposals

If one always adopts proposals that would raise welfare conditional on adopting them, one ignores the cost of evaluating proposals. There are real costs for people to continually pay attention to new proposals, watching for ones that might threaten their interests. But when speculators estimate the benefits of a proposal conditional on adopting it, versus not adopting it, they ignore this cost of evaluating the proposal, since this cost will have already been incurred either way. If it were easy for anyone to create a market that proposed a policy change, more good changes might result, but we might lose overall from paying more attention to the proposal process.

While this argues for limits on the rate at which proposals are considered, it does not argue for any particular way to choose proposals. As discussed above, many reasonable approaches seem feasible.

20. You Might Not Catch Buggy Decisions Quickly Enough

If a bad decision were made due to bad info, then market decision advice should be reversed the moment speculators became aware of this fact, as they would have authority to reverse the bad decision. But what if a bad decision were instead due to a bug in the welfare function, i.e., an unintended consequence of oversimplifying some aspect of our values? For example, invading and enslaving a neighboring nation might be good for currently-defined welfare, but bad if a more realistic weight were placed on the welfare of neighbors.

To fix this kind of problem, the democratic part of this form of government would have to vote to change the welfare function. But this process might be too slow to avoid harm from bad decisions that are expensive to undo; a war might be over before welfare was redefined to prevent it.

Since legislatures today can also make mistakes that they are slow to correct, it isn't clear that this problem would be any worse than today. Even so, it could be addressed via the future-welfare market vetoes included in the reference proposal above. In it, even if markets clearly estimate that a proposal will increase welfare as it is defined today, that proposal is vetoed if other markets clearly estimate a reduction in welfare as it will be defined later, say in one year. So if someone

then spotted a bug, they could bet that elected representatives would agree that it is a bug and fix it within one year. If speculators agreed, the proposal would not be implemented.

21. This Is Too Big A Change All At Once

This form of government sets up a rule for changing the status quo, but allows for any status quo and any rate of change. Thus a nation might keep its current political system and modify its constitution only to allow policy to change via speculative markets. Once a bill was passed declaring the first definition of welfare and the agencies responsible for measuring it, policy would change only as fast as speculators deemed appropriate for achieving welfare as so defined.

22. How Could One Number Encode All Our Values?

A national welfare function could be arbitrarily complex, and could refer to not just to measured economic, business, entertainment, or environmental quantities, but also to any expressed opinions. For example, to reduce the sense of shame we might feel after a war, national welfare could refer to broad surveys on shame taken after a war, or to more detailed evaluations by random juries after the war.

Legislatures could build a welfare measure incrementally. They could start with say GDP, and then entertain proposals to change this measure. Those in favor of a change could point to choices, real or hypothetical, where the new definition seems to choose better than the current definition. Those opposed would then point to choices where the old definition seems to choose better than the new. Representatives would then decide which errors seemed most tolerable.⁴

23. Easy to Measure Values Would Be Over-Valued

A teacher who is rewarded only on the basis of student test scores may neglect aspects of teaching that do not much influence test scores. If other outcomes of teaching are important and hard to measure, it can be better to pay teachers a flat salary than to pay them only for outcomes that can be easily measured (Holmstrom & Milgrom, 1991).

This new form of government could similarly cause harm if important things we valued did not contribute to national welfare as formally measured. And some things we value are certainly much harder to measure than others. The similarity with teachers ends there, however. It is risk-aversion, together with exerting hidden effort, that makes us reluctant to pay teachers on noisy measures of the real outcomes of interest. Market speculators, in contrast, are to a good approximation risk-neutral, and exert no relevant hidden effort when making estimates.

There is no good reason for a national welfare measure not to include every factor that contributes substantially to real welfare, even if the best available measures of those factors are rather noisy. For example, we might care about the beauty and comfort of our city centers beyond how such things contribute to city land values or tourism revenue. So we might create many independent commissions to rank the beauty and comfort of city centers around the world, and include a median commission ranking in a city's welfare measure.

24. You Can Not Pay Off Bets If Earth Is Destroyed

⁴ Another approach (suggested by Peter McCluskey) would be to use as the official welfare function the function that speculators estimated would give the smallest error rate in predicting later jury choices on random hypothetical test choices between national outcomes.

Any betting market implicitly gives estimates that are conditional on that market continuing to function and pay off bets. Betting markets cannot therefore directly estimate how a proposed policy would affect the chances all humans are destroyed, or any other event that would prevent paying off the bets. We do care, however, about how our policies might affect the chance of such events.

If this bias is big enough to bother with, we might augment national welfare to include official-panel-made estimates of the chances we have avoided so far of various types of destruction. Or we might more directly include in national welfare the average ticket prices for refuges, places where people might hope to survive an extreme disaster (Hanson, 2008).

25. Risk Premia Distort Estimates

Average investment returns today vary with a few dimensions that are said to pay “risk premia.” While in general asset prices should equal an expected value of future asset prices, averaging over different possible states of the world, states in which future returns are high are in effect given a lower weight than their probability would suggest. Ideally the definition of national welfare (and the base asset) would try to correct for this effect.

Conclusion

This paper has explored the idea of using the successes of speculative markets as info aggregation mechanisms to help mitigate the failures of democracies as info aggregation mechanisms. Under a new form of governance, we could formally defer to betting markets on matters of fact, while retaining representative democracy on matters of value. That is, we could vote on values but bet on beliefs. This paper has taken an engineering-style approach to this concept. After reviewing democratic failures, speculative market successes, and the measurement of welfare, this paper has considered twenty-five design issues with this concept, and sketched a more specific proposal to deal with many of these issues.

If, after this examination, the concept passes the low engineering threshold of promising, then the next appropriate engineering-style step would be to test simple prototypes in simple test environments. Mathematical and simulation models could be constructed, and laboratory experiments could compare simple versions of this concept to a more familiar democracy in simple info environments. Laboratory successes might then prompt larger trials, such as for key decisions in a firm or other organization. Eventually, some government agency might try agency-level decision-making, such as with monetary policy.

Even later, some nation might try to change its form of government. This would probably be a relatively peripheral nation, as such nations tend more to be the source of democratic innovations (Markoff 1999). Perhaps, at such a moment, such a nation might even consider what betting markets have to say about how this change would impact its future GDP and other available welfare measures.

References

Abramowicz, M. (2004) Info markets, administrative decisionmaking, and predictive cost-benefit analysis. *University of Chicago Law Review*, 71(3):933-1020.

- Abramowicz, M. (2008) *Predictocracy: Market Mechanisms for Public and Private Decision Making*, Yale University Press.
- Althaus, Scott. (2003) *Collective Preferences in Democratic Politics: Opinion Surveys and the Will of the People*. Cambridge: Cambridge University Press.
- Ascher, W. (1999) *Why Governments Waste Natural Resources: Policy Failures in Developing Countries*. Johns Hopkins University Press, Baltimore.
- Barker, A., & Peters, B. G. (1993) *The Politics of Expert Advice*. University of Pittsburgh Press, Pittsburgh.
- Bell, T. (1997) Idea futures: Making the marketplace of ideas work. Tech. rep., Chapman University School of Law.
- Berg, J., Nelson, F., Rietz, T. (2008) Prediction Market Accuracy in the Long Run. *International Journal of Forecasting*, 24:283-298, April-June.
- Berg, J. E., & Rietz, T. A. (2003) Prediction markets as decision support systems. *Info Systems Frontiers*, 5(1):79-93.
- Boskin, M. J. (2000) Economic measurement: Progress and challenges. *American Economic Review*, 90(2):247-252.
- Brennan, G., & Lomasky, L. (1993) *Democracy and Decision: The Pure Theory of Electoral Preference*. Cambridge University Press, Cambridge.
- Brenner, R., & Brenner, G. A. (1990) *Gambling and Speculation*. Cambridge University Press, Cambridge.
- Brunner, J. (1975) *The Shockwave Rider*. Harper & Row, New York.
- Buchanan, J. M. (1988) Contractarian political economy and constitutional interpretation. *American Economic Review*, 78(2):135-139.
- Camerer, C. (1998) Can asset markets be manipulated? A field experiment with racetrack betting. *Journal of Political Economy*, 106:457-482.
- Caplan, Bryan. (2007) *The Myth of the Rational Voter: Why Democracies Choose Bad Policies*. Princeton, NJ: Princeton University Press.
- Caplan, B. (2002) Systematically biased beliefs about economics: Robust evidence of judgmental anomalies from the survey of Americans and economists on the economy. *Economic Journal*, 112:433-458, April.
- Caplan, B. (2003) The idea trap: the political economy of growth divergence. *European Journal of Political Economy*, 19(2):183-203, June.
- Ceci, S. J., & Liker, J. K. (1986) A day at the races: a study of IQ, expertise, and cognitive complexity. *Journal of Experimental Psychology: General*, 115:255-266.

- Chen, K.-Y., & Plott, C. R. (2002) Info aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Tech. rep. 1131, California Institute of Technology.
- De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990) Noise trader risk in financial markets. *Journal of Political Economy*, 98(4):703-738.
- Delli-Carpini, M. X., & Keeter, S. (1997) *What Americans Know about Politics and Why It Matters*. Yale University Press, New Haven.
- Dixit, A. K. (1997) *The Making of Economic Policy: A Transaction-Cost Perspective*. MIT Press, Cambridge, Massachusetts.
- Epley, Nicholas; Keysar, Boaz; Van Boven, Leaf; Gilovich, Thomas (2004) Perspective Taking as Egocentric Anchoring and Adjustment, *Journal of Personality and Social Psychology*. 87(3):327-339, September.
- Erikson, R. Wlezien, C. (2008) Are Political Markets Really Superior to Polls as Election Predictors? *Public Opinion Quarterly*, 72(2):190-215.
- Figlewski, S. (1979) Subjective information and market efficiency in a betting market. *Journal of Political Economy*, 87(1):75-88.
- Forsythe, R., & Lundholm, R. (1990) Information aggregation in an experimental market. *Econometrica*, 58(2):309-347.
- Forsythe, R., Nelson, F., Neumann, G. R., & Wright, J. (1992). Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142-1161.
- Forsythe, R., Rietz, T. A., & Ross, T. W. (1999) Wishes, expectations and actions: A survey on price formation in election stock markets. *Journal of Economic Behavior and Organization*, 39(1):83-110.
- Hanson, R. (1995) Could gambling save science? Encouraging an honest consensus. *Social Epistemology*, 9(1):3-33.
- Hanson, R. (1999) Decision markets. *IEEE Intelligent Systems*, 14(3):161-169.
- Hanson, R. (2006) Decision markets for policy advice. In Patashnik, E., & Gerber, A. (Eds.), *Promoting the General Welfare: American Democracy and the Political Economy of Government Performance*, pp. 151-173. Brookings Institution Press.
- Hanson, R., Oprea, R., & Porter, D. (2006) Info aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization*, 60(4):449-459.
- Hanson, R. (2008) Refuge Markets. Overcoming Bias blog post, July 25, 2008.
<http://www.overcomingbias.com/2008/07/refuge-markets.html>
- Hanson, R., & Oprea, R. (2009) Manipulators increase info market accuracy. *Economica*, 76(302):304-314, April.

- Hausch, D. B., Lo, V. S., & Ziemba, W. T. (1994) *Efficiency of Racetrack Betting Markets*. Academic Press, San Diego.
- Hirshleifer, J. (1971) The private and social value of info and the reward to inventive activity. *American Economics Review*, 61(4):561-574.
- Holmstrom, B., & Milgrom, P. (1991) Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization*, 7(0):24-52.
- Jorion, P., & Goetzmann, W. N. (2000) A century of global stock returns. Tech. rep. 7565, NBER.
- Kindleberger, C. P. (1984). *A Financial History of Western Europe*. George, Allen, & Unwin, London.
- Kraus, N., Malmfors, T., & Slovic, P. (1995) Intuitive toxicology: Expert and lay judgments of chemical risks. *Risk Analysis*, 12(2):215-232.
- Kruglanski, A. W., & Freund, T. (1983) The freezing and unfreezing of lay-inferences: Effects on impressional primacy, ethnic stereotyping, and numerical anchoring. *Journal of Experimental Social Psychology*, 19:448-468.
- Leamer, E. E. (1986) Bid-ask spreads for subjective probabilities. In Goel, P., & Zellner, A. (Eds.), *Bayesian Inference and Decision Techniques*, Elsevier Science Publications, pp. 217-232.
- Lichtenstein, D. R., Kaufmann, P. J., & Bhagat, S. (1999) Why consumers choose managed mutual funds over index funds: Hypotheses from consumer behavior. *Journal of Consumer Affairs*, 33(1):187-205.
- Lo, A. W. (1997) *Market efficiency: Stock market behaviour in theory and practice*. Elgar, Lyme.
- Lupia, A. (1998) *The Democratic Dilemma: Can citizens learn what they need to know?* Cambridge University Press.
- Maddison, A. (1995) *Monitoring the World Economy 1820-1992*. OECD, Paris.
- Maloney, M. and Mulherin, J. H., "The Complexity of Price Discovery in an Efficient Market: The Stock Market Reaction to the Challenger Crash," *Journal of Corporate Finance*, 9(4):453-79.
- Markoff, J. (1999) Where and When Was Democracy Invented? *Comparative Studies in Society and History*, 41(4):660-690, October.
- Metzger, M. A. (1985). Biases in betting: An application of laboratory findings. *Psychological Reports*, 56:883-888.
- Noeth, M., Camerer, C. F., Plott, C. R., & Webber, M. (1999). Info aggregation in experimental asset markets: Traps and misaligned beliefs. Tech. rep. 1060, California Institute of Technology.
- Nordhaus, W. D. (2000). New directions in national economic accounting. *American Economic Review*, 90(2):259-263.

- Olson, M. (1996). Big bills left on the sidewalk: Why some nations are rich and others are poor. *Journal of Economic Perspectives*, 10:324.
- Oprea, R., Porter, D., Hibbert, C., Hanson, R., & Tila, D. (2007). Can manipulators mislead market observers. Chapman University, E.S.I. Working Papers 08-01.
- Oswald, A. (1997). Happiness and economic performance. *Economic Journal*, 107:1815-1831.
- Page, B. I., & Shapiro, R. Y. (1992). *The Rational Public: Fifty Years of Trends in American Policy Preferences*. University of Chicago Press, Chicago.
- Pennock, D. M., Giles, C. L., & Nielsen, F. A. (2001). The real power of artificial markets. *Science*, 291:987-988.
- Roll, R. (1984). Orange juice and weather. *American Economic Review*, 74(5):861-880.
- Rothschild, David (2009) Comparing Prediction Markets, Polls, and Their Biases, *Public Opinion Quarterly*, 73(5):895–916.
- Sachs, J. D., & Warner, A. M. (1995). Economic convergence and economic policies. NBER Working Paper 5039.
- Sears, D., & Funk, C. (1990). Self-interest in Americans political opinions. In Mansbridge, J. (Ed.), *Beyond Self-Interest*, pp. 147-170. University of Chicago Press, Chicago.
- Servan-Schreiber, E.J. (2012). Prediction Markets: Trading Uncertainty for Collective Wisdom. In Landemore, H. & Elster, J. (Eds) *Collective Wisdom: Principles and Mechanisms*. Cambridge University Press.
- Shiller, R. J. (2000). *Irrational Exuberance*. Princeton University Press.
- Surowiecki, James (2004) *The Wisdom of Crowds*, Random House, New York.
- Spencer, J. (2004). New ICAP-NYMEX derivatives have U.S. gas markets number. *Wall Street Journal*.
- Stiegler, M. (1999). *Earthweb*. Baen Books, New York.
- Stiglitz, J. (1998). The private uses of public interests: Incentives and institutions. *Journal of Economic Perspectives*, 12(2):322.
- Sunder, S. (1995). Experimental asset markets. In Kagel, J. H., & Roth, A. E. (Eds.), *The Handbook of Experimental Economics*, pp. 445-500. Princeton University Press, Princeton New Jersey.
- Taylor, J. B. (Ed.). (1999). *Monetary Policy Rules*. University of Chicago Press, Chicago.
- Wittman, D. (1995). *The Myth of Democratic Failure*. University of Chicago Press, Chicago.

Wolfers, J., & Zitzewitz, E. (2004). Prediction markets. *Journal of Economic Perspectives*, 18(2):107-126.

Zeckhauser, R. J., & Viscusi, W. K. (1990). Risk within reason. *Science*, 248:559-564.

I thank the Center for Study of Public Choice and the Mercatus Center for financial support. For comments, I thank Michael Abramowicz, Kathryn Aegis, Tom Bell, Peter Boettke, Nick Bostrom, Tom Breton, Damein Broderick, James Buchanan, Bryan Caplan, Ed Clarke, Joseph Coffey, Roger Congleton, Tyler Cowen, Eric Crampton, Wei Dai, Hal Finney, David Friedman, William Fischel, Andrew Gelman, Amihai Glazer, Katja Grace, Tim Groseclose, Karl Hallowell, Thomas Hanson, Paul Hewitt, Bernardo Huberman, Craig Hubley, Peggy Jackson, Ron Johnson, Michael Kremer, Ken Koford, Hassan Masum, Peter McCluskey, Jim McKinney, Steve McMullen, Joshua Miller, Mencius Moldbug, Eli Lehrer, Florenz Plassmann, John Salvatier, Edward Stringham, Nick Szabo, Alexander Tabarrok, Earl Thompson, Nicolaus Tideman, Norm Tucker, Hal Varian, Karen Vaughn, Gary Wagner, Eliezer Yudkowsky, Richard Zeckhauser.