

Shall We Vote on Values, But Bet on Beliefs?

Robin Hanson*
Department of Economics
George Mason University†

October 2007 (First Version September 2000)

Abstract

Democracies often fail to aggregate information, while speculative markets excel at this task. We consider a new form of governance, wherein voters would say what we want, but speculators would say how to get it. Elected representatives would oversee the after-the-fact measurement of national welfare, while market speculators would say which policies they expect to raise national welfare. Those who recommend policies that regressions suggest will raise GDP should be willing to endorse similar market advice. Using a qualitative engineering-style approach, we present three scenarios, consider thirty-three design issues, and finally a more specific design responding to those concerns.

*For their comments, I thank Kathryn Aegis, Tom Bell, Peter Boettke, Nick Bostrom, Tom Breton, Damein Broderick, James Buchanan, Bryan Caplan, Ed Clarke, Joseph Coffey, Roger Congleton, Tyler Cowen, Wei Dai, Hal Finney, David Friedman, William Fischel, Andrew Gelman, Amihai Glazer, Tim Groseclose, Karl Hallowell, Bernardo Huberman, Craig Hubley, Peggy Jackson, Ron Johnson, Michael Kremer, Ken Koford, Hassan Masmum, Peter McCluskey, Jim McKinney, Steve McMullen, Eli Lehrer, Florenz Plassmann, Alexander Tabarrok, "TheRob," Earl Thompson, Nicolaus Tideman, Norm Tucker, Hal Varian, Karen Vaughn, Gary Wagner, Eliezer Yudkowsky, Richard Zeckhauser, and participants of the Monomedia Berlin: Value conference, the 2000 GMU Public Choice Outreach Seminar, and the 2001 Public Choice Society annual meeting. I thank the Center for Study of Public Choice and the Mercatus Center for financial support, and Edward Stringham for research assistance.

†rhanson@gmu.edu <http://hanson.gmu.edu> 703-993-2326 FAX: 703-993-2323 MSN 1D3, Carow Hall, Fairfax VA 22030

Introduction

The space of possible forms of government is vast and largely unexplored. This fact is illustrated by the new form of governance explored in this paper. This new form seems attractive if we accept three assumptions:

1. Current democracies fail largely by not aggregating enough available information.
2. Speculative markets are the best general institutions we know for aggregating information.
3. It is feasible to, after the fact, tell rich happy nations from poor miserable ones.

The plausibility argument goes as follows. Democracies often fail by adopting bad policies, i.e., policies that in effect hurt most voters. If these policies would not have been adopted had all voters known they were bad policies, and if someone somewhere did know or could have known they were bad policies, we can say that such democracies failed because they did not induce people to acquire and share this information. If speculative markets do very well at inducing people to acquire and share information, but are little used in current governmental institutions, then we should consider changing our government institutions to rely more on speculative markets.

In order to get speculative markets to tell us which policies are bad before the fact, we require some ability to see after the fact whether we have been on net hurt or helped by the sum total of all our policies and other contributing causes. But we can roughly tell after the fact whether nations have been prosperous and happy, or deprived and miserable. Economists even frequently recommend policies that macroeconomic data suggest are causally related to increasing GDP. This suggests that economists think that GDP is often a good enough outcome measure for drawing such conclusions.

To make use of speculative markets, we can “vote on values, but bet on beliefs.” We now use democracy both to decide what we want, and to decide how to get what we want. We might instead still have democracy say what we want, but let speculative markets say how to get what we want.¹ That is, elected representatives might define a formal measure of “national welfare” (analogous to GDP) and manage its measurement after the fact. Market speculators would then say which proposed policies they expected to raise national welfare as so defined. The basic rule of government would then be this: when speculative markets clearly estimate that a proposed policy would increase expected national welfare, that policy becomes law. (As we will explain 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.)

I offer neither a formal theoretical model nor a statistical analysis of this proposal; I instead take an “engineering” approach. That is, I seek to identify relevant design issues, and then seek a basic design that, by seeming to address enough design issues, seems plausible enough to be worth testing at the next level of prototype realism. This paper will therefore qualitatively review the plausibility of the above three assumptions, describe some scenarios of how this approach might work, consider thirty-three design issues that have been raised, and finally present a more specific design intended to deal with many of these design issues.

¹For no particular reason, I’ve named this alternative “futarchy.”

Information 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 skeptics of democracy.

Most people made up their minds long before the election ... 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 high levels of ignorance continue today (Delli-Carpini & Keeter, 1997). For example, only 29 percent of U.S. adults can name their congressman, and only 24 percent can identify the first amendment of the U.S. constitution.

Formal analysis has identified many potential problems with democracy, including instability, rent-seeking, coordination failures, and commitment failures (Holcombe, 1985; Besley & Coate, 1998). Formal analysis has not, however, tended to validate what many considered the most serious problem: ignorant policies due to low incentives for voters to become informed (Wittman, 1995). This is because while it is true that voters have very little incentive to become informed, for rational voters this problem can be mitigated by other features of democracy.

For example, in theory it should be enough to have citizens vote selfishly and retrospectively, and to have policies centrally determined. That is, if it is clear who is responsible for policy, and if citizens voted for incumbents only when their personal lives seemed better than expected, then citizens would not need to understand much about abstract policies. Incumbents would then have the power and incentive to make voters feel good about their lives. Real voters, however, do not in fact seem to vote very selfishly (Sears & Funk, 1990). Thus in order to get informed voter-driven policy, we seem to need most voters to obtain information about the broader consequences of policy.

In this case ignorant voters, such as those who do not understand the consequences of import quotas, can indeed induce bad policies, such as by voting for politicians who support import quotas. Democracy seems reasonably capable of overcoming such ignorance, however, at least in theory given rational voters. After all, voters learn many things as a side effect of just living, large elections can average out random errors by individual voters, and political entrepreneurs can take the initiative to inform voters via advertising. Also, a small fraction of informed citizens can determine elections if uninformed citizens either abstain, infer what the informed know from opinion polls, or trust political parties, news media, and other sources that rely on the informed voters.

More recent empirical work has similarly found less to complain about in voters.

Collective public opinion is rational [meaning] ... real, stable, differentiated, consistent, coherent; reflective of basic values and beliefs; and responsive (in predictable and reasonable ways) to new information and changing circumstances ... On most domestic matters, about which elites often compete and provide multiple sources of

%	In U.S. Agree With This Opinion	Cite
47	God created humans in present form in last 10,000 years	(Gallup, 1999)
52	'astrology has some scientific truth'	(Davis, Smith, & Marsden, 1996)
72	believe in angels	(Gallup, 1998)
80	'US govt. is hiding that it knows of the existence of Aliens'	(CNN, 1997)
85	'Jesus Christ was born to a virgin'	(Group, 1994)

Table 1: Contrarian Public Opinions

information, the public can ... form ... opinions that approximate fully and correctly informed preferences (Page & Shapiro, 1992).

This is not to deny that there are information problems, however. The above quote continues,

In foreign affairs, on the other hand, government monopolies of information (and consensus among elites) may sometimes lead the public astray from preferences it would hold if fully informed (Page & Shapiro, 1992).

And in fact these authors go on to suggest that the U.S. public was lead far astray regarding World War II and the Cold War.

Democracy can also suffer from information failures due to long delays in information getting to many people. Such delays (or irrationality) might explain the fact that time series of public opinions do not tend to look like random walks. Even on questions of fact such as the risk of a nuclear power accident, it seems that one can predict future average opinions from trends in past average opinions (Page & Shapiro, 1992).

More relevant evidence comes from systematic deviations between expert and public opinion, about both facts and policy. Economists, for example, are much less concerned about foreign aid and down-sizing, and more concerned about reduced productivity growth (Caplan, 2002), and toxicologists tend to estimate lower risks and higher benefits from chemicals (Kraus, Malmfors, & Slovic, 1995). These tendencies also exist, at a reduced level in (men and) the more highly educated, and these differences are not attributable to factors such as age, income, or ideology. The straightforward interpretation of this data is that experts and those who are better educated actually know more than the general public about which policies are better.

Differences between public and expert opinions might be attributed to public ignorance about expert opinion. But the public also seems to have many contrarian public opinions, such as those listed in Table 1. In these cases, the public seems largely aware of how expert opinion differs from public opinion.

Contrarian public opinions suggest not just ignorance but also irrationality more directly. In theory, on matters of fact rational agents should not knowingly disagree (Aumann, 1976). For example, as two rational agents alternate telling each other their new opinion, neither one should be able to predict the direction in which the other agent will next disagree with them (Hanson, 2002). Such results generalize to agents are just trying to be rational (Hanson, 2003b). But people do in fact knowingly disagree with each other on factual claims, especially regarding

politics. This suggests that people are not just ignorant, but more fundamentally non-truth-seeking, and are self-deceived about this fact (Cowen & Hanson, 2002). This is consistent with the view that otherwise puzzling policies and political behavior can be understood in terms of weak but positive voter preferences for irrational beliefs (Caplan, 2001).

Of course misinformed and irrational voters need not prevent informed policy, if voters allow policy to be determined by informed rational elites such as perhaps academic advisors. In many areas, however, such as tariffs or immigration, policy often seems closer to what public opinion would suggest than to what relevant experts advise (Dixit, 1997; Stiglitz, 1998). This suggests that often, and perhaps usually, the public does not defer to experts. Furthermore, many case studies suggest that when governments do use academic experts, they often do so to legitimate predetermined policies, rather than to gain information to help determine policy (Barker & Peters, 1993).

If familiar political institutions suffer from information failures, how large are those failures? The huge variation in economic growth rates across nations over the last two centuries suggests that the effects of information failures may be very large (Maddison, 1995). While there are surely other important reasons why some nations are rich and others poor, an important fraction of the variation seems attributable to some nations adopting policies which relevant experts knew to be bad, and thereby becoming poorer than nations which adopted better policies (Olson, 1996; Sachs & Warner, 1995; Ascher, 1999). Since the subgroup of democratic nations also have large variations in growth rates, democracies also seem to often fail to adopt good policies.

Many factors contribute to bad policy, including commitment and coordination failures. But it seems hard to imagine that nations would adopt bad policies nearly as often as they do if it were common knowledge that such policies were bad. Thus at some level bad policy seems to be fundamentally due to a failure to aggregate relevant information.

Information Successes of Speculative Markets

While democratic policy seems to suffer from information failures, speculative markets show striking information successes.

Most markets for stocks, bonds, currency, and commodities futures are called *speculative markets* because they allow people to bet on future prices by buying or selling today in the hope of reversing such trades later for a profit. Such opportunities to “buy low, sell high” occur when identical durable items are frequently traded in a market with low transaction costs. Given such opportunities, everyone is in essence invited to be paid to correct the current market price, by pushing that 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 too much money in previous bets.

Betting markets are speculative markets that trade assets that are specifically designed to allow people to bet on particular matters of fact, such as which horse will win a race. The final values of such assets are defined in terms of some official final judgment about the fact 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 many centuries, and for many decades economists have studied the ability of such markets to aggregate information. The main finding of this research is that such markets tend to be relatively “efficient” in the sense that it is hard to find information that has not been incorporated into market prices (Lo, 1997; Hausch, Lo, & Ziemba, 1994). The main possible exceptions seem to be long-term aggregate price movements, and a long-shot bias in high-transaction-cost betting markets.

Many have suggested that asset markets have too much long-term aggregate price variation, such as stock market “bubbles” (Shiller, 2000). Risk and delay most discourage speculators from correcting such pricing errors, and irrational traders can actually gain superior returns (though not utility) from irrationally-large risk-taking (De Long, Shleifer, Summers, & Waldmann, 1990). Long-term aggregate prices, however, are also where it is hardest to empirically distinguish irrationality from rational information about fundamental economic change (Barsky & De Long, 1993), and where selection effects most pollute our data (Jorion & Goetzmann, 2000). Even if speculative markets are distorted by irrational bubbles, it is not clear that any of our other information institutions do better. For example, no other information institution in our society, such as academia or news media, consistently predicted that we were over-investing during the “dotcom” bubble. Yes some individual academics or reporters so predicted, but so did some individual stock investors.

Over the last few decades economists have also studied speculative markets in laboratory experiments, where they have more control over trader information and preferences. Such experiments find that speculative markets aggregate information well, even with four traders trading \$4 over four minutes, and even when such traders know little about their environment or other traders (Sunder, 1995). For example, traders can aggregate information well when they are experienced in their role and abstractly know the payoffs of players in other roles (Forsythe & Lundholm, 1990). If the structure of traders’ information is complex enough relative to the number of assets available to trade, however, information “traps” can occur where individual traders have no direct incentive to reveal their information (Noeth, Camerer, Plott, & Webber, 1999). Such problems are typically, though not always, reduced by allowing trading of more kinds of related assets.

Absolute accuracy levels, however, are not the key issue. The key policy question about any institution is how it performs *relative* to alternative institutions dealing with the same situation. A few studies have presented field data on this question, directly comparing real world speculative markets with other real world institutions for aggregating information. For example, racetrack market odds improve on the prediction of racetrack experts (Figlewski, 1979). Florida orange juice commodity futures improve on government weather forecasts (Roll, 1984), Oscar markets beat columnist forecasts (Pennock, Giles, & Nielsen, 2001), and gas demand markets beat gas demand experts (Spencer, 2004). Betting markets beat major national opinion polls 451 out of 596 times in predicting U.S. presidential election results (Berg & Rietz, 2002). Finally, betting markets beat Hewlett Packard official forecasts 6 times out of 8 at predicting Hewlett Packard printer sales (Chen & Plott, 2002; Plott, 2000).

Unfortunately, no studies have directly compared estimates from speculative markets to estimates from academic-style institutions. We do know, however, that those who do best at betting on horse races are smart in ways they can not articulate, and in ways unrelated to I.Q. (Ceci & Liker, 1986). Academic-style institutions, in contrast, seem largely limited to

aggregating articulated knowledge from those with high I.Q. Academic institutions put a great deal of weight on the opinions of experts relative to ordinary people. And while speculative markets may put less weight on experts, it does not seem that they place too little; if anything, they seem to put too much weight on experts, both public and private (Figlewski, 1979; Metzger, 1985; Lichtenstein, Kaufmann, & Bhagat, 1999).

How can betting markets beat opinion polls when they use the same fallible human sources? A study of election betting markets found that traders overall suffered from standard biases such as expecting their favored candidate to win, and seeing that candidate as having won debates. “Market makers,” however, were found to be on average much less biased. These were traders who made offers that others accepted, rather than accepting offers made by others. Compared to other traders, market makers invested twice as much, traded more, earned higher returns, and made one sixth as many errors. They also tended to be more highly educated, and experienced at trading (Forsythe, Nelson, Neumann, & Wright, 1992; Forsythe, Rietz, & Ross, 1999). Betting markets seem to meet or beat competing institutions in part because of the disproportionate influence such markets give to rational and informed traders. We also know more generally that people with stronger incentives to be accurate show fewer cognitive biases (Kruglanski & Freund, 1983).

There are costs to create and run markets, so there is a limit to the number of markets that can be created. However, while it was once thought that speculative markets could only be viable if they annually traded millions of dollars, say 10,000 trades of \$100 each (Carlton, 1984), it is now clear that much smaller markets are viable. For example, laboratory experiments consistently show the viability of very small markets. Low internet transaction costs are also now spurring a burst of innovation exploring a great many new market forms (Varian, 1998; Shiller, 1993; Hanson, 2003a). Play money web markets are now available where anyone can create new betting topics, and where a handful of traders betting play pennies once every few weeks are typically successful at aggregating information into prices (see, for example, *hsx.com*, *ideosphere.com* (Kittlitz, 1999; Pennock et al., 2001)).

Gambling and securities regulations make it very difficult, however, to create real money markets like these play money markets. This regulatory block on financial innovation should not be surprising, because all of our familiar financial institutions were once prohibited by laws against gambling and usury. For example, 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).

The history of financial regulation can thus be roughly summarized as everything being banned as gambling (or usury) until an exception was granted for some newly legitimized higher purpose. For each purpose, such as capitalizing firms, insuring idiosyncratic risk, or insuring common risk, laws and regulations were created to ensure that generic gambling could not slip in. We may thus reasonably hope to someday legitimate, and thereby legalize, markets whose main function is to aggregate information on questions that matter (Bell, 1997).

Measuring Welfare

Since speculative markets do so well at a task that democracies have troubles with, it is tempting to try to improve democracy by making them rely more on speculative markets. However, in order to use speculative markets to advise national policy, it helps to be able to settle bets about how well off a nation has been.

Many have argued that we typically know too little about individual preferences to help people via detailed interventions into their daily lives (Buchanan, 1988). As we aggregate more across people, time, and contexts, however, many of these uncertainties seem to cancel, allowing us to make more accurate judgments about how well off people are on average. It seems reasonably clear, for example, that on average people in Ethiopia today are less satisfied with their lives than people in Sweden today. This can be true even if we would find it difficult to advise any random Ethiopian on what he should eat for lunch.

The development of national accounts, i.e., the collection of statistics such as GDP, has been justifiably called one of the greatest economic innovations of the twentieth century. While such measurements are made with error, and leave out a great deal, they seem to be a sufficient basis for many policy recommendations. For example, most researchers in empirical growth economics seem willing to presume that policies are good if they causally induce sustainably higher GDP. And frequent travelers find it hard to escape the impression that, comparing nations with large differences in measured GDP, most people who live in the high GDP nations are richer and better off than most people who live in the low GDP nations.²

Furthermore, economists spending modest budgets have already explored many promising extensions to current GDP measures, such as lifespan, home production, leisure, environmental assets (Boskin, 2000; Nordhaus, 2000). Economists also have well developed frameworks for discounting future welfare, and for dealing with risk aversion and inequality. Direct measures of happiness can even be included (Oswald, 1997). It therefore seems plausible that with greatly increased funding and the full attention of our elected representatives, we could devise even more accurate and robust measures of national welfare.

Consider a person who is willing to recommend a certain policy in a given situation because statistical analysis suggests that those policies tend to cause GDP to rise in such situations. Such a person should also be willing to recommend this policy if speculative markets were to estimate that this policy was expected to increase GDP in this situation. Speculative market estimates are probably no less accurate than a typical statistical analysis. After all, if a statistical analysis is persuasive then speculators should be persuaded by it, and if speculators disagree with a statistical analysis they probably have good reasons.

Therefore, if one is willing to recommend policies that statistical studies suggest will increase (a time average of future) GDP, one should be willing to recommend policies that speculative markets estimate will increase GDP, and so one should be willing to consider a form of policy making which relies more on such market estimates in choosing policies. While GDP leaves out many things, nations that consistently adopt the best possible policies for increasing (a time average of future) GDP would in a few decades probably be much richer (and happier) than they would otherwise have been. And one should be even more willing to consider basing

²Similar claims apply to other rough welfare measures, such as lifespan and other development indices; I focus on GDP here for concreteness.

policy on speculative market estimates of their effects on GDP+, i.e., an improved measure of national welfare.

Decision Market Mechanics

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

Consider a simple market where one can bet on the event D , i.e., the Democratic party wins the 2008 U.S. presidential election. In such a market a bank could without risk accept \$1 in payment for the pair of contingent assets, “Pays \$1 if D ” and “Pays \$1 if not D .” This transaction carries no risk because exactly one of these assets will be worth \$1 in the end. Since the expected dollar value of the asset “Pays \$1 if D ” is $p(D)$, if someone is willing to buy this asset for \$0.70, we can interpret this willingness as this person saying that the chance that Democrats will win is at least 70%. And a market price of \$0.70 can be interpreted as a consensus among potential traders that $p(D) \approx 70\%$.³ After averaging in their minds over plausible scenarios, traders would have judged that Democrats win in about 70% 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 C , i.e., Hillary Clinton is the nominee of the Democratic party in the 2008 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 this time they should only consider scenarios consistent with the event C , that Clinton is the nominee. A market price of \$0.60 here can be interpreted as a consensus that about 60% of these scenarios have the Democrats winning. This would be a consensus that $p(D|C) \approx 60\%$, i.e., that the *conditional* probability of Democrats winning, given that Clinton is the nominee, is 60%.

Another way to get the same sort of conditional probability estimates is to have a market where, instead of trading the asset “Pays \$1 if D ” for some fraction of the asset \$1, people trade the asset⁴ “Pays \$1 if D and C ” for some fraction of the asset “Pays \$1 if D .” Since this is trading an asset worth $p(D \& C)$ for some fraction of an asset worth $p(C)$, the market-price-fraction here can be interpreted as an estimate of the ratio $p(D \& C)/p(C)$, which is by definition equal to the conditional probability $p(D|C)$. Thus if the Democratic party wanted to know which nominee would give them the best chance of beating Bush, they could compare the market estimate of $p(D|C)$ to similar estimates for other Democratic candidates.

Speculative markets can estimate not only probabilities, but also expected values. Let the variable x denote the Democratic party’s fraction of the total vote in the presidential election. A bank can accept \$1 in payment for the pair “Pays \$ x ” and “Pays \$1 – x ,” and then the

³This interpretation neglects possible differences in the value of money to traders between the sets of states D and not D .

⁴A bank can without risk exchange \$1 for the set of four assets “Pays \$1 if D and C ,” “Pays \$1 if not D and C ,” “Pays \$1 if D and not C ,” “Pays \$1 if not D and not C .”

fraction of \$1 required in trade for the asset “Pays \$ x ” can be interpreted as a market estimate of $E[x]$. Similarly, the fraction of the asset “Pays \$1 if C ” traded for the asset “Pays \$ x if C ” can be interpreted as a market estimate of $E[x|C]$.

Since 1996, the Iowa Electronic Markets have actually had markets whose prices could be combined to represent estimates of each party’s expected vote fraction conditional on various possible nominees for the other party (Berg & Rietz, 2003). For example, 1996 prices indicated that Dole, the actual Republican nominee, was a weak candidate against Clinton; there was always another candidate who speculators said would have received a higher expected vote fraction than Dole.

One can similarly create markets to estimate whether a new policy N is an improvement over the status quo Q in increasing national welfare W . Given a measure of national welfare W , normalized to be between zero and one, markets that trade assets “Pays \$ W if N ” for some fraction of “Pays \$1 if N ” give a market price estimate of $E[W|N]$. Just as people considering what fraction of \$1 to pay for “Pays \$ x ” will estimate the average value of x across all plausible scenarios, people considering what fraction of “Pays \$1 if N ” to pay for “Pays \$ x if N ” will average only over plausible scenarios consistent with the event N . When the market estimate of $E[W|N]$ is clearly greater than $E[W|Q]$, speculators are saying that this new policy is expected to increase national welfare.

The Engineering of Institutions

It is tempting to use the success of betting markets as information institutions to solve the problems of democracies as information institutions. But do we know enough about either type of institution to be proposing new forms of government based on this idea?

It depends on whether one thinks like a scientist or like an engineer. 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 some other 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 government, which can not possibly have direct empirical support, and which are too complex for our theories to make direct predictions.

An engineer, on the other hand, is more interested in improving systems than theory. An engineer is happy to work on a concept with a five percent chance of success, if the payoff from success would be thirty times the cost of trying. An engineer uses theory-informed intuitions to consider a wide range of design issues, and if the concept still seems promising, moves on to increasingly realistic “proof of concept” *prototypes*, from simulations, to “wind tunnel” models, to field tests. Scientists have little use for prototypes and their tests, being neither theory nor data that tests theory, but prototypes are what make the engineers’ world go round.

While academic study of social systems is now mostly dominated by a scientific style, this paper takes an engineering style to consider a new form of governance. The purpose of this paper is not to induce high confidence that this concept would work well, but merely to raise readers’ confidence up to a level that would justify further exploration via the next level of prototype. This paper thus, from this point forward, mainly takes on an engineering tone, qualitatively identifying and addressing a wide range of design issues.

Scenarios

This general idea of using speculative markets more widely to aggregate information 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 information (Zeckhauser & Viscusi, 1990; Hanson, 1995, 1999; Berg & Rietz, 2003; Hanson, 2006). While such markets could be placed in an informal advisory role, the following three scenarios illustrate how they might be given a more formal role in corporate, agency, and national governance.⁵

Corporate Governance

Most corporations are in effect small democratic governments. The corporate charter is the constitution, and the owners of stock are the citizens. The charter authorizes citizens to vote in proportion to their holdings for a board of directors, who choose and compensate the CEO, who runs the company. The CEO compensation plan usually gives the CEO a strong incentive to maximize the stock value, but the directors usually have weak incentives, and are often chosen by the CEO. As a result, boards of directors seem to consistently wait too long before firing a bad CEO.

Once smaller scale experiments were successful, corporations might consider using speculative markets to advice big corporate decisions. Imagine that the corporation OldTek amended its corporate charter so that the continued tenure of its CEO depended on certain market prices. A financial house agreed to ensure that there are always dump-the-OldTek-CEO stock markets, trading OldTek stock conditional on dumping the current CEO at the end of the current quarter.

In an ordinary stock market people agree to trade cash for stocks, and so the ratio of cash to stock in such trades is the market price, with the bid and ask prices being the prices at which people offer to buy and to sell stock. In a dump-the-CEO 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 by the end of the current quarter.⁶ The OldTek corporate charter might say that if the dump-the-CEO price were clearly higher than the keep-the-CEO stock price, the board must dump him within the next quarter. (The price might be deemed "clearly higher" if the dump-the-CEO bid price were above the keep-the-CEO ask price for 90% of the last week of a quarter.)

If non-corporate-insiders were guaranteed trading anonymity, then Disgruntled Jones, a junior OldTek employee, could avoid retribution if he made offers in the dump-the-CEO market to buy OldTek stock. The CEO might accept sell offers in this market in order to try to keep his job, but the OldTek charter might say that his trades and offers must identify him, and that his compensation must be public knowledge. (These rules might apply to all OldTek insiders.) Thus while the CEO might want to bias the dump-the-CEO stock price, observers

⁵One could also imagine international decision making scenarios. For example, a uniform international carbon emissions tax level might be set by markets that estimate future world product conditional on different tax levels.

⁶Alternatively, the assets could be contingent instead of the trades. That is, assets of the form "Pays \$1 if CEO dumped" could be traded for assets of the form "Pays one OldTek stock if CEO dumped."

would have plenty of information to help them correct for such a bias. If enough traders agreed with Disgruntled Jones to dump the CEO, and if OldTek prospered because of this, traders like Jones would be directly rewarded.

Agency Governance

After some successful examples of using speculative markets in corporate governance, a federal agency might consider trying it.⁷

For example, the central banks of most nations periodically make changes to short-term interest rates, reserve requirements, currency conversion rates, and other monetary policy parameters. It is widely agreed that such actions are made in an attempt to smooth out business cycles, limit inflation, and promote high average growth rates. In fact, the future values of GDP growth rates, inflation, and unemployment are generally considered to describe of the main outcomes of interest.

While specialists in this area sometimes disagree about the relative weight to give to smoothing unemployment versus lowering inflation versus increasing growth, what they mostly disagree on is *how* to change monetary policy in order achieve these outcomes. For example, some argue that central banks should hire knowledgeable people, keep them independent, and then let them choose as they see fit. Others argue that central banks should be more politically accountable, and still others argue that central banks should just announce some policy rule and then stick to it (Taylor, 1999).

Monetary policy is thus an area where people largely agree on a few statistics by which one can tell, after the fact, the quality of the policy outcome, and where they agree on a few control variables to use to get a good outcome. Monetary policy is thus a good candidate for getting advice from speculative markets. To get such advice, the government would declare an explicit function describing good outcomes. This might, for example, be a weighted sum over each of the next forty quarters of that quarter's GDP growth rate minus its unemployment rate minus its inflation rate, minus the absolute value of how much each parameter changed quarter.

Outcome assets would be created whose final value is linear in this outcome function. That is, the outcome measure would end up being some specific number, and each outcome asset could then be exchanged for that number of a certain chosen 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 complementary value (the maximum amount minus the outcome measure value).

Each week there could be three markets, up, down and same, to determine whether interest rates should change at week's end. In the up market, people would trade some amount of cash for the outcome asset, except that these trades would be called off if interest rates were not raised one unit that week's end. Trades in the down market would be similar, being called off unless interest rates were lowered one unit. Trades in the same market would be called off if

⁷The author was an architect of the "Policy Analysis Market." Cancelled in July 2003 after being labeled "Terrorism Futures", this market would have given estimates of geopolitical trends in the middle east conditional on US policy choices. More at: <http://hanson.gmu.edu/policyanalysismarket.html>.

either change was made. If either the up or the down market had a consistently higher price over that week than the same market, interest rates would be raised or lowered accordingly.

Information about the state of the economy that is now given privately to monetary officials could instead be made public, and hence available to market speculators. Monetary experts would have to persuade market speculators to influence monetary decisions. And market speculators would have to decide which, if any, of these experts to believe.

National Governance

Once some government agencies were successful in getting decision advice from speculative markets, a nation might consider changing its constitution to put such markets at the core of its form of government. With such a government, we would vote on values, but bet on beliefs. The job of elected representatives would be limited to defining and overseeing an after the fact measurement of $GDP+$, a measure of national welfare.

At any one time there would be a market estimating $GDP+$ given the status quo Q of no policy change. There would be some official process for proposing new policies, perhaps involving a proposal fee. And for each new proposal N there would be a market estimating $GDP+$ if that proposal were adopted. The basic rule of government would then be:

When an approved betting market clearly estimates that a proposed policy would increase expected $GDP+$ ($E[W|N] > E[W|Q]$), that proposal immediately becomes law.

Democracy would still say what we want, but betting markets would now say how to get it.

A very simple definition of $GDP+$ would be a few percent annually discounted average (over the indefinite future) of the square root of GDP each period. A not quite as simple $GDP+$ definition would substitute a sum over various subgroups of the square root of a GDP assigned to that subgroup. Subgroups might be defined geographically, ethnically, and by age and income. (Varying the group weights might induce various types of affirmative action or discrimination policies.) A more complex $GDP+$ could include measures of lifespan, leisure, environmental quality, cultural prowess, and happiness.

One might also want to include consequences for foreigners and animals in a national welfare function. In fact, treaties between nations might take the form of each nation agreeing to explicitly giving the other nation's welfare function some particular weight in its own welfare function. Finally, some of the most difficult choices in defining welfare may well be how to treat changes in national population due to births, immigration, and emigration.⁸

In addition to defining $GDP+$, we would have to choose some base asset in which payoffs are denominated. (The "base asset" in "Pays \$1 if B" is the asset \$1.) This base asset should have competitive returns over the long term, if people are to hold them over the long term. Long-term inflation-indexed government bonds or shares of broad stock index funds are reasonable choices. The national welfare assets traded would then be of the form "Pays $GDP+$ units of the base asset." So if the base asset were index fund shares, then once $GDP+$ was determined to be 0.127, each welfare asset could be exchanged for 0.127 index fund shares.

⁸We might, for example, give a 10% relative weight to any person anywhere, and a 50% weight to our allies, possible children, and likely immigrants and emigrants.

Design Issues

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 thirty-three design issues, each expressed as an objection to this overall approach. After this we will present a more detailed proposal for a form of government.

Maybe Low Information Is Not Democracy's Major Failing

What if democracy's biggest problem were not its ability to aggregate information, but some other problem, such as a 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. We would still 'vote on values,' and so we could change this democratic component to adjust for other problems. For example, some governments might face the problem of paying off powerful groups who could otherwise overthrow the government. Instead of paying them off via direct financial payments or pork barrel programs, we might pay them off by raising the welfare weight on outcomes they favor.

Democracy Might Become Unstable

In theory, democracies seem subject to policy instability, or "cycling," since for any policy there is typically another policy that some majority favors. This does not seem to be a large problem in practice (Tullock, 1981), but this may be in part due to voter ignorance, as some theories suggest (Hanson, 1997). If so, this proposal might make instability worse, by better informing voters.

This might be a reason to lean toward more stable democratic institutions. We might prefer one or more large decentralized legislatures whose members have staggered terms, and we might shy away from powerful presidents or proportional representation systems in which coalitions rule as a group. If values change more slowly than beliefs about policy consequences, and if less expertise is required to make value judgments, then the democratic part of this form of government might reasonably be allowed to be slower and less expert than today.

People Prefer Their Comfortable Illusions

Many people may enjoy their illusions about politics and policy. If so, we might hurt them if we forced them to forgo such illusions via better information aggregation. Fortunately, voters need not be directly confronted with speculator beliefs. The prices that determine policy need not directly explain speculators' reasoning, and news media that catered to viewers with cherished illusions could avoid explaining this reasoning. Even today, media often avoid telling viewers distasteful details about how legislative sausage is made.

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, can distribute in controlled amounts, and know that most people value. Only some form of money now seems to satisfy these requirements. “Play” money might be used, if enough people valued it as bragging rights.

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 result in 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. If such speculators believed that such policies would best achieve national welfare, then betting for these policies would be the best way for each speculator to be as rich as possible, given whatever policies are approved. Each speculator typically has little influence over policy in general, but has far more influence over his personal betting wins or losses. This proposed form of government should thus appeal to those who can accept the values democracy would favor, and who think that speculators would agree with their assessment of how to best promote such values.

People Would Reveal Less And Lie More

Perhaps many people now often reveal their politically relevant information, and refrain from lying about such topics, because they feel it is their civic duty to do so. 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 information, in order to benefit from related bets. Perhaps info shared via bets would not make up for lost info in spoken words.

Those who feel that civic duty makes them speak honestly about politics are often self-deceived, however. Not only can betting markets discourage such self-deception, but once people run out of ways to bet on their secrets, they should want to reveal those secrets and persuade other speculators to believe them, so that they can reverse their earlier trades for a profit. If secret-holders never think they have run out of ways to bet on their secrets, this may be a problem. But even so, the studies mentioned earlier showing speculative markets beating other institutions suggest that other institutions often have worse problems.

Time-Inconsistency Might Be A Problem

The inability to commit has been identified as a major reason, in theory and some practice, for government failures (Levy & Spiller, 1994; Besley & Coate, 1998). The basic rule of gov-

ernment described above does not overcome this problem, but it also does not obviously make this problem any worse. So perhaps we should continue with the current approach, which is to mostly only allow governments to commit if they can arrange to do so themselves, such as via transaction costs, posting bonds, external reputation, or other external institutions. Any advantages from creating more direct ways to avoid time inconsistency problems must be weighed against the problem of allowing earlier generations to take more advantage of later generations.

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 why this would make expressive voting more of a problem overall.

The Rich Would Get More Influence

The rich have more influence in most every area of life, including politics in every known political system. After all, the point of being rich is to be able to get more of whatever it is you want. If we do not want the rich to get as much as they do, let us tax them more, and they will not be as rich. Perhaps the concern here is that the rich should have proportionally less influence in politics than they do in other areas of life. But in this new form of government the rich have proportionally less influence in politics, if their relevant information is proportionally less, as seems likely. Betting markets are not opinion polls where the rich get more votes. The rich have more *potential* votes, but they will lose those votes if they use them without having proportionally more information. And the rich do not tend to throw their money away easily; those who do, do not stay rich very long. Also, we have the option to limit how much money each person can use to the market to trade on policy proposals.

One Profits Little By Supporting Unlikely Proposals

Imagine that you thought well of some policy proposal, but that few people 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. In your eyes it seems unfair that you have never been rewarded, and they have never been punished, for your being right. Of course from each trader's point of view there should always be some chance each proposal will be adopted, since other traders might jump in and support it before the deadline. (And the smaller the chances of the proposal passing, the cheaper it might be to try this.) So there is always some possible reward from trading on unlikely proposals, though with finite transaction costs this incentive might not be enough.

Every known political system has a similar problem. Supporters of contrarian proposals might first try them out on smaller scales, perhaps by getting other speculators to allow local regimes where such proposals can be tried. Also, the insights that tell you why your proposal

is good might be useful in predicting the consequences of proposals that are actually adopted. If so, you might first bet to convince others of your insights, then convince them to support your proposal.

Markets Might Be Too Thin To Make Good Estimates

The proposal here is to use market prices to change the status quo. So if some markets failed to exist, the worst that could happen would a failure to change the status quo when it should change. Very “thin” markets, however, with few traders and noisy prices, might possibly cause harm by accidentally seeming to clearly favor a bad proposal. This might be a good reason to use a conservative definition of whether a market “clearly” favors a proposal.

A market that influences important policies, however, simply cannot be very thin (Hanson & Oprea, 2004). 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 were similarly funded, their combined trades would thicken the market. Alternatively, if one side were willing to spend much more, then speculators who knew this fact could make easy profits by trading against that better funded side, again thickening the market. So either way, the market would not end up thin.

Bad Guys Might Pay To Add Noise

Some of our policies effect foreign nations, and so the markets that estimate the consequences of such policies might be tempting targets for foreign agents. Even if they couldn’t bias the estimates in their favor on average, couldn’t foreign agents just trade randomly and so reduce our ability to estimate the consequences of our policies? The thicker such markets were, the more it would this strategy would cost. And in fact if speculators anticipated such attempts, adding more noise should make prices *more* accurate. This is a general feature of speculative markets; all else equal, more “noise” traders, who trade for reasons other than having information about the asset, make prices more accurate. Other speculators scale up their trades and their information gathering efforts.

One Rich Fool Could Do Great Damage

Even Bill Gates, the richest person alive today, has only a small amount of wealth compared to that available to all market speculators. So if most speculators were confident that Bill Gates were wrong about some market estimate, even Bill Gates could not substantially change that estimate. If Gates pushed the price one way, all those speculators would, by trading, push it right back. More likely, however, speculators would allow Gates to move the price some because they reasonably suspect that he had access to relevant information. Given his successes, this would not be an unreasonable assumption.

People Could Buy Policy Via Trades

If prices determined policy, and if trades move prices, then could someone buy a favorable policy just by making the right trades? No; a new trade typically moves prices because other traders

suspect that this new trader has new information. When other traders with deep pockets can clearly observe that a particular person is trading for non-information reasons, such traders will not allow the price to change. These non-information reasons can include entertainment, idiocy, needing to cash out, *and* wanting to manipulate the price. The same thing happens if other traders cannot observe which trades being made by such “noise traders,” but know the total quantity and direction of noise trades. If traders can predict a bias, they can and will correct for it.

If other traders are sure that a certain trader has relevant information, they will let him move the price, and will try to not trade against him. A manipulator can thus only influence the price if other traders both think he might have relevant information, *and* if they are unsure about the sign and strength of manipulation motives. Here other traders will respond to an average of an informed trader and a simple manipulator. Thus on average interested parties cannot bias the price; they can at worst only add noise.

Sometimes observers will under-estimate how badly a manipulator wants to raise the price, and this will allow him to raise the price. But such cases should be exactly balanced by other cases where observers *over*-estimate how badly he wants to raise the price, which *lowers* the price. Furthermore, noise traders who lose their bets on average attract speculators to study this topic more carefully. In fact, when other traders have deep enough pockets, prices in markets with more manipulators (or more noise traders of any kind) should be no less accurate, and can be *more* accurate! This is the implication of standard theory (Hanson & Oprea, 2004), and has been verified in laboratory experiments (Hanson, Oprea, & Porter, 2006; Oprea, Porter, Hibbert, Hanson, & Tila, 2007), and in field (Camerer, 1998; Wolfers & Zitzewitz, 2004).

For example, interested insiders who happen to have an unusually strong desire to raise the market estimate of a certain stadium’s profitability might succeed. But if so some other stadium’s estimate would be reduced because their insiders happened to have an unusually *weak* desire to raise that profitability estimate. And the existence of such insiders would attract more speculators to study stadium profitability.

We Need To Verify If A Proposal Was Implemented

To implement a rule requiring market-approved proposals to be adopted, we want it to be clear whether approved proposals have in fact been adopted. This includes being clear about what old policies each new policy would invalidate. After all, it is hard to enforce ambiguous laws. It seems sufficient, however, to treat proposals like contracts, in the sense of being responsible for saying how to resolve such ambiguities. A contract can specify its rules of interpretation or the arbitrator who should do the interpretation, and contract law can lie in the background to fill in remaining ambiguities. Similarly, official policy proposals could specify their interpretation regime as well as who would do the interpreting.

If it were not clear to speculators how a proposal would be interpreted, such speculators would have to average over likely interpretations in estimating the consequences of such policies. An unclear proposal might be a sloppy but sincere attempt to improve national welfare, or an insincere attempt to slip in policies that would not be approved if proposed clearly.

People Could Do Harm to Win Bets

Insurance companies worry about people deliberately setting fires to collect fire insurance, or killing their spouses to collect life insurance. In principle, public corporations should also worry about people selling their stock short and then damaging their factory or delivered product. This latter scenario is almost never a problem, however. Suspicions about the Tylenol poisonings case and the World Trade Center attacks were never substantiated. The closest example I can find is Roger Duronio, a PaineWebber employee who in 2002 set off a logic bomb in one thousand company computers after investing \$20,000 in options betting that the stock price would fall. The damage totaled \$3 million, but system redundancy prevented any loss of data, the stock price did not fall, and Duronio was soon caught. (Geller, 2002) Apparently, few people can substantially damage an entire company, and they are easy enough to monitor. Criminals with an ability to harm seem to prefer simple extortion, as in protection rackets.

Not only is the ability to damage national welfare even rarer than the ability to damage a corporation, but there are already many speculative markets where people could profit enormously from such an ability. Furthermore, damaging actions should only favor market evaluations of a particular policy 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 damaging actions, 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 yourself. This seems an unlikely scenario.

Welfare Measurement Might Be Corrupted

In business, the more that rides on a measurement the more one must work to avoid measurement corruption. If you're going to reward your sales staff based on how happy customers say they are, for example, it is best not to have the sales staff manage the customer happiness survey. The former Soviet Union was full of examples of bad business measures, such as the proverbial factory that was rewarded for the tons of nails it produced, which then produced a few enormous nails. However, a great deal now rides on GDP measurements, yet most rich countries seem to have limited corruption in long-term GDP measurements to tolerable levels. Also, many businesses seem to have found ways to limit corruption in measuring many important business statistics. One robust approach is redundancy; several different agencies could collect similar data and process it in similar ways, with the official estimate being the median agency estimate.

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, you would have 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 deep financial pockets or convince deep-pocket speculators that you had so committed, without revealing your corruption to enforcers. This seems an unlikely scenario.

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. Similarly, we should be wary of allowing markets to change policies that change the political process that leads to changes in the welfare definition.

A treaty between two nations might, among other things, include a provision where each nation gives the other nation's outcomes more weight in its own national welfare definition. But if we asked speculative markets to evaluate such a treaty, we would be allowing markets to influence welfare definitions. We might thus prefer to separate such treaties into two components, one approved by markets, and a welfare definition change approved by the "vote on values" side of government.

Policies Might Be Encoded In The Welfare Definition

Elected representatives could in principle directly get whatever decisions they wanted by encoding them in the definition of welfare. 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. But more generally we might have to rely on the public to consider such welfare terms as scandalous. At worst, this new form of government would just have reverted to our something like familiar democracy.

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 factors excessively influence policy, we can require that markets "clearly" favor a proposal over the status quo, via consistently higher prices over a substantial period of time. But then how could small changes ever be approved? The answer is that 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.

An interesting general policy would be to have laws "recurse," via derivative legal regimes with more focused welfare measures. So 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). A stadium that would not noticeably effect national welfare may noticeably effect regional welfare. Similarly, a general defense policy might base defense decisions on the chances of deterring or winning a war.

The Military Needs Secrecy, While 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 as a policy a regime in which military decisions are made in private. That private regime could be a recursion 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 such cases, speculators would take into account the fact that such overrides would be public, and therefore known to potential enemies. If they still approved the policy, it would be because they estimated the secrecy loss to be outweighed by some greater gain.

We Must Define When A Market “Clearly” Estimates

There are many detail questions about how exactly to implement the betting markets. What kind of “cash” do people bet, who can issue betting assets, and what marketplaces count for setting prices? What information about offers and trades is made public? Who can trade in these markets, and how much money can they each bring in? Who can make proposals, and is there a fee to make proposals or a reward for approved proposals? How large a price difference lasting for how long is a “clear” price difference? Are market trades taxed or subsidized? What are the maximum and minimum national welfare values?

The recursion approach mentioned above can deal with many of 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. If there ends up being a clearly best way to answer these questions, different fundamental conservative rules might result in nearly the same in-practice rules.

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 would also all have to pay more attention to the proposal process.

To adjust the rate at which proposals are made, we could choose a fee that people must pay to make a proposal official. (This fee might be refunded, perhaps even at a reward multiple, if the proposal were accepted.) As described above, legal recursion could set this fee by having a high fee at the base level, and inviting the creation of regimes at other levels with lower fees. If the basic rule required a \$10 million fee for making a proposal official, then within this regime speculators might estimate that the additional changes that would result from lowering this fee to \$10 thousand would on average raise national welfare. In this case, speculators would approve as a policy a new regime which authorized proposals using this lower fee.

You Might Not Catch Buggy Decisions Quickly Enough

If a bad decision were made due to bad information, then market decision advice should be reversed the moment speculators became aware of this fact, giving legal 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? 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.

So far we have suggested that a proposal would be approved today if a market clearly estimated good consequences for welfare as it is defined today. To deal with this welfare bug problem, we could also allow such proposals to be vetoed if another market clearly estimated bad consequences for welfare as it will be defined in the future, say in one year. 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. This system might give more power to representatives who have private information about likely welfare definition changes. But this may be a reasonable price to pay to avoid bad decisions due to welfare definition bugs.

This Is Too Big A Change All At Once

Speculators should take into account the negative costs of too quickly disrupting established processes when they estimate the consequences of each proposed policy change. 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 were 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.

Basic Rights and Freedoms Are Not Guaranteed

Rights and freedoms could be preserved by putting them directly into the welfare function. That is, the government might collect statistics regarding the number and types of violations of basic rights and freedoms, and the definition of national welfare might give a large negative weight to such outcomes. In addition, the constitution might empower courts to reject policies that violate certain stated rights.

It Seems Hard To Make One Measure Encode All Our Values

The democracy part of government could build a welfare measure incrementally. They would 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 the most tolerable.⁹

Decisions Would Magnify Measurement Errors

Under this form of government, speculators should raise measured welfare as much as possible. This can be done not only by raising real welfare, but also by raising the positive error between measured and real welfare. The easier it is for policies to move those errors in a consistent direction, the more harm might be done by a crude welfare measure. Businesses and other institutions have long had to deal with related problems, however, including what incentives to give a CEO, a plant manager, a salesperson, a teacher, or a student. Measured performance always differs from real performance in some ways, and that difference can cause problems. But those problems do often seem to be manageable. Note also that a similar problem occurs with *any* way in which things learned from looking at crude welfare measures are allowed to influence policy. Any such influence on policy is also at risk of responding to, and increasing, welfare measurement error rather than real welfare.

Easy to Measure Values Would Be Over-Emphasized

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 might 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 the case of the teacher ends there, however. It is risk-aversion, together with exerting hidden effort, that makes teachers reluctant to be paid based on very 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 therefore no 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.

⁹A different approach would be for elected representatives to pick a distribution over hypothetical test choices between national outcomes, instead of directly choosing a welfare function. Juries or the legislature would later have to make choices for a random subset of these hypothetical cases. Before this, private groups could then propose welfare functions, and the official welfare function at any one time would be the one that speculators estimated would give the smallest error rate on test choices. (This general approach was suggested by Peter McCluskey.)

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.

You Can Not Pay Off Bets If Earth Is Destroyed

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 effect the chances of "destruction," such as all known intelligent life being destroyed, or any other event that would prevent paying off the bets. We do care, however, about how our policies might effect the chance of such events.

If this bias were low enough, we might prefer the simplicity of ignoring it. If the bias might be high, however, we might prefer to expand the ex post measurement of national welfare to include estimates of the chances so far of various types of destruction, chances we have avoided but want to take into account. Welfare could then be a weighted average of welfare as we see it and welfare given various types of destruction, weighted by chances assigned to each type of destruction.

But Infinity Never Arrives

The proposed initial welfare definition was a few percent annually discounted average of some annual welfare measure, such as square root GDP. This is an infinite sum, however, and so would require welfare-measuring-agencies to, ever year, measure and declare annual welfare not only according to the current definition, but also according to every previous definition ever used. An alternative would be to define total welfare as an average over annual welfares for, say, the next twenty years, plus some weight times welfare as it will be officially defined twenty years hence. Even easier, total welfare could now be defined as a combination of this year's annual welfare and total welfare as it will be defined next year.

Risk Premia Distort Estimates

Investments along one or perhaps a few dimensions of aggregate asset prices are paid a risk premium. This means that while asset prices should equal an expected value of the future asset prices, averaging over different possible states of the world, states in which future aggregate asset prices are high would be 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.

Correlation Is Not Causation

When you are in full control of a decision, standard decision theory tells you to make the choice that, according to your values and information, gives the best expected outcome conditional on your making that choice. If other people are watching you, however, then even if they exactly share your values they do *not* necessarily want you to make the choice that they estimate will give the best outcome conditional on your making that choice. This is because they have to

worry that the fact you make a particular choice could tell them something about things you know that they do not.

For example, imagine a company's board of directors were going to keep or dump their CEO based on their estimate of the expected value of the stock conditional on keeping or dumping the CEO. Imagine further that speculators knew less than the board. There are some possible distributions over outcomes where speculators will estimate a higher stock value conditional on keeping the CEO than on dumping him, even though the information speculators have says the company is better off dumping him.

This sort of effect can distort betting market decision advice when the market speculators are not as well informed as the agents who are making the decision. To avoid this distortion, it is sufficient either to have the informed insiders trade in the markets, so that prices reflect what they know, or to put market estimates directly in control of decisions, rather than making such price estimates advisory. In either case it is important to make the time of the decision clear, by specifying a short time duration during which the proposal can be considered. This is important because if market speculators are making a decision now, but think that the decision may actually be made later, then they have to worry about what they might know in the future that they do not know now.

A More Specific Proposal

We have now considered thirty-three concerns regarding the idea of using information aggregation abilities of speculative markets to fix information problems with democracy. Upon further examination, some of these issues seem more serious than others, and we have identified variations that might mitigate some of these problems. Let us now collect these various fixes together into a revised design concept for "futarchy," a form of government where we vote on values, but bet on beliefs.

This form of government would have a democratic component capable of approving bills, either by electing representatives who approve bills, or via some more direct democracy. This democratic component need not be as expert or time-responsive as existing democracies, and should instead be designed to be more robust to other problems, such as revolution or cycling. Bills approved would primarily define national welfare, how to measure it after the fact, and what agencies would further define and measure parts. Bills could also include treaties that coordinate welfare definitions across nations.

At any one time there would be a current official definition of national welfare. While new bills would be allowed to change that definition, the relevant agencies must eventually publish values for all past official definitions of national welfare where there remain bets regarding the current year. 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 in other institutions would not substantially and systematically change what values agencies eventually declare.

National welfare is to be defined primarily in terms of outcomes desired, though "outcomes" can include actions which are directly considered to be good or bad, beyond the other outcomes they produce. Courts should declare a national welfare definition change invalid if it seemed too

directly an attempt to implement specific policy whose rationale is based on other measurable outcomes.

There would at any one time be a set of nested status quo policies, with more fundamental policies authorizing less fundamental policies. This is similar to the way that a law today authorizing the creation of some government agency implicitly authorizes whatever regulatory rules that agency creates. This new form of government would begin with the previous set of laws and regulations as the status quo policies, except that the most fundamental policy would be this rule:

Futarchy’s Basic Rule: *When* a valid market clearly estimates that, conditional on approving a certain valid proposed policy, national welfare as defined today would be higher (compared with not approving this policy), *and* if no valid market clearly vetoes this by estimating that, conditional on approving this policy, national welfare as defined a certain delay later (e.g., a year) would be lower, *then* that proposal immediately becomes law, at one level removed from this basic rule, and overrides any conflicting non-constitutional laws.

Each valid proposed policy declares a specific short time window during which it can be approved, and says, directly or indirectly, what people must do to comply with it, what happens if they do not comply, and what previous policies it invalidates. Such a proposed policy can also specify how and by whom it is to be interpreted. Courts need only fill in any remaining ambiguities, and rule on whether the policy is valid. Valid policies do not violate any constitutional guarantees, such as a bill of rights, and may not substantially and systematically influence how welfare is defined or measured.

At the basic level there are some very conservative rules regarding what is a valid market, when such a market clearly estimates a price difference, and how long is the veto delay. These rules say who can propose policies with what fees and rewards, who can bet on them and how much they can bring in, what assets and markets count for determining prices, and how long and distinctly prices must differ in order to say that one price is “clearly” above another.¹⁰ Futarchy’s basic rule is allowed to approve other policies which look like the basic rule one level removed, except that they have different answers to these questions.

While not yet specific enough to officially propose, this proposal may be specific enough to evaluate in more detail.

Conclusion

This paper has explored the idea of using the successes of speculative markets as information aggregation mechanisms to help reduce the failures of democracies as information aggregation mechanisms. In this new form of governance, we could formally defer to betting markets on matters of fact, while retaining democracy on matters of value. That is, we could vote on values

¹⁰For example, these conservative rules might let anyone bet, let anyone propose a policy if they pay a \$100,000 fee (refunded ten times over if the policy is approved), require for approval a difference in CFTC-market-approved prices of at least twice both bid ask spreads averaged over each of two consecutive weeks, and have a veto delay of one year.

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, we considered thirty-three design issues with this concept, and finally sketched a more specific proposal to deal with some 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 a simple versions of both this concept and a more familiar democracy in simple information environments. Laboratory successes might then prompt larger trials, such as for important decisions in a corporation or other large organization. Eventually, we might consider trying agency-level decision making, such as with monetary policy, and later still, we might even be ready to consider changing a national form of government. Perhaps, at such a moment, we might even consider what a betting market has to say about how making this change would impact future GDP and other welfare measures.

References

- Abramowicz, M. (2004). Information markets, administrative decisionmaking, and predictive cost-benefit analysis. *University of Chicago Law Review*, 71(3).
- Ascher, W. (1999). *Why Governments Waste Natural Resources: Policy Failures in Developing Countries*. Johns Hopkins University Press, Baltimore.
- Aumann, R. (1976). Agreeing to disagree. *The Annals of Statistics*, 4(6), 1236–1239.
- Barker, A., & Peters, B. G. (1993). *The Politics of Expert Advice*. University of Pittsburgh Press, Pittsburgh.
- Barsky, R. B., & De Long, J. B. (1993). Why does the stock market fluctuate?. *Quarterly Journal of Economics*, 108(2), 291–311.
- Bell, T. (1997). Idea futures: Making the marketplace of ideas work. Tech. rep., Chapman University School of Law.
- Berg, J., & Rietz, T. (2002). Accuracy and forecast standard error of prediction markets. Tech. rep., University of Iowa, College of Business Administration.
- Berg, J. E., & Rietz, T. A. (2003). Prediction markets as decision support systems. *Information Systems Frontiers*, 5(1), 79–93.
- Besley, T., & Coate, S. (1998). Sources of inefficiency in a representative democracy: A dynamic analysis. *American Economic Review*, 88(1), 139–56.
- 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, B. (2001). Rational irrationality and the microfoundations of political failure. *Public Choice*, 107(3/4), 311–331.
- 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.
- Carlton, D. W. (1984). Futures markets: Their purpos, their history, their growth, their successes and failures. *The Journal of Futures Markets*, 4(3), 237–271.
- 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). Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Tech. rep. 1131, California Institute of Technology.
- CNN (1997). Poll: U.S. hiding knowledge of aliens.. <http://www.cnn.com/US/9706/15/ufo.poll/>.
- Cowen, T., & Hanson, R. (2002). Are disagreements honest?. Tech. rep., George Mason University Economics.
- Davis, J. A., Smith, T. W., & Marsden, P. V. (1996). General social survey.. <http://www.icpsr.umich.edu/GSS/>.
- 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.

- 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. *The Journal of Economic Behavior and Organization*, 39(1), 83–110.
- Gallup (1998). Gallup ufo poll: Some want to believe, some don't.. <http://www.parascope.com/articles/0597/gallup.htm>.
- Gallup (1999). Americans support teaching creationism as well as evolution in public schools..
- Geller, A. (2002). 'pained' webber: Geek tried to sink stock with cyber 'bomb'. *New York Post*.
- Group, B. R. (1994). Barna research online archives.. <http://www.barna.org/cgi-bin/PageCategory.asp?CategoryID=6>.
- 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), 16–19.
- Hanson, R. (2002). Disagreement is unpredictable. *Economics Letters*, 77, 365–369.
- Hanson, R. (2003a). Combinatorial information market design. *Information Systems Frontiers*, 5(1), 105–119.
- Hanson, R. (2003b). For savvy bayesian wannabes, are disagreements not about information?. *Theory and Decision*, 54(2), 105–123.
- 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. (2004). Manipulators increase information market accuracy. <http://hanson.gmu.edu/biashelp.pdf>.
- Hanson, R., Oprea, R., & Porter, D. (2006). Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization*, 60(4), 449–459.
- Hanson, R. D. (1997). *Four Puzzles in Information and Politics: Product Bans, Informed Voters, Social Insurance, and Persistent Disagreement*. Ph.D. thesis, California Institute of Technology.

- 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 information and the reward to inventive activity. *American Economics Review*, 61(4), 561–74.
- Holcombe, R. (1985). *An Economic Analysis of Democracy*. Southern Illinois University Press, Carbondale.
- 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.
- Kittlitz, K. (1999). Experiences with the foresight exchange. *Extropy Online*.
- Kraus, N., Malmfors, T., & Slovic, P. (1995). Intuitive toxicology: Expert and lay judgments of chemical risks. *Risk Analysis*, 12(2), 215–32.
- Kruglanski, A. W., & Freund, T. (1983). The freezing and unfreezing of lay-inferences: Effects on impression 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*, pp. 217–232. Elsevier Science Publications.
- Levy, B., & Spiller, P. T. (1994). The institutional foundations of regulatory commitment: A comparative analysis of telecommunications regulation. *Journal of Law, Economics, and Organization*, 10(2), 201–46.
- 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.
- Maddison, A. (1995). *Monitoring the World Economy 1820-1992*. OECD, Paris.
- 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). Information 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, 3–24.
- Oprea, R., Porter, D., Hibbert, C., Hanson, R., & Tila, D. (2007). Can manipulators mislead market observers. <http://hanson.gmu.edu/judges.pdf>.
- 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 America's 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.
- Plott, C. R. (2000). Markets as information gathering tools. *Southern Economic Journal*, 67(1), 2–15.
- Roll, R. (1984). Orange juice and weather. *American Economic Review*, 74(5), 861–880.
- Sachs, J. D., & Warner, A. M. (1995). Economic convergence and economic policies. Tech. rep. 5039, NBER.
- 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.
- Shiller, R. J. (1993). *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks*. Clarendon Press, Oxford.
- Shiller, R. J. (2000). *Irrational Exuberance*. Princeton University Press.
- Spencer, J. (2004). New icap-nymex derivatives have u.s. gas market's 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), 3–22.
- 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.

- Tullock, G. (1981). Why so much stability. *Public Choice*, 37(2), 189–202.
- Varian, H. R. (1998). Effect of the internet on financial markets. Tech. rep., University of California, Berkeley.
- 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.