Shall We Vote on Values, But Bet on Beliefs?

Robin Hanson *
Department of Economics
George Mason University†

July 2000

Abstract

A key question to ask about any social institution is how well it generates, aggregates, and distributes information. Speculative markets seem to do well at this, while familiar democratic institutions, relying in part on academic institutions, seem to fail in many ways. So perhaps we should consider "futarchy," a form of government where betting markets become our primary common source on matters of fact. Democracy would say what we want, while speculators would say how to get it. That is, elected representatives would formally define and manage an after-the-fact measurement of national welfare, while market speculators would say which policies they expect to raise national welfare. If we are willing to recommend policies that macroeconomic data suggest are causally related to GDP, it seems we should be willing to consider futarchy. Using an qualitative engineering-style approach, this paper considers thirty design issues with futarchy, and then presents a relatively specific proposal which responds to those concerns.

*For their comments, I thank Peter Boettke, Nick Bostrom, Tom Breton, Bryan Caplan, Roger Congleton, Tyler Cowen, Hal Finney, David Friedman, Karl Hallowell, Bernardo Huberman, Peggy Jackson, Hassan Masum, Peter McCluskey, Jim McKinney, Eli Lehrer, Karen Vaughn, Eliezer Yudkowsky, Richard Zeckhauser, and participants in the Monomedia Berlin: Value conference and the 2000 GMU Public Choice Outreach Seminar. I thank the Mercatus Center for financial support, and Edward Stringham for research assistance.

†*rhanson@gmu.edu* http://hanson.gmu.edu 704-993-2326 FAX: 704-993-2323 MSN 1D3, Carow Hall, Fairfax VA 22030

Introduction

A fundamental social problem seems to be this: which social institutions are better at generating, aggregating, and distributing knowledge? This is a key question at three levels: economics, politics, and abstract information.[1] That is, this question seems to be a key issue for understanding which economic institutions do better, which political institutions do better at regulating them, and which abstract information institutions, like academia, are better at informing ordinary citizens about economic and political institutions.

If we call this topic "social epistemology," then we should want our position on social epistemology to be consistent across these three levels. For example, a standard position is that current economic institutions do best, that current political institutions best regulate them, and that current academic and media institutions reliably tell people these things. A common alternative position is that academia, while usually reliable, has misestimated the harms that democratic interventions inflict on economies, a mistake that advocates within academia will eventually overturn.

This paper will explore another position, that betting markets are in many ways superior to familiar academic-style institutions for aggregating and distributing knowledge on many important topics, both specific and abstract. If so, perhaps those who disagree with the current academic consensus on political institutions should work to have betting markets displace such academic institutions, at least in part, as sources that ordinary people rely on for evaluations of political institutions.

This third position also suggests that we give betting markets a more direct role in the way political institutions evaluate the promise of economic interventions. While there are some more limited ways in which this could be done, this paper will explore an extreme proposal, to see just how far we can take this idea. After reviewing the centrality of information institutions and how democracies may have failed while betting markets have succeeded as such institutions, this paper will present an "engineering-style" proposal for a new form of government called "futarchy," and will discuss thirty issues and objections to the concept.

In futarchy, the public would vote on values, but bet on beliefs. The democratic process would be limited to managing an after the fact measurement of "GDP+", a measure of national welfare that they would define. The basic rule of government would be that when a betting market estimates that some proposed policy would increase expected GDP+, that proposal becomes law. Democracy would still tell us what we want, but betting markets would tell us how to get it. If we are willing to recommend policies that macro data suggest are causally related to GDP, it seems we should be willing to consider futarchy.

[1] In this paper, "knowledge" and "information" are treated as synonymous.

Information Institutions

> The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use ... exists ... solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess. ... Which of [central planning or competition] is likely to be more efficient depends mainly on whether we are more likely to succeed in putting at the disposal of a single central authority all the knowledge which ought to be used but which is initially dispersed among many individuals, or in conveying to the individuals such additional knowledge as they need in order to enable them to fit their plans in with those of others. (Hayek, 1945).

Economic theorists today wholeheartedly accept Hayek's once-revolutionary claim that an uneven distribution of knowledge, now called asymmetric information, is the key problem in social systems. They also accept that policy interventions are mainly constrained by a very wide dispersion of relevant knowledge, and by the fact that those who know may not want to tell. Most even accept Hayek's claim that unregulated competition is better than pure central planning. Most economists today do not, however, accept Hayek's stark choice being these polar extremes. Hayek's argument that our choice is effectively binary was based on human cognitive biases toward the seen over the unseen:

> When we decide each issue solely on what appear to be its individual merits, we always over-estimate the advantages of central direction ... [Thus] freedom can be preserved only if it is treated as a supreme principle which must never be sacrificed (Hayek, 1973).

Most economists, however, are not persuaded that they suffer from such severe cognitive biases.[2]

Instead, most economists now study what can go wrong in economies with asymmetric information, and how they might best intervene given their knowledge constraints. That is, they study how to pursue policy goals when one understands the overall structure of some policy area, but is ignorant of many details. And few economists believe that the exact best policy is always exactly zero intervention.[3]

An institutional critique of this perspective is that our choices are much more limited. Ordinary citizens can not choose from all possible intervention policies; they can at best only choose from feasible stable *institutions*, which will then choose more specific policies.[4] The question then

[2] Many economists instead argue that cognitive biases are *against* government intervention.

[3] One might deny that we can know much about the overall structure of an area of life. But at best this argues for using more abstract economic theory, which makes fewer specific assumptions.

[4] A "constitutional" critique appears to go further, arguing that we can at best choose among constitutionally-defined forms of government, which will then choose more specific government institutions.

becomes which institutions, including a "no intervention" institution if that is feasible, produce better policies on average.

Public choice critiques of government intervention argue further that familiar high-intervention democratic political institutions suffer greatly from problems like rent-seeking, coordination failures, commitment failures, adverse selection, and especially low incentives for voters to become informed (Holcombe, 1985). Defenders of democratic interventions, in contrast, suggest that such problems seem to be manageable in practice, as rational agent theory suggests that they should be (Wittman, 1995).

Whichever side has the better arguments, such arguments seem too complex for ordinary citizens to evaluate personally. So most citizens must again choose institutions, instead of specific policies. That is, ordinary citizens with little time to directly study political institutions must instead defer to some indirect information sources for answers to such abstract questions. And since citizens also have little time to evaluate the reliability of information sources, a handful of such sources must serve them for all such topics. Thus at the abstract information level, citizens must again typically choose between a few broad long-lived institutions, perhaps as specific as particular academic disciplines or news media, but more likely academia and mainstream news media taken as wholes.

Academic opinion on the effectiveness of political institutions has for many decades seemed to lean more toward democracy's defenders than its detractors. Mainstream media has on the whole leaned the same way. Critics who think democracy results in too much intervention, as well as those who think it results in too little, often explain this mistaken judgment as due to various institutional failures in academia and news media, failures which reduce the informativeness of any current consensus (Yeager, 1997; Rosen, 1997).

Many academic studies do in fact suggest real failures in academia-style institutions (Redner, 1987), such as universities, research labs, and expert advisory committees. In statistical studies, these failures include overconfidence in variance estimates (Henrion & Fischhoff, 1986), strong unacknowledged selection biases (Long & Lang, 1992), and very high error rates (Dewald, Thursby, & Anderson, 1986). In peer review, failures include very low levels of agreement (Chubin & Hackett, 1990), strong biases against obscure institutions (Peters & Ceci, 1982), and biased objections to methods that give disliked results (Mahoney, 1977).

Critics, however, mostly continue to work to change academic opinions, apparently accepting academia as the appropriate forum for such debate. (Even Hayek granted that "a body of suitably chosen experts may be in the best position to command all the best [scientific] knowledge available" (Hayek, 1945).) The alternative information institutions that critics offer, such as specific factions of academic disciplines (e.g., Austrian economics) or specific "ideological" news media, seem far too small to gain a reputation among ordinary citizens as reliable information sources.

If the currently dominant information institutions suffer from serious failures, however, critics of

standard views on political institutions might do better to promote better information institutions. After all, even if they succeed in changing academic consensus, their criticisms of academia could then come back to haunt them. In contrast, they should expect widespread acceptance of better information institutions to make their political views more likely to be widely accepted. They should also find common cause in creating better information institutions with those who hold other non-standard views. Finally, since economic and political institutions *are* in part information institutions, better information institutions may also directly allow better economic and political institutions. This possibility should be particularly interesting to those who think current political institutions do poorly as information institutions.

To put it another way: a fundamental social problem seems to be *social epistemology*. Epistemology has traditionally been concerned mostly with the conscious strategies that a rational person should use when deciding what to believe given primary sources of evidence. Most real problems in deciding *what* to believe, however, concern how the social institutions people are embedded in influence *who* we should believe. Social epistemology was first defined as,

> How should the pursuit of knowledge be organized, given that ... knowledge is pursued by many human beings ... each equipped with roughly the same imperfect cognitive capacities ...? (Fuller, 1988)

For our purposes, we can take social epistemology to be the study of how effective various social institutions are at generating, aggregating, and distributing knowledge and information. This seems to be the central subject of economics, politics, *and* of abstract persuasion regarding such topics. We should prefer a *reflexive* position on social epistemology, i.e., a position consistent across economics, politics, and the crucial subject of why ordinary people should believe this abstract position.

For example, a standard roughly-consistent position is that decentralized markets do well as information institutions, that they do even better with limited regulation, that familiar democracy does sufficiently well as an information institution to provide such regulation, and that academia is reliable enough as an information institution for ordinary people to rely on it to answer abstract questions about economics and politics. This position considers academia reliable enough to be given important political roles in deciding policy.

Another position, with perhaps more internal tensions, is that while academia is the most reliable known information institution which ordinary citizens can turn to for answers to abstract questions about economics and politics, it is not reliable enough to be given a large role in deciding policy, and it has made a serious error on the subject of democratic interventions into the economy. This position suggests that the academic minority that knows of this error should continue to push academia to correct it, and that until they succeed little else can be done to change the opinions of ordinary people.

Several other common positions could be described, but this paper will instead present a new position, that another general information institution is superior to current academic and news media institutions at many aspects of delivering policy-relevant answers to abstract questions, that the relevant academic research endorses this conclusion[5], and that the reason ordinary citizens have not yet been persuaded of its robust reliability is both that technology made it very cheap only recently, and that regulatory error has kept it illegal. This regulatory error is of a sort one expects to be common, a broad activity ban intended to reduce a "seen" problem, which then prevents an "unseen" new industry from developing.

This alternative institution is *betting markets*, which are decentralized markets that avoid many of the perceived disadvantages of markets in general. Betting markets can not only be relied on to tell us what sorts of economic interventions will get us the things we want, but can also be formally embedded in our political institutions in order to directly and officially fulfill that role.

One refinement of this position says that economic interventions are almost never a good idea, and that a better abstract information institution will reveal this fact. Another refinement, however, says that a better information institution will reveal that we want even more intervention than we see now. *Both* of these positions have a common interest in promoting a better information institution.

Information Failures of Democracy

If familiar political institutions were effective at producing informed policy choices, there would be relatively little reason to seek better information institutions, either to imbed within political institutions, or to inform ordinary people about the effectiveness of their political institutions. So how well-informed is current democratic policy?

Half a century ago, empirical research on individual U.S. voters seemed to confirm the worst fears of democratic skeptics.

> 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).

[5]This position thus accepts academia as a reasonable forum in which to make the argument that another institution is better than academia.

Such high levels of ignorance continue today (Delli-Carpini & Keeter, 1989). For example, only 29 percent of U.S. adults can name their congressman, and only 24 percent can identify the first amendment.

Formal analysis found many problems with democracy as well, including instability, rent-seeking, coordination failures, and commitment failures (Holcombe, 1985; Besley & Coate, 1998). Formal models did not, however, as easily validate what many considered the most serious problem: low incentives for voters to become informed (Wittman, 1995).

In theory democracies with centralized responsibility could be very well informed if voters just determined policy by voting selfishly and retrospectively, i.e., voting for the incumbents only when their personal lives seemed better than expected. However, voters 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 broader policy consequences.

Democracy seems reasonably capable of achieving this, however, at least in theory given rational voters. After all, voters learn many things as a side effect of just living, large elections 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 and other sources regarding how to cast their votes.

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

This is not to say there are not information problems with democracy, however. The above quote goes on to say,

> 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 goes on to claim 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, can plausibly explain the fact that time series of

%	In U.S. Agree With This Opinion	Cite
47	God created humans in basically present form in last 10,000 years	(Gallup, 1999)
52	'astrology has some scientific truth'	(Davis et al., 1996)
72	believe in angels	(Gallup, 1998)
80	'US government 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

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 many specific contrarian public opinions, such as those listed in Table 1. These opinions are not just cases where the public is ignorant of expert opinion, such as they might be on foreign trade. Instead, these are cases where the public seems largely aware that expert opinion disagrees with them.

Contrarian public opinions suggest not only ignorance or a failure to share information, but also irrationality more directly. In theory, rational agents with common priors, who would have the same estimates were it not for having different information and computation, should use the fact that someone disagrees with them in a certain direction to update their beliefs. As two rational agents alternate telling each other their new opinion, each one should not be able to predict in which direction the other will disagree with them next time, and the only stable endpoint of this process is to eliminate such disagreements (Aumann, 1976; Hanson, 1997). Yet people seem to persistently disagree on factual claims, especially on political topics, even when these people seem well aware of the directions that others disagree. This lends some support to theories that try to explain otherwise puzzling policies and political behavior in terms of weak but positive voter preferences for irrationality (Caplan,).

Of course misinformed and irrational voters need not prevent informed policy, if voters allow policy to be determined by informed elites such as 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 the public does not defer to experts in many areas. Furthermore, many case studies suggest that when governments do use academic experts, they do so frequently for legitimation of predetermined policies, rather than for information to help determine policy (Barker & Peters, 1993).

Even when democratic policy does defer to academic experts, the failings of academic-style institutions may lead to poorly-informed policy. In academic-style policy information institutions,

such as government research labs, advisory agencies, or advisory committees, experts simply declare their policy advice, without much in the way of clear or direct incentives to be right. Even when such experts are selected on the basis of their "academic reputation," this basically means they are selected by other experts, and those other experts again have little clear or direct incentive to choose those who would offer accurate policy advice. (Data suggesting real failures in academia was mentioned in the last section.)

If familiar political institutions suffer from information failures, how large are those failures? The huge variation in economic growth rates across nations, depicted in Figure 1, suggests that the effects of information failures may be very large (Maddison, 1995). There are surely some random and uncontrollable reasons why some nations are rich and others poor. An important fraction of the variation, however, 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 have serious failures 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 are bad. Thus at some level bad policy seems to be fundamentally due to a failure to aggregate and distribute relevant information.

Information Successes of Speculative Markets

While democratic policy seems to suffer from serious information failures, speculative markets have shown some dramatic information successes.

Most markets for stocks, bonds, currency, and commodities futures are called "speculative markets" because they allow people bet on future prices by buying or selling today in the hope of later reversing such trades for a profit. Such opportunities for speculation 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 their money in previous bets.

"Betting markets" are speculative markets which 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.

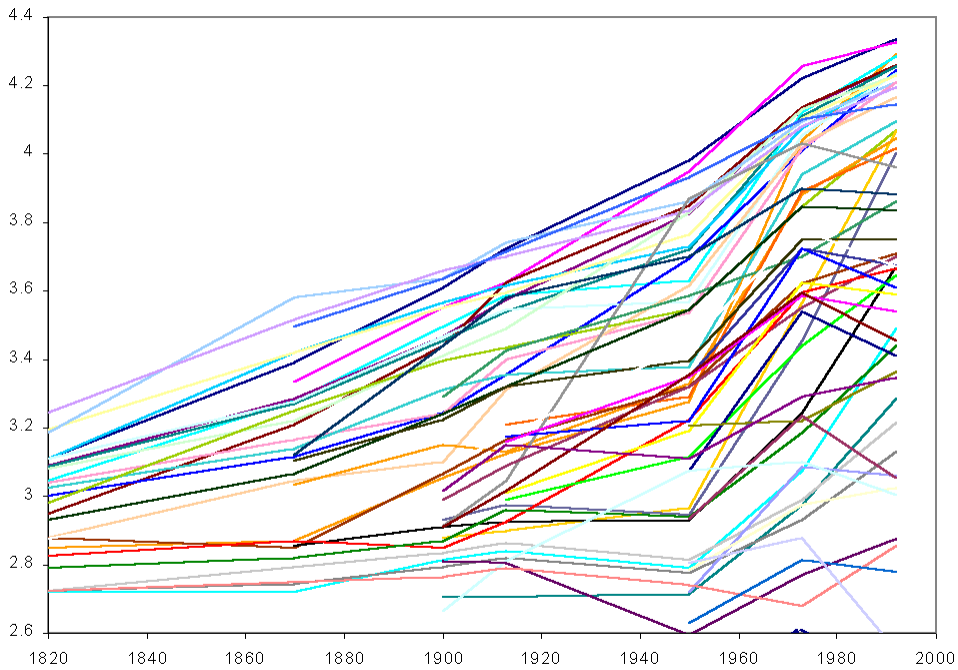

Figure 1: Log GDP Per Capita vs. Year For 56 Countries

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 apparent exceptions seem to be long-term aggregate price movements in real asset markets, and a long-shot bias in high-transaction-cost betting markets.

Compared to what seems rational, asset markets seem to have too much long-term aggregate price variation, such as stock market "bubbles" (Shiller, 2000). These price movements are where risk and delay most reduce the extra returns to speculators for correcting pricing errors, and where theory says that "noise" traders can actually gain superior returns (though not utility) from their irrationally-large risk-taking (Long, Shleifer, Summers, & Waldmann, 1990). Long-term aggregate prices, however, are also where it is hardest to empirically distinguish irrationality from rational shared information about changes in the nature of the economy (Barsky & Long, 1993), and where selection effects most pollute the available data (Jorion & Goetzmann, 2000). So it remains unclear just how irrational long-term aggregate price movements are.

Speculators have a much stronger incentive to eliminate price biases that vary over shorter terms and that are independent of the few aggregate price dimensions that command a "risk-premium."

High transaction costs, however, can deter speculators from eliminating such biases. Parimutuel betting markets, for example, typically charge about an 18% transaction fee, and also typically display a "long shot bias," over-estimating the chances of unlikely outcomes (Hausch et al., 1994). This bias occurs because those who favor a low probability outcome need to spend proportionally less money to have the market price reflect their views, and transaction costs deter speculators from correcting this bias.

Over the last few decades economists have also studied speculative markets in laboratory experiments, where they have more control over trader information and preferences. The name "Hayekian hypothesis" has even been given to the roughly confirmed thesis that speculative markets aggregate information well even when traders know little about their environment and other traders (Sunder, 1995). Experimenters have found, for example, that 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, since that allows more kinds of arbitrage, i.e., risk-less profits from correcting market prices.

The key policy question about any institution is how it performs *relative* to alternative institutions in the same situation or environment, or averaged over a given range of environments. A few studies have looked at this empirically, directly comparing speculative markets with other 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). Betting markets beat opinion polls at predicting U.S. election results (Forsythe, Nelson, Neumann, & Wright, 1992). Finally, betting markets consistently beat Hewlett Packard official forecasts at predicting Hewlett Packard printer sales (Chen & Plott, 1998; Plott, 2000).

Unfortunately, no studies have directly compared estimates from speculative markets to estimates from academic-style institutions (though one has been proposed (Hanson, 1995a)). 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 place a great deal of weight, perhaps too much, on the opinions of experts relative to more ordinary people. One might worry that betting markets where anyone could join would place too little weight on experts relative to ordinary people. We know, however, that speculative markets seem to, if anything, put too much weight on advice from experts, both public

and private (Figlewski, 1979; Metzger, 1985; Lichtenstein, Kaufmann, & Bhagat, 1999). Thus one cannot reasonably claim that while academic institutions may place too much weight on experts relative to ordinary people, betting markets place too little weight. Speculative markets thus seem to do better at this weighting choice.

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 unbiased on average. These are traders who make offers that others accept, rather than accepting offers made by others, or making offers that others do not accept. Compared to other traders, market makers invest twice as much, trade more, earn higher returns, and make one sixth as many errors. They also tend to be more highly educated, and experienced at trading (Forsythe et al., 1992; Forsythe, Rietz, & Ross, 1999). Betting markets seem to beat opinion polls and other competing institutions because of the disproportionate influence such markets give to more rational and informed traders.

The number of topics on which markets can create estimates is limited by the number of markets one can create. There are fixed costs to create and run markets, and dividing attention among more markets also raises trader liquidity and volatility costs, i.e., costs due to difficulties in finding trading partners and fluctuations in prices (Pagano, 1989). 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 that markets with a couple of traders are viable. Very low internet mechanical transaction costs are also now spurring a burst of innovation exploring a great many new market forms, many of which are small (Varian, 1998; Shiller, 1993). 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)).

It remains widely illegal, however, to create real money markets like these play money markets, and so most speculative markets still trade millions of dollars a year. This regulatory block on financial innovation should not be surprising, however, 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).[6]

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 anywhere near enough about either type of institution to be proposing radical new forms of governments 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 empirically-supported theory. A scientist bases policy recommendations only on relatively direct data, or on well-supported theory. A scientist who studies systems tends to assume that existing systems are functional, and uses that as data to refine theory. A scientist therefore 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 about.

An engineer, on the other hand, is more interested in improving systems than in improving theory. An engineer is thus willing to make cruder judgments, farther removed from 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 explicitly as far as it will go, but also uses theory-informed intuitions to more informally think about a wide a range of design issues. An engineer then typically moves on to a series of increasingly realistic and expensive "proof of concept" *prototypes*, from computer simulations, to laboratory "wind tunnel" models, to field tests. While scientists have little use for prototypes and their tests, being neither basic theory nor data that tests theory, prototypes are what make the engineers' world go round.

Social science is now mostly dominated by a scientific, rather than an engineering, style. The most respected proposals for new economic and political institutions are thus based firmly on established theory. For example, authorities were persuaded to consider new institutions such as pollution emissions trading and new forms of auctions largely on the basis of theoretical arguments and endorsement by economic theorists. Efforts to test prototypes of such institutions, in contrast,

[6]General betting has been legal, if highly regulated, in England for several decades now. The number of betting topics, however, is limited because bookies are expected to set stable odds, rather than allow odds to fluctuate and so be set by bettors. Bookies thus mainly allow questions they feel they can afford to estimate well (Sharpe, 1994).

have gained relatively little academic prestige.

Economic theorists of "mechanism design" and "implementation" have considered, often in quite some generality, what complex mechanisms would let you achieve your arbitrary policy goals by inducing people to tell you what they know. Assuming that human behavior is exactly described by standard game theories, such theorists have found that it often seems possible to get people to tell everything they know relevant to achieving almost arbitrary policy goals.

For example, regarding a simple "public goods" decision, i.e., should the government implement a project and if so who should pay how much, one should in theory be able to get people to tell everything relevant they know as long as their information is correlated in some way, and no one runs out of money (D'Aspremont & Gerard-Varet, 1979). More generally, it should be enough that policy makers are just more patient than ordinary people (Kalai & Ledyard, 1998), or that anything any one person knows is also at least implied by a combination of things other people know (Postlewaite & Schmeidler, 1986; Duggan, 1997), or that the probability of some observable event varies depending on each relevant thing people know (Bergin & Sen, 1998). Few economists, however, believe that these complex mechanisms would work with real people, and at least one experiment has supported this expectation.

Regarding political institutions, formal voting models have supported alternative voting rules such as "approval voting," where voters can vote for as many or few candidates as they like (Brams & Fishburn, 1983). Formal models have also supported several simple institutions for making public goods decisions (Tideman & Tullock, 1976; Groves & Ledyard, 1977), one of which has even done well in some laboratory experiments (Chen & Tang, 1998).

A few social scientists have informally proposed more radical institutional changes, such as random selection of legislators (Carson & Martin, 1999; Callenbach & Phillips, 1985; Burnheim, 1985), or allowing fine grain private choice of legal systems (Friedman, 1989). Prototypes of these proposals have even been explored to varying degrees. But none of this has gained much academic prestige.

The radical new form of government proposed in this paper, "futarchy," is intended to be taken in the engineering spirit. While motivated in part by theory, it seems difficult to use established theory to in non-trivial detail compare futarchy to existing institutions, and in any case no such models are presented here. The purpose of this paper is not to induce high confidence in readers that futarchy 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.

Proposal Concept

Directly or indirectly, academic-style institutions are today important information institutions that ordinary people defer to, when they defer to anyone, for answers to abstract questions about political institutions and policy.[7] The resulting democratic policies often seem less well-informed than they could be. Yet betting markets seem to do a remarkably good job of aggregating information. So perhaps we should consider substituting betting markets for academic-style and other familiar institutions in answering abstract questions about policy. There are four levels at which we might consider such substitution: parameter advice, decision advice, agency decision control, and entire government control.

Parameter Advice

At the simplest level, policy makers could get into the habit of deferring to betting markets, instead of commissioning studies or expert committees, to predict parameters such as budget surpluses, lifespan changes, trust fund depletion rates, global warming rates, AIDS infection rates, or probabilities of asteroid strikes. The elites who would have otherwise been on the elite committee, or influenced its conclusions, would instead have to gain influence by persuading market speculators. This general idea of using speculative markets more widely for information aggregation has been discussed before, both in journals (Hirshleifer, 1971; Leamer, 1986) and in science fiction (Brunner, 1975), as has the more specific application of aggregating policy-relevant information (Zeckhauser & Viscusi, 1990; Hanson, 1995b).

The thicker a market is, the more offers and trades it has, and thus the harder it is for one person to change the price. In order to ensure that such speculative markets are thick enough to have good estimates, one could either directly subsidize a market maker, or credibly commit to having the market price influence policy. One can subsidize a market maker in a way that eventually gives money to whomever turns out to be right (Hanson, 1992). This can allow even a single other trader to be paid to create informed estimates on some topic. We will later show that an influential market can not be very thin.

Decision Advice

Since one can ask betting markets to estimate just about any parameter that one might reasonably measure after the fact, one can ask betting markets to make conditional estimates. So, for example, not only could a betting market estimate future sea levels, it could also estimate sea levels conditional on future CO2 levels. (One way to do this is with called-off trades, i.e., undoing trades if

[7] Most of this deference is indirect, via intermediating institutions such as news media.

the condition is not met.) And since the choices we make can be treated as conditions, we can ask betting markets to estimate the observable consequences of decisions. We could, for example, ask betting markets to estimate murder rates conditional on passing gun control laws, and conditional on not passing gun control laws. If market estimates of murder rates were clearly higher given gun control than given not, that would argue against passing gun control laws.

Thus, at the decision advice level, policy makers could get into the habit of deferring to betting markets, instead of commissioning studies or expert committees, to estimate the consequences of policies. Markets might estimate student test scores given school vouchers, lifespans given universal health care, attendance if a new stadium is built, or the chance of war given a treaty. In general, one would ask a market to estimate an outcome of interest, conditional on choices one might make. This idea has been described before, both in journals (Hanson, 1999) and in science fiction (Stiegler, 1999).

Agency Decision Control

The above approach could help if government agency decisions are uninformed due to agencies lacking access to relevant information. Agency heads, however, might instead lack incentives to make their decisions reflect available information. They might get away with this behavior if legislators have limited abilities to oversee the agency, or if voters have limited abilities to oversee legislators. To deal with this problem, we might put betting markets more directly in control of policy.[8]

For example, monetary policy seems to be a policy area where people largely agree on a few statistics by which one can tell, after the fact, how effective a policy was, and on a few control variables one might use to get good outcomes. Policymakers try to achieve high and steady GDP and employment, and they control short term interest rates and reserve requirements. The dispute seems to be mostly about what control variable choices, in what contexts, lead to good outcomes. Thus monetary policy seems a good candidate for more direct "futarchy," or government by betting markets.[9]

In monetary futarchy, the government would choose some explicit function describing good outcomes. This function would presumably be increasing in future GDP and employment, and be concave in these in order to reward stability. The rule of monetary policy would then be that we would raise or lower interest rates whenever a betting market clearly estimated a higher value for the good outcome function conditional on such a change. Information that is now given privately

[8] We later discuss how non-in-control betting markets can also be biased due to fears that decision makers have private information.

[9] The name "futarchy," comes from government by futures markets, since one can thing of bets as claim-judgment futures.

to monetary officials would instead be made public, and hence available to market speculators.

Other possibilities for agency futarchy would be to approve drugs based on estimates of costs and lived saved versus lives lost, to set crime policy based on estimated costs and crime rates, or to set organ transplant policy based on estimates of costs and life years gained. In general for each agency one would officially choose some value function that describes the goal of that agency, and then officially let market speculators determine which policies best achieve that goal.

Entire Government Control

A more ambitious level at which we might substitute betting markets for other information institutions is that of controlling an entire government. We would vote on values, but bet on beliefs. That is, elected representatives would mainly define and oversee the ex post (after the fact) measurement of some "GDP+," a measure of national welfare. The basic rule of government would be:

> When an approved betting market clearly estimates that a proposed policy would increase expected GDP+, that proposal becomes law.

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

Futarchy seems to be a promising form of government, *if* we accept three assumptions:

1. It is not that hard to tell rich happy nations from poor miserable ones.
2. Existing democracies fail primarily due to not aggregating available information.
3. Betting markets are the best known institutions for aggregating information.

Most empirical research in growth economics seems willing to presume that policies that induce high GDP are good policies. Of course we know that GDP leaves out many things we value, and is often imperfectly measured. But it seems hard to believe the errors are very large; 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. Thus researchers have typically recommended policies which their research suggest causally raise GDP.

But this same attitude also suggests that GDP-futarchy, i.e., a government which does whatever betting markets say would raise measured GDP, would do well for most citizens. More specifically, imagine that national welfare were defined as a few-percent-annually-discounted average of future log GDP, with GDP defined and measured as it is now in the U.S. or similar countries.

It seems plausible that most citizens of most nations would do better under this GDP-futarchy than under their current forms of government. If nations are typically poor because of adopting known-to-be-bad policies, then GDP-futarchies which grow their GDP as fast as the best-informed

policies allow should become rich. And futarchy policy would be well-informed, both because in betting markets those who know they are ignorant have strong incentives to stay away, and because those who do not know they are ignorant will decline in influence as they lose their money.

If a GDP-futarchy would do well, then futarchy based on GDP+ should do even better. That is, if elected representatives, freed from thinking about how to get what they want, can develop substantially better ex post measures of what citizens want, then most citizens should be even better off. And since economists spending modest budgets have already found many promising extensions to current GDP measures (Boskin, 2000; Nordhaus, 2000), we might do even better with both increased funding and the full attention of elected representatives.

Futarchy is intended to appeal to a wide range of ideologies and political factions, remaining neutral on most such disputes. Futarchy could in principle result in either extreme socialism or extreme minarchy, depending on what elected representatives say they want, and on which policies speculators believe will achieve those goals. Futarchy does require that there be some community of speculators who have assets they are free to bet, and that some extra benefits consistently accrue to those who win their bets. These benefits need not be large, however, nor need the community need be very large. Futarchy should appeal most to those who can accept the values democracy would favor, and who think that speculators could be persuaded to agree with their assessment of how to best promote such values.

Design Issues

Many ideas seem promising at first glance, but then seem less so after one considers some details. Let us therefore, in engineering fashion, identify and briefly discuss thirty design issues, each expressed as an objection to this overall approach. After this we will present a more detailed proposal for "futarchy" as a form of government.

The Rich Would Get More Influence

A basic fact of life is that in general the rich get more influence. After all, the point of being rich is to be able to get more of whatever it is you want. If we do not like that, we can tax the rich more, and then they will not be as rich.

Perhaps the concern here is that the rich should not have proportionally as much influence in betting markets as we accept them having elsewhere. Betting markets are not opinion polls where the rich get more votes, however. The rich have more *potential* votes, but they should lose money if they use those extra votes without having proportionally extra information. And if there's one thing we know about the rich it is that they don't easily throw their money away; those who do,

do not stay rich very long.

It Would Be Better To Appeal To Higher Motives Than Money

If one could reliably pay off bets in units of "higher" forms of value, then we might want to do so. We want to induce people to tell us what they know, and we need to offer them something that we can reliably produce, can distribute in controlled amounts, and know that most people value. Money satisfies these requirements, and no other forms of value yet seem to. We do not seem to know how to design productive information institutions based on appealing to other values.

People Would Volunteer Less information And Lie More

Consider someone who holds information relevant to a policy choice. Under familiar political institutions it may seem that such a person often has little way to benefit from their information other than to honestly reveal it because they feel it is their civic duty. If a betting market were available, however, such a person might be tempted to keep quiet about their information in order to benefit by betting on it. They might even publicly lie in order to deceive other traders.

Of course betting reveals information directly to markets, and once people with secrets 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 could reverse their earlier trades for a profit. Secret-holders might, however, never think they have run out of ways to bet on their information.

This is indeed a reasonable concern, at least in theory. The studies mentioned earlier showing that speculative markets beat other institutions at aggregating information suggest that institutions which now try to rely on weak incentives and civic duty must have other serious drawbacks. But perhaps a new institutions will one day be discovered to better tap this potential.

People Prefer Their Comfortable Illusions

The beliefs that speculators would use to estimate policy consequences would probably not be the same beliefs as the comfortable illusions that voters prefer. So if futarchy forced voters to accept specific speculator beliefs, it might thus force them to forgo their comfortable illusions.

Fortunately it seems that voters need not be directly confronted with speculator beliefs, and therefore need not forgo their illusions. Voters need to keep abreast of news about government policy, and stay informed about choices made by their elected representatives, but this can all be done without becoming much aware of speculator beliefs.

Speculator estimates of welfare that determine policy would not directly explain speculators' reasoning, and specific approved policies could typically be rationalized in many ways. Other markets and forums might more concretely explain speculator reasoning, but news media that catered

to viewers with cherished illusions would likely not confront their viewers with uncomfortable beliefs. Even today, media avoid telling viewers how legislative sausage is made, when viewers find that distasteful.

Maybe Low Information Is Not Democracy's Major Failing

Futarchy has been designed primarily to fix information problems with democracy. But perhaps, contrary to the impression one might get from the evidence presented earlier, democracy's biggest problem is some other problem, such as a coordination or commitment failure.

In this case the big question would be whether futarchy makes that bigger problem any worse, or whether the changes one might want to make to fix that problem are incompatible with futarchy in some way.

Democracy Might Become Unstable

In theory, democracies seem to have problems of policy instability, or "cycling," since for any policy there is always another policy that some majority likes more. In practice, however, this does not seem to be a large problem. If voter ignorance is part of the reason for such stability, as some theories suggest (Hanson, 1997), perhaps futarchy would make instability worse, by better informing citizens about the consequences of their choices.

This might be a reason to lean toward more stable democratic institutions within futarchy. So we might prefer one or more large decentralized legislatures whose members change gradually, 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 the consequences of policy, and if less expertize is required to make value judgments than to make full policy choices, then we might reasonably accept a democratic part of futarchy that is slower and less expert than today's democratic processes. Thus futarchy might allow more experimentation with non-other democratic forms, such as demarchy or more direct democracy.

Time-Consistency Might Be A Problem

The inability to commit has been identified as a major reason, in theory and in at least some practice, for government failures (Levy & Spiller, 1994; Besley & Coate, 1998). The basic rule of futarchy described above does not seem to overcome this problem, but it also does not seem to make this problem any worse.

While a constitution might declare that the government could commit itself to future choices, it is not clear that courts could be relied on to enforce such commitments. It is also possible that

allowing earlier generations to take advantage of later generations creates worse problems than commitment failures. Perhaps we should keep the standard approach, which is to typically allow governments to commit only if they can arrange to do so themselves, such as via transaction costs, posting bonds, external reputation, or other external institutions.

Expressive Voting Could Still Cause Problems

Voters often seem to vote expressively, i.e., to care about other things when voting besides influencing what policies the winners will promote (Brennan & Lomasky, 1993). For example, voters want to take sides, to show themselves and their friends that they are knowledgeable and that they care, and to have people they like and respect represent them. And while these may be worthwhile goals, policy choices can suffer as a result.

Futarchy, by limiting democracy primarily to values, would presumably turn voters more toward showing that they care, and away from showing that they are knowledgeable. It is not clear that futarchy makes expressive voting any more of a problem overall, but if this is a concern, we might consider forms of democracy that are harmed less by expressive voting, if such can be identified.

Markets Might Be Too Thin To Make Good Estimates

In general, there are an infinity of possible markets, and so finite costs of creating markets ensure than most possible markets do not actually exist. And even if one created some arbitrary market, there is no guarantee that people would trade much there; the market might end up being too "thin."

A market that influences important policies, however, cannot be very thin. After all, if a very thin market were to influence policy, interested parties could then by definition 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 was willing to spend much more, then speculators who knew that fact could make easy profits by trading against that better funded side, again thickening the market. So either way, the market would end up being thick.

People Could Buy Policy Via Betting Markets

Imagine that speculators knew that a certain interested party had a strong interest in a price moving in a certain direction, but had no special information about that topic. If this interested party then began trading large amounts in order to move the price, speculators would see easy profits in moving the price back, which they would do given low transaction costs.

Thus the only way for such an interested party to substantially influence the price is to possibly have substantial special information, or to trade via parties that might have such information. And this means that all the interested party can really do is choose whether to hide or reveal some private information they hold.

For example, assume some proposed city stadiums would be profitable, while others would be unprofitable, and that many speculators know how profitable proposed stadiums are on average. For each stadium, assume only few interested insiders know how profitable that stadium would be, and that these insiders would benefit from the stadium being built even if it were not profitable.

Imagine first that speculators initially drove market profitability estimates for all proposed stadiums to the average value. Then imagine that many insiders tried to buy a favorable impression of their stadium by bidding up its market estimate of profitability. Speculators could then profit by driving some prices back down so that the average price returned the known average stadium profitability.

If it was known that all insiders bid up their estimate equally, speculators would push all estimates back to the same average value. Speculators would only let estimate differences stand if they expected a correlation between insider trading and stadium profitability. If insiders tend to push more when they know their stadium is more profitable, they would thereby reveal some information through their trades, and market estimates would correlate with stadium profitability. But if there were no correlation between insider trading and profitability, all estimates would end up being the same.

Therefore, as claimed, in this case all that insiders can really do is either reveal their inside information, or try to hide it. And it turns out that all cases are basically of the same form as in this stadium example. When insiders face speculators with deep pockets and low transaction costs, they can only influence market-estimate-determined policy by revealing or hiding information.

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. If most all those speculators were therefore confident that Bill Gates were wrong about some estimate, even Bill Gates could not substantially change a betting market estimate. If Mr. Gates pushed the price one way, all those speculators would push it right back.

More likely, however, speculators would allow Mr. Gates to move the price some because they reasonably suspected that he had access to hidden information. And given his previous successes, this would not be an unreasonable assumption.

You Need A Way To Tell If A Proposal Was Implemented

It seems reasonable to treat proposals like contracts, in the sense of being responsible for indicating how ambiguities should be settled. A contract typically specifies a jurisdiction whose law is to interpret it, and can also specify other rules of interpretation or the arbitrator who is to do the interpretation. Contract law remains in the background to fill in remaining ambiguities. Similarly, futarchy 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 and implemented, then speculators would have to average over the kinds of interpretations they think likely. An unclear proposal could be a sloppy but sincere attempt to improve national welfare, or it could be an insincere attempt to slip in policies that would not be approved if proposed clearly. Insincere proposals probably reduce national welfare, while sloppy sincere proposals might increase it.

Speculators must thus estimate the relative contribution of sincere and insincere proposals to the class of unclear proposals in any given policy area. When insincere proposals dominate, speculators should be wary of endorsing unclear proposals, just as anyone is wary of signing a contract he or she does not understand. But when sloppy but sincere proposals dominate, speculators may accept the occasional insincere proposal as the price to be paid for valuable and sincere but sloppy proposals.

A Policy Might Influence How Welfare Is Measured

If speculators expected a policy to raise estimates of national welfare via influencing how welfare is measured, they would approve that policy, even though it did not actually raise welfare. It is therefore important that welfare measurement be a relatively independent process, so that we can forbid markets from approving proposals that substantially bias the measurement of welfare.

Similarly, betting markets should not be allowed to approve policies that substantially bias the political process leading to changes in the welfare definition. Otherwise futarchy might, for example, naturally tend to approve proposals that prevent changes to the current welfare measure.

People Could Do Harm to Win Bets

Insurance companies have to worry about people deliberately setting fires in order to collect their fire insurance, or people killing their spouses to collect life insurance. In principle, public companies also have to worry about someone selling their stock short and then damaging their factory or product, such as by poisoning a few drug or food products on store shelves. This latter concern seems to be much less of a problem, however, because it is harder for most people to substantially damage an entire company, and easier to monitor the few people who can.

Ordinary people should find it even harder to substantially damage national welfare. If they could, however, the straightforward way to profit from this would be to bet against GDP rising

as much as other people expect. Even in this case, however, such damaging actions should only influence market comparisons between proposals and the status quo if speculators expect a correlation between such actions and some policy; speculators will favor the policy that makes damaging actions less likely.

Thus in order to bias a decision via 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 based on this information, or credibly convince deep-pocket speculators that you had so committed yourself.

Welfare Measurement Might Be Corrupted

As is familiar in business, the more that rides on a measurement the more care one must take to look for and otherwise deter corruption in measurement. If you're going to reward your sales staff on how happy customers say they are, for example, it is best not to have the sales staff manage the customer survey. The former Soviet Union was full of examples of bad business measures, such as the proverbial factory rewarded for the tons of nails it produced, which then produced a few enormous nails.

A great deal now rides on GDP measurements. For example, administrations look good if the economy is measured to grow a lot during their watch. Yet most rich countries seem to have limited corruption in GDP measurements to tolerable levels. Also, many businesses seem to have found ways to limit corruption in measuring many important business statistics.

As with actions that actually damage welfare, the straightforward way to profit from actions that change welfare measurements would be to artificially raise or lower some variable that GDP depends on, after one has bet that the GDP estimate would rise or fall. In order to bias a futarchy decision via corruption, one would have to commit to corrupting an important measurement substantially if a decision went one way, but not if a decision went another way, and either have deep financial pockets to bet based on this information, or credibly convince deep-pocket speculators that you had so committed yourself without revealing your corruption to enforcers.

Redundancy is a reasonable general strategy for reducing measurement corruption. Several different agencies could collect the same sort of data and process it the same way and the official estimate could be the median of set of agency estimates.

Policies Might Just 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. If they wanted a certain road built, for example, they could put a term in the welfare definition which takes on a large positive value if the road were built

as specified, and zero otherwise. Speculators would then have to agree that building the road as specified would raise national welfare.

One might be able to rely on courts in part to enforce this, as courts now enforce the rule the 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. The simpler the total welfare function is, the more citizens would be able to review it for their approval, and the more likely it is that representatives would be deterred from directly encoding decisions into the welfare function.

Most Decisions Can Not Detectably Impact Welfare

Speculative markets do not create infinitely precise prices, and not all small price fluctuations are due to new information. Instead, many random contingencies influence exactly who trades what when. To avoid letting such factors excessively influence policy, the proposed futarchy rule requires that markets "clearly" favor a proposal over the status quo, via consistently higher prices or bid and ask offers over a substantial period of time. The flip side of avoiding noise, however, is conservatism. So how could small changes ever be approved?

A lot of small changes added together, however, may no longer be small. Therefore if one proposes a general policy regarding how a large number of small decisions are to be made, that general policy may be large enough to make a detectable impact on national welfare. For example, no single stadium seems big enough to be noticeable, but a general policy regarding all stadiums might be. So one might, for example, propose that all stadium decisions be made by markets estimating measures of regional welfare. This would be an example of futarchy "recursion," where a general futarchy rule authorizes other more localized futarchies.

The Military Needs Secrecy, While Markets Are Public

Militaries often prefer to make their decisions in secret, but futarchy allows public decisions to be made about any aspect of military policy.

Futarchy does not require all decisions to be public, however. The general futarchy rule could be used to approve some policy regime in which military decisions are made in private. That private regime could be a form of futarchy limited to selected approved traders, or it could be any other institution that speculators expect to produce good military policy.

Futarchy would still allow ordinary speculators to override that general military policy in any specific case where they estimate that doing so would raise national welfare. Such speculators would take into account the fact that such overrides would be public, and therefore known to potential enemies. If they still approve an override, it would be because they estimate the secrecy loss to be

outweighed by some greater gain.

You Must Define When A Market "Clearly" Estimates

There are many detail questions about how exactly to implement the betting side of futarchy. Who can issue assets? What maximum and minimum values should be assigned to national welfare? What marketplaces count for setting prices? What information is public about offers there? Should ordinary people, foreigners, or compulsive gamblers, be allowed to trade in these markets? Who can make proposals? Should a proposal fee be required? How large a price difference lasting for how long is a "clear" price difference? Who should be allowed to see proposals and prices, especially regarding national security issues? Should we subsidize market makers?

The recursion approach mentioned above can be used to decide most of these detail questions. If the most basic futarchy rule is conservative, only approving changes that get clear support from trusted speculators, then people can propose within that framework particular less conservative policies. So, for example, if the basic futarchy policy were that the market estimate of welfare given a proposal must beat the status quo estimate for an entire year, one could propose instead that a week is enough. If that week-only proposal itself lasted a year, then new proposals after that would only need to last a week.

If there ends up being some clearly best way to answer these questions, then many different fundamental conservative rules should end up with the same in-practice rules. Someone would propose that best way, and speculators would then endorse it as being most likely to get good policies enacted.

You Might Not Catch Buggy Decisions Quickly Enough

If a bad decision were due to bad information, then the moment new information became available to speculators, the market decision advice should be reversed, giving legal authority to reverse a bad decision. What if a bad decision were instead due to a bug in the welfare function, and hence was an unintended consequence of oversimplifying some aspect of our values? To fix this problem, the democratic part of futarchy would have to vote to change the welfare function. Could they do this quickly enough, before too many bad things happened as a result of the bad decision?

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. One way to deal with this welfare bug problem would be to veto the proposal if another market clearly estimated bad consequences for welfare as it will be defined in the future, say a month from today. Under this system, the moment someone spotted a bug, they could bet that elected representatives would agree with it being a bug and fix it within a month. If speculators agreed with this assessment, then the proposal would not be

implemented.

A danger of this approach is that it might give too much veto power over decisions to representatives who have private information about what is likely to happen soon in the democratic process that changes the welfare definition. This may be a reasonable price to pay, however, to catch bad decisions before they become expensive to undo. And some democratic systems, such as perhaps direct democracy, may allow less such private process information.

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. Futarchy sets up a rule for changing the status quo, but allows for any status quo.

Thus a nation might keep its current political system, modifying its constitution only to allow policy to change via futarchy. A bill would be passed declaring the first definition of welfare, and declaring some agency responsible for measuring welfare as so defined. Then policy would change only as fast as speculators deemed appropriate for achieving welfare as so defined. If speculators deemed a proposal a good idea in general, but also thought that the nation couldn't handle one more change at the moment, they would veto the proposal.

Hidden Losers Will Fight Transition To Futarchy

Any policy change induces winners and losers, and we expect losers to oppose a change if they know about it and are organized enough to act. If we can identify such losers, and if the policy is efficient in the sense that the winners win more than the losers lose, then we should be able to combine the policy with a cash transfer from the winners to the losers to make everyone a net winner. So it is the hidden losers who we expect to consistently oppose a policy, often by arguing that the policy is not in fact efficient (Stiglitz, 1998). If we take opposition as a signal of losing, and compensate opposers more, then more groups will oppose in order to be compensated.

This argument applies to any policy change, and yet policies do change, often for the better. Perhaps it applies with special force to large changes, where broader uncertainty seems to make it easier to argue for possible inefficiency. Hence such problems may prevent adoption of futarchy anytime soon. But large changes do sometimes happen, nonetheless.

The engineering approach accepts a large chance that a promising concept may not overcome all obstacles to deployment. As long as a concept like futarchy seems among the most promising available, it is worth pursuing, even if success could take a long time.

It Seems Hard To Make One Function Encode All Our Values

We could build a welfare function by trial and error. We would start with say GDP, and then entertain proposals to change it. Those arguing in favor of a change would point to choices, real or hypothetical, where they thought that the current welfare function choose badly, but where the proposed new welfare function would choose better. Those opposed to the change would then try to point to choices where they think the old function chooses better than the new one.

Representatives would then have to decide which set of errors seemed the most tolerable, and perhaps how much to trade off decision errors for simplicity. Also, welfare function designers could use these mentioned decision errors to drive the design of new welfare functions that avoided most of both sets of decision errors.

A rather different approach would be for elected representatives to pick some distribution over hypothetical test choices between national outcomes, instead of directly choosing a welfare function. Test choices would be continually selected randomly from the chosen distribution, and then juries, legislative committees, or perhaps the entire legislature would be asked to make a choice on the test case. 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. Speculators would be paid based on how well the welfare function they bet on did in predicting the test choice.[10]

So Many Things Would Have To Go Into A Welfare Definition

Many things could go in to a welfare definition, but even a simple version might do better for citizens than our current forms of government.

Even a very simple GDP-futarchy, one where welfare was defined as GDP fifty years into the future, would need to choose some base asset in which payoffs were paid.[11] To bet on future GDP, one would pick some largest and smallest GDP values to be considered, and create assets which paid off an amount of the base asset given by where in that range future GDP turned out to be.

For each asset that was created, an inverse asset would also be created, so that the sum of the value of the asset and its inverse always equaled the same fixed amount of the base asset. If the future GDP value turned out to be out of the defined GDP range, the asset would pay as if it were at that end of the range. Thus one needs to choose a base asset which is to be bet, and a maximum and minimum plausible welfare value.

A reasonable welfare function would also presumably include time discounting and risk-aversion.

[10] This general approach was suggested by Peter McCluskey.

[11] To induce long-term investment, such base assets should have competitive market returns. Long-term government bonds or broad stock index funds are reasonable options to consider.

For example, welfare might be defined as a few percent annually discounted average of future log GDP. Inequity-aversion also seems desirable. For example, if one could consider national GDP to be a sum of GDPs assigned to various subgroups, and then replace log GDP in welfare with an average over the log GDP for each subgroup. (Varying the group weights might induce various types of discrimination or affirmative action policies.)

Other reasonable extensions to GDP that economists are currently working on include directly considering lifespans, leisure, and environment assets (Boskin, 2000; Nordhaus, 2000). We might also modify or replace GDP with more direct measures of happiness (Oswald, 1997).

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, the most difficult choices in defining welfare may well be how to treat changes in national population due to births, immigration, and emigration.

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 if all known intelligent life were destroyed. More generally destruction is any event that would prevent the paying off of 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 want 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 so far but want to account for. Welfare could then be a weighted average of welfare as we see it and welfare given various types of destruction, weighted by probabilities assigned to each type of destruction.

But Infinity Never Arrives

The proposed initial welfare definition was a few percent annually discounted average of future log GDP. This is an infinite sum, however, and so defining welfare like this, i.e., as an average over "annual welfare" measures like GDP, would basically commit the welfare-measuring-agencies to, ever year, measure annual welfare as defined in all prior official definitions. Eventually this could be more bother than it was worth.

An alternative would be to define total welfare as an average over annual welfares for, say, the next fifty years, plus some weight times welfare as it will be officially defined fifty years hence. Or, continuing farther in this direction, 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. This last plan would only require the welfare agency to measure one kind of annual welfare at a time.

You Have Not Guaranteed Basic Rights and Freedoms

A direct way to preserve rights and freedoms would be to put them into the welfare function. That is, one might collect statistics regarding the number and types of violations of basic rights and freedoms, and give a large negative welfare weight to such outcomes.

If the preservation of rights were considered a time-consistency problem, so that we did not trust future voters to place a large enough negative weight on rights violations, then we might consider a "constitutional bill of rights," whereby we require all approved policies to respect stated rights, and we make it difficult for voters to change the constitution. This approach may help in some cases, but it requires continuing political support that may not materialize, and is limited due to the ambiguity of language.

Risk Premia and Excess Volatility 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 price, averaging over different possible states of the world, states in which future aggregate asset prices are high get 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, if some simple clean way can be found to do this.

Irrational noise traders should induce excess volatility along such dimensions via their shared errors in estimating future prices (Long et al., 1990), although just how much depends on the fraction of such traders in real markets, which we do not know. If a proposed policy will not change errors in noise traders' estimates, then markets should not be biased in comparing the consequences of a policy to the status quo; both markets would have the same error. If we can identify ways in which noise trader errors are correlated with which policies are implemented, however, then we might look for ways to define national welfare (and the base asset) that correct for this bias.

There Is A Decision Selection Bias

When you are in full control of a decision, then you want to make the choice that you estimate will give you the best outcome conditional on your making that choice. If someone else will make the choice, however, then you have to worry about what the fact that they made a choice tells you about what they know. Your estimate conditional on you making the choice can therefore be very

different from your estimate conditional on them making the choice.[12]

This effect can cause biases in how betting markets estimate the consequences of choices. If speculators know that they are making the decision right now, then there is not much of a problem. But if the market does not control the decision right now, then speculators have to worry about what they can infer from the fact that a particular choice is made. This is a reason to put markets directly in charge of decisions, as in futarchy, rather than limiting them to an advisory role.

This potential bias is also a reason to not spread the decision process too far out in time. So proposals should specify some short duration during which they will be considered. And one should not require market prices to agree on a proposal for too long a period before it can be officially approved. Otherwise speculators early in the period may worry that speculators later in the period will know substantially more, and bias their estimates thereby.

This Would Magnify Measurement Errors

When a crude measure of welfare does not substantially influence social actions, then errors between that measure and real welfare may be random enough to allow that crude measure to show us what other social variables influences real welfare. When a crude measure of welfare does substantially influence actions, however, things could change.

In futarchy, speculators would 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 more scope there is for actions which move those errors in a consistent direction, the more harm might be done by a crude futarchy.

Businesses and other institutions have long had to deal with related problems, 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 also often seem to be manageable.

Note also that a similar problem occurs with *any* way in which things learned from looking at

[12] For example, consider an x-y graph where the x-axis represents your payoff if you make a certain choice, and where the y-axis represents your payoff if you make the opposite choice. Your uncertainty about the payoff of your choices would then be described by some probability distribution over this x-y space. The diagonal line $x = y$ would divide the graph into two regions, one where $x > y$, making the x-choice better for you, and the other where the y-choice is better. If the center of mass (i.e., mean) of your probability distribution were in the $x > y$ region, then you expect the x-choice to on average be better. If, however, the decision will be made by someone who knows where the point is in the x-y space, and who has your preferences, then your expectation of your payoff conditional on making the x-choice is the x-component of the center of mass of the probability distribution limited to the $x > y$ region. Your expectation of payoff conditional on making the y-choice is the y-component of the center of mass of the probability distribution limited to the $x < y$ region. It is easy to find distributions such that these x-y-components are on the opposite side of the $x = y$ line from the center of mass of the whole probability distribution. Thus estimates conditional on someone else's choice are not a reliable guide to how you would want to choose based on your information.

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. When you take action, you risk making a mistake.

A Specific Proposal

After reviewing the importance of information institutions, and how democracies may not do well as such institutions while speculative markets do very well, we have discussed thirty issues surrounding the idea of using betting markets to fix democracy. We are thus ready to consider a specific proposal for "futarchy," a form of government where we vote on values, but bet on beliefs.

Futarchy would have a democratic component capable of approving bills, either by electing representatives who approve bills, or via some more direct democracy. The most important of these bills would be those that define what national welfare is, how to measure it later after the fact, and who is to measure what parts of it. At any one time there would be a current official definition of national welfare, and while new bills would be allowed to change that definition, futarchy is committed to eventually measuring and publishing the values of all past official definitions of national welfare.

The agencies that are 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 systematically change what values they eventually declare.

National welfare is to be defined primarily in terms of outcomes desired. Some exceptions can be allowed to embody ethical judgments about actions which are not desired no matter what outcomes they lead to. The court of futarchy should, however, declare a national welfare definition change invalid if it seemed too directly an attempt by legislators to implement specific policy.

A reasonable initial definition of national welfare would be a few percent annually-discounted average of log GDP. Welfare could average the log GDP of as many demographic groups as was feasible. This measure could also be corrected for agency-estimated chances avoided so far of destroying the nation so much that bets would not be paid off. There are also many other directions for improving the definition of welfare.

In futarchy, there would at any one time be a set of nested status quo policies, with more fundamental policies authorizing less fundamental policies. This is like the way that today a law authorizing the creation of some agency implicitly authorizes whatever regulatory rules that agency creates. Futarchy 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* an valid market clearly estimates that, conditional on approving a certain valid proposed policy versus not approving it, national welfare as defined today would be higher, *and* if no valid market clearly vetoes this by estimating that, conditional on approving this policy versus not approving it, 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 overriding any conflicting non-constitutional laws.

Each valid proposed policy declares a short time window during which it can be approved, and says, directly or indirectly, things like what people must do to comply with it, how they are to be punished if they do not comply, and what previous policies it invalidates. Such a proposed policy should specify how and by whom it is to be interpreted. Thus the one official court 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. Thus the democratic process is immune to modification by this futarchy rule.

There are some conservative rules regarding what is a valid market, when such a market clearly estimates something, and how long the veto delay is. These rules say who can propose policies, who can bet on them, what assets and market places count for determining prices, and how long and distinctly prices must differ in order to say that one price is "clearly" above another. Futarchy's basic rule is allowed to approve other policies which look like futarchy's basic rule one level removed, except that they have less conservative answers to these questions.

While not yet specific enough to officially propose, this proposal seems specific enough to critically evaluate.

Conclusion

What should ordinary people believe about the effectiveness of various economic institutions, and of the political institutions that regulate them? Since ordinary people do not have the resources to attend to detailed arguments on such topics, this question becomes: what at-most-handful of general information sources should ordinary people turn to answer this and almost all other abstract questions relatively far removed from their personal experience?

Today a common source, directly or indirectly, is academic-style institutions, i.e., institutions where people show themselves to be experts by long years of focused study and especially by the approval of closely-related similarly-chosen experts. Academia now largely endorses both current political institutions and the levels and types of economic interventions that such institutions pro-

duce. This puts into an awkward position those who endorse academia as our best abstract information institution, while simultaneously claiming that academia has made a serious error regarding current political institutions, or that academic can not be trusted to advise political institutions on the value of specific economic interventions.

This paper has explored a new position, that betting markets in many ways beat academic-style and other familiar institutions as information institutions. This position has been defended, though far from conclusively demonstrated, by reference to the academic literatures on both democratic failures and speculative market successes as information institutions. This position suggests that ordinary people should turn to betting markets to learn about both general political institutions and specific economic interventions. It also suggests that we consider having betting markets displace academic-style institutions in giving policy advice, and perhaps even give such markets a more formal role in this task.

This paper has primarily focused on this later possibility, that of creating a new form of government around the idea of formally deferring to betting markets on matters of fact, while retaining democracy on matters of value. In "futarchy," we would vote on values but bet on beliefs. This paper has taken an engineering-style approach, by considering many design issues with such a concept and then sketching a relatively specific proposal with those issues in mind.

If, after all that has been considered here, the concept passes the low engineering threshold of "promising," then the next appropriate engineering-style step would seem to be to test simple prototypes in simple test environments. A laboratory experiment, for example, might compare a simple version of futarchy to a simple version of more familiar democracy in some simple but non-trivial information environment. Laboratory successes might then prompt larger trials, such as perhaps using futarchy to make important decisions in a corporation or other large organization.

Eventually, we might consider trying agency-level futarchy, such as with monetary policy, and later still, we might even be ready to consider giving futarchy full control of a government. Perhaps, at such a moment of great decision, we might even consider what a betting market has to say about how implementing futarchy would impact future GDP.

References

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., & Long, J. B. D. (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.

Bergin, J., & Sen, A. (1998). Extensive Form Implementation in Incomplete Information Environments. *Journal of Economic Theory*, *80*(2), 222–56.

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.

Brams, S. J., & Fishburn, P. C. (1983). *Approval Voting*. Birkhauser, Stuttgart.

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.

Burnheim, J. (1985). *Is Democracy Possible? The Alternative to Electoral Politics*. Polity Press, London.

Callenbach, E., & Phillips, M. (1985). *A Citizen Legislature*. Banyan Tree Books, Berkeley, California.

Caplan, B. Rational Irrationality and the Microfoundations of Political Failure. *Public Choice*. forthcoming.

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.

Carson, L., & Martin, B. (1999). *Random Selection in Politics*. Praeger.

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. (1998). Prediction markets and information aggregation mechanism: Experiments and application. Tech. rep., California Institute of Technology.

Chen, Y., & Tang, F.-F. (1998). Learning and Incentive-Compatible Mechanisms for Public Goods Provision: An Experimental Study. *Journal of Political Economy*, *106*(3), 633–62.

Chubin, D. E., & Hackett, E. J. (1990). *Peerless Science: Peer Review and U.S. Science Policy*. SUNY Press, Albany.

CNN (1997). Poll: U.S. hiding knowledge of aliens.. http://www.cnn.com/US/9706/15/ufo.poll/.

D'Aspremont, C., & Gerard-Varet, L.-A. (1979). Incentives and Incomplete Information. *Journal of Public Economics*, *11*(1), 25–45.

Davis, J. A., Smith, T. W., & Marsden, P. V. (1996). General Social Survey.. http://www.icpsr.umich.edu/GSS/.

Delli-Carpini, M. X., & Keeter, S. (1989). *What Americans Know about Politics and Why It Matters*. Yale University Press, New Haven.

Dewald, W. G., Thursby, J. G., & Anderson, R. G. (1986). Replication in Empirical Economics: The *Journal of Money, Credit and Banking* Project. *The American Economic Review*, *76*(4), 587–603.

Dixit, A. K. (1997). *The Making of Economic Policy: A Transaction-Cost Perspective*. MIT Press, Cambridge, Massachusetts.

Duggan, J. (1997). Virtual Bayesian Implementation. *Econometrica*, *65*(5), 1175–99.

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.

Friedman, D. (1989). *The Machinery of Freedom* (2nd edition). Open Court, La Salle, Ill.

Fuller, S. (1988). *Social Epistemology*. Indiana University Press, Bloomington, Indiana.

Fuller, S. (2000). *The Governance of Science*. Open University Press, Buckingham.

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..

Group, B. R. (1994). Barna Research Online Archives.. http://www.barna.org/cgi-bin/PageCategory.asp?CategoryID=6.

Groves, T., & Ledyard, J. O. (1977). Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem. *Econometrica*, *45*(4), 783–809.

Hanson, R. (1992). Idea Futures. *Extropy*, *3*(2), 7–17.

Hanson, R. (1995a). Comparing Peer Review to Information Prizes – A Possible Economics Experiment. *Social Epistemology*, *9*(1), 49–55.

Hanson, R. (1995b). 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. 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). *Efficiencey of Racetrack Betting Markets*. Academic Press, San Diego.

Hayek, F. A. (1945). The Use of Knowledge in Society. *The American Economic Review*, *35*(4), 519–530.

Hayek, F. A. (1973). *Law, Legislation and Liberty*. University of Chicago Press, Chicago.

Henrion, M., & Fischhoff, B. (1986). Assessing uncertainty in physical constants. *American Journal of Physics*, *54*(9), 791–798.

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.

Jorion, P., & Goetzmann, W. N. (2000). A Century of Global Stock Returns. Tech. rep. 7565, NBER.

Kadane, J., & Winkler, R. (1988). Separating Probability Elicitation from Utilities. *Journal of the American Statistical Association*, *83*(402), 357–363.

Kalai, E., & Ledyard, J. O. (1998). Repeated Implementation. *Journal of Economic Theory*, *83*(2), 308–17.

Kindleberger, C. P. (1984). *A Financial History of Western Europe*. George, Allen, & Unwin, London.

Kittlitz, K. (1999). Experiences with the Foresight Exchange. *Extropy Online*. http://www.extropy.org/eo/articles/fx.htm.

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.

Long, J. B. D., & Lang, K. (1992). Are all Economic Hypotheses False?. *The Journal of Political Economy*, *100*(6), 1257–1272.

Long, J. B. D., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise Trader Risk in Financial Markets. *Journal of Political Economy*, *98*(4), 703–738.

Maddison, A. (1995). *Monitoring the World Economy 1820-1992*. OECD, Paris.

Mahoney, M. J. (1977). Publication Prejudices: An Experimental Study of Confirmation Bias in the Peer Review System. *Cognitive Therapy and Research*, *1*.

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 Misalligned 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.

Ortner, G. (1998). Forecasting Markets - An Industrial Application. Tech. rep., University of Technoloy Vienna. http://ebweb.tuwien.ac.at/apsm/fmaia2.pdf.

Oswald, A. (1997). Happiness and Economic Performance. *Economic Journal*, *107*, 1815–1831.

Pagano, M. (1989). Endogenous Market Thinness and Stock Price Volatility. *The Review of Economic Studies*, *56*(2), 269–287.

Page, B. I., & Shapiro, R. Y. (1992). *The Rational Public: Fifty Years of Trends in America's Policy Preferences*. University of Chicago Press, Chicago.

Peters, D., & Ceci, S. (1982). Peer Review Practices of Psychological Journals: The Fate of Published Articles, Submitted Again. *The Behavioral and Brain Sciences*, *165*, 187–195.

Plott, C. R. (2000). Markets as Information Gathering Tools. *Southern Economic Journal*, *67*(1), 2–15.

Postlewaite, A., & Schmeidler, D. (1986). Implementation in Differential Information Economies. *Journal of Economic Theory*, *39*, 14–33.

Redner, H. (1987). Pathologies of Science. *Social Epistemology*, *1*(3), 215–247.

Roll, R. (1984). Orange Juice and Weather. *The American Economic Review*, *74*(5), 861–880.

Rosen, S. (1997). Austrian and Neoclassical Economics: Any Gains From Trade?. *Journal of Economic Perspectives*, *11*(4), 139–152.

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.

Sharpe, G. (1994). *William Hill's Bizarre Bets*. Virgin Books, London.

Shiller, R. J. (1993). *Macro Markets: cReating Instituitions for Managing Society's Largest Economic Risks*. Clarendon Press, Oxford.

Shiller, R. J. (2000). *Irrational Exuberance*. Princeton University Press.

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.

Tideman, N., & Tullock, G. (1976). A New and Superior Process for Making Social Chocies. *Journal of Political Economy*, *84*(6), 1145–1159.

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.

Yeager, L. B. (1997). Austrian Economics, Neoclassicism, and the Market Test. *Journal of Economic Perspectives*, *11*(4), 153–165.

Zeckhauser, R. J., & Viscusi, W. K. (1990). Risk Within Reason. *Science*, *248*, 559–564.