Anthony J. Evans and Jeffrey Friedman

"SEARCH" VS. "BROWSE": A THEORY OF ERROR GROUNDED IN RADICAL (NOT RATIONAL) IGNORANCE

ABSTRACT: *Economists tend to view ignorance as "rational," neglecting the possibility that ignorance is unintentional. This oversight is reflected in economists' model of "information search," which can be fruitfully contrasted with "information browsing." Information searches are designed to discover unknown knowns, whose value is calculable ex ante, such that this value justifies the cost of the search. In this model of human information acquisition, there is no primal or "radical" ignorance that might prevent people from knowing which information to look for, lacking omniscience. Unlike ignorance that is rationally chosen on the basis of an accurate cost/benefit calculation, radical ignorance can explain human error. An account of error as grounded in radical ignorance bypasses the need to appeal to irrationality in order to explain economic (and other) mistakes.*

Anthony J. Evans, Associate Professor of Economics, ESCP Europe, 527 Finchley Road, Hampstead, London, NW3 7BG, United Kingdom, anthonyjevans@gmail.com, delivered previous versions of this paper at a conference entitled "From Vienna to Virginia" (New York, September 2008) and at the Southern Economic Association Annual Meeting (Washington, D.C., November 2008). Jeffrey Friedman, Max Weber Fellow of the Institute for the Advancement of the Social Sciences, Boston University, and Visiting Scholar, Department of Government, University of Texas at Austin, is the editor of *The Rational Choice Controversy* (Yale, 1996) and *What Caused the Financial Crisis* (Penn, 2010); and is the coauthor of *Engineering the Financial Crisis: Systemic Risk and the Failure of Regulation* (Penn, 2011). The authors thank Peter Boettke, Jim Cowen, Steven Horwitz, Daniel Kuehn, Mateusz Machaj, John Meadowcroft, Mark Pennington, Jochen Runde, and Nikolai Wenzal for comments on previous drafts

Critical Review 23(1–2): 73–104
ISSN 0891-3811 print, 1933-8007 online

DOI: 10.1080/08913811.2011.574471

> The worth of new knowledge cannot begin to be assessed until we have it. By then it is too late to decide how much to spend on breaching the walls to encourage its arrival.
>
> —G.L.S. Shackle 1972, 272-73

In 1976, the Chicago-school economist George Stigler (1976, 216) advised that economists bracket the very existence of mistakes for purposes of economic modeling, maintaining that error "will not become a useful concept until we have a theory" of it. Stigler offered this advice fifteen years *after* he had produced what is still the seminal statement of the mainstream neoclassical economic approach to the question of *ignorance*: his 1961 paper, "The Economics of Information." There, Stigler had maintained that economic ignorance must be rational in the very strong sense that the ignorant party has *accurately* calculated that knowing a given datum would not be worth the cost of "searching" for it. This strong concept of rational ignorance is incompatible with a theory of error, as he implicitly recognized in 1976, because the rational searcher *knows* that the information of which he or she chooses to remain ignorant would not justify the search cost. Such a searcher cannot make a mistake.

The spread of the (strong) idea of rational ignorance into political science (e.g., Somin 1998 and 2012), where it is mistakenly used to explain why voters are so abysmally ignorant about politics,[1] is enough to justify our attempt to fulfill Stigler's request for a theory of error. (Indeed, our theory will be a theory of ignorance-based error.) However, it is just as important in economics, where widespread and significant errors based on ignorance would seem to be commonplace: otherwise, corporations and investors would never lose money.

A Theory of Ignorance-Based Error

Our theory of error, in a nutshell, is that people often make serious mistakes because of their ignorance.

One might think it obvious that ignorance is a source of human error, but if so, this fact has been obscured from economists by their extremely narrow idea of rational ignorance, where "ignorance" is *deliberately chosen* on the basis of an accurate cost/benefit calculation. "If the cost of search is equated to its expected marginal return, the optimum amount of search will be found," Stigler (1961, 216) tells us.

What Stigler called "the cost of search" constitutes the *resource constraint* on information acquisition. Without dismissing it entirely, we would like to suggest, as a heuristic for our larger argument, the increasing irrelevance of the resource constraint in the age of the Internet. Information that *one knows that one needs* is increasingly available free of charge.

It might seem paradoxical, then, that we appear recently to have experienced one of the greatest bouts of mass economic error in history: the financial crisis.[2] Our theory of error tries to resolve the paradox by starting at a point that is diametrically opposed to that of search theory—namely, the point at which one is inadvertently, unwittingly ignorant of information that one does *not* know that one needs; or that one knows one needs but does not know exists; or that one knows one needs but cannot find. Inadvertent ignorance, because it is not *chosen*, is neither rational nor irrational: rationality and irrationality are properties of a decision, but inadvertent ignorance is "there" whether or not one is making a decision, and it persists despite making a rational decision. Inadvertent ignorance is simply a basic aspect of the human condition, grounded in the fact that human beings are not omniscient.

We thus propose an *epistemological* theory of error that applies to all fields, including economics. According to this theory, people make genuine errors when they (a) are ignorant of relevant information (for any of the reasons just specified); (b) are misled by false information; (c) are misled by true but irrelevant information; or (d) misinterpret true and relevant information.

Noneconomists might view such a theory as truistic, but economists have trouble with it because they tend to view decision makers not only as rational agents in the weak sense of intending their actions to serve some further purpose (which we do not dispute here),[3] but as agents who are rational in the strong sense that they know what they need to know if they are to achieve their purposes. In search theory, this strong rationality manifests itself in the assumed ability of an ignorant agent to *know in advance* the cost and the benefit of learning that of which he or she is "ignorant." We ask how one could know the cost and the benefit of learning something without, in fact, already having learned it.

Stiglerian search theory presupposes that "ignorant" agents are in fact omniscient about the costs and benefits of acquiring the information they do not yet know. This assumption is derived from contemporary economists' foundational theory of general equilibrium, which posits "that every agent has all the information that is needed about every good"

(Hahn 1981, 134). The exceptions are (1) when agents know that the cost of learning this information is too high to justify the benefit of acquiring it (Stigler 1961), i.e., rational ignorance; (2) when information is unverifiable or costly to verify; and (3) when certain information is being hidden from the agent by someone with an interest in perpetuating ignorance of this data, i.e., an interest in perpetuating a situation of "asymmetrical information" (Akerlof 1970; Greenwald and Stiglitz 1986). These cases get as close as mainstream economists are willing to go in acknowledging ignorance.

The superficiality of this approach to ignorance is suggested by the fact that asymmetrical information produces error on the side of the party who is kept in the dark—but *not* on the side of the party who has the information and is hiding it. Similarly, *rational* learning cannot lead to or be based upon error. So genuine, "unforced" errors must be denied or else explained away as irrational.

In the wake of the financial crisis, we can expect an upsurge of interest among economists in the notion that decision makers err even though (by common assumption) they *do* have the relevant information and have *not* misinterpreted it. In other words, since many economists have conflated rationality with non-error, they will probably conflate error with irrationality.

Irrationality theories might at least seem to have the advantage of providing what Stigler called for, since irrationality, unlike rational ignorance and unverifiable or asymmetrical information, might account for inadvertent, accidental errors. Perhaps, as is so often said about the financial crisis, the bankers "did not want to know" the truth (Stiglitz 2010, 14), even though they allegedly possessed the "the information" and did not misinterpret it. Because it is built on the same assumption of omniscience as the orthodox economics of rational information search, we will argue, it is implausible at best, incoherent at worst.

First, however, let us examine how the omniscience assumption trivializes the information "imperfections" contemplated in search theory.

The Omniscience Assumption vs. Surprise

The problem with the economists' attempts to explain error is that they all presuppose that if relevant information exists, the default position of an economic decision maker is to possess this information and to interpret its significance accurately. This includes information about the

value of a given ''piece'' of information, such that the putatively ''ignorant'' agent, before undertaking a search for datum X, knows that the benefits of knowing X will outweigh the costs of learning X. Therefore, if the information is *not* possessed by the decision maker, somebody's deliberate decision must have interfered with the default position of what would otherwise be ''All-Knowing Beings'' (Bhidé 2010, ch. 5). The theory of rational information search thus implicitly builds into the definition of rationality not merely the desire to achieve an objective, but the knowledge of *how* to achieve it. It is no wonder that these models have little room for genuine blunders.

If economic agents were allowed to be rational but *ignorant* of how to achieve their objectives, such that they could blunder, it would be difficult for the economist to predict the results. If a rational agent may not, in any given instance, know what he or she needs to *learn* if he or she is to achieve an objective, then how could we forecast what he or she might *do* in pursuit of that objective? Such agents might as well be acting randomly, and the mainstream economic pretense of being a predictive science would come to an end. Thus, economists have never been able to do justice, for example, to uncertainty of the type that Keynes (1921) and Knight (1921) famously contrasted against calculable risk.[4] If agents faced genuine, incalculable uncertainty, their empirical behavior would be difficult or impossible to model.

Tractable modeling, however, is in this case purchased with glaring unrealism. Any sentient human being is aware of his or her own unchosen ignorance, and only the paranoid assume that what we do not know and did not choose not to learn must have been hidden from us by someone with an interest in maintaining our ignorance.

The signal to sentient human beings that they have been inadvertently ignorant is *surprise* (Kirzner 1997, 81). If we are surprised by something, it is almost certainly the case that we did not know it; or that we thought we knew something that has suddenly turned out to be false; or that we knew something that, it has suddenly turned out, we misinterpreted. Only surprises that are deliberately arranged for us (such as surprise parties) resemble the type of ''ignorance'' that mainstream neoclassical economics is prepared to accept (in this case, asymmetrical information). But our lives are peppered with unarranged surprises about matters that we didn't know existed or that we thought we had a handle on but turned out to be wrong about. The absence of conceptual space for surprise, we maintain, is the missing dimension in

mainstream economics, without which it cannot produce the needed theory of error.

The missing dimension has been described as "radical ignorance" (Ikeda 2003) and "inadvertent ignorance" (Bennett and Friedman 2008). Both terms refer, in an economic context (if we stipulate instrumentally rational but possibly inaccurate choice),[5] to ignorance of the very fact that one would need to learn some datum or body of information in order to achieve one's ends. One cannot learn something—except inadvertently, as when "browsing"—that one does not know in advance would be rewarding to learn. And if we fail to learn it, we will be unpleasantly surprised to find out about it—for example, by experiencing losses rather than profits.

The Immunity of Radical Ignorance to Incentives to Learn

We believe that there are many important cases in which one cannot know ex ante that it would be valuable to learn something—without first learning it. For instance, in 2005 one could not know the value of learning the proposition that *the housing boom is not being sustained by the growing wealth of Americans and their desire to live in bigger houses* if one did not already have reason to think that this proposition was indeed true. People who did not already suspect that a bubble (as opposed to a boom) was developing were radically ignorant about it; they could not very well have deliberately set themselves the task of finding out about it.[6] They could have found out about it only *accidentally*: by hearing or reading a presentation of the bubble hypothesis in glancing through the business pages, for instance. If they found the presentation persuasive, they would be pleasantly surprised by it, assuming that they could now do something about it.

It is possible that a factor in economists' resistance to recognizing radical ignorance is that it cannot be affected by one of the main weapons in economists' conceptual arsenal, incentives. Only if one already knows (or, rather, thinks) that something would be valuable to learn would one recognize the *incentive* to learn it; but in order to know that, one would somehow have to know the crucial aspect—the value—of the "unknown" information ex ante, before one has learned it. This borders on the impossible, since even if one knew that somebody out there was claiming that the whole economy was about to collapse, one could not

know how valuable it would be to read the arguments and evidence in support of this *outré* claim until one had actually read the arguments and evidence. Before doing that, when one was still ignorant of the arguments and evidence (as opposed to the claim), one could be aware of the *potential* value of the information embodied in the claim, but, lacking the information (the evidence and arguments) itself, might mistakenly discount the actual value of learning it and therefore decide not to learn it. Lots of people make lots of potentially important claims, and nobody can pursue all of them. The Internet merely makes plain what was previously just as true: the problem facing those who wish to take well-informed actions is usually not the scarcity of information, but the overabundance of it.[7]

Where information is overabundant, the value of acquiring known-about but not-yet-learned information is an unknown unknown. Therefore, the objective incentive to learn the information is impotent to affect the deliberator's action, because the incentive consists of the actual value of the information, which the deliberator cannot know without first learning the information. Similarly, in the more clear-cut case of information whose very *existence* is an unknown, the ex-post incentive to know it goes unrecognized ex ante, and the would-be information searcher is immune to the ''incentive'' to search for it. In either case, he or she may therefore encounter an *unpleasant* surprise when (say) the bubble bursts—as will those who have acquired relevant information that they misinterpret, and those who have acquired inaccurate or irrelevant information (such as information about the growing prosperity and population of the United States) that they misinterpret as being relevant.

Investors who were putting their money into mortgage-backed bonds can be said, with the luxury of hindsight, to have had an *objective* ''incentive'' to learn that the mortgages were part of a housing bubble. But because they were ignorant of the existence of the bubble claim; or ignorant of the arguments and evidence supporting the claim; or because they were were predisposed to view the claim as specious or implausible due to their previous acquisition of inaccurate or misinterpreted information (i.e., they were ignorant of the relevant true information), or due to their use of previously successful but, in this case, irrelevant information-search heuristics[8]—these investors did not (subjectively) know beforehand that they should learn the information, or it did not occur to them to learn it, so the incentive to learn it was a dead letter.

The Rationality of Inaccurate Information Searches

Are cases such as the financial crisis rare?

It would be ridiculous to try to quantify how often people's ignorance is too radical to allow them to know beforehand the costs and benefits of searching for a specific datum or body of knowledge. That quantity is a known unknown, because the universe of cases is still under development as modern life and, in this case, capitalism continue to change in some ways and remain the same in others. However, we can try to imagine cases in which the information-search model *would be* applicable so as to suggest how rare these cases are in comparison to cases of radical ignorance.

The easiest case in which one could know in advance the costs and benefits of acquiring certain information might be if one were directly paid to learn it. Then the value would be the predetermined payment, known ex ante, and the cost would the effort of learning it, which is, or at least might be, known ex ante.

Such situations are relatively infrequent. Usually neither the cost nor the benefit of learning something is clear in advance, because the instrumental benefit of learning is not often external to the information learned—i.e., it is not often a fixed payment—but is rather the *change* that the information itself will make to the actions that one would have taken if one had remained in ignorance. (Stepping back from economics, the value of learning is the change in one's thinking that it brings about, regardless of whether one can act on the basis of the new thoughts.) In effect, the information itself, if valuable, will be surprising, and will thus substitute a pleasant intellectual shock for, say, the unpleasant shock of losing one's investment.

Moreover, the surprising aspects of learning are not confined to the benefits; there are also surprises in the form of costs. Consider grades for classroom performance. A student may place a certain value on getting an A, and may know that according to the course syllabus, getting an A requires mastery of a certain body of information. The student cannot, however, know how difficult it will be to master the information before it has been mastered. Mastery at a level that will procure an A is not the same thing as merely reading a certain number of pages of material at a certain number of minutes per page. The difficulty of mastering the material cannot be fully known until one has gone through the process of trying to master it. One may guess how difficult it will be, based on past experience or any other heuristic. It is not *irrational* to make a calculation

based on a guess, but since the guess may be wrong, this also means that it is not irrational to make a mistake in under- or overestimating the cost of learning. Rationality has nothing to do with accuracy.

These examples, moreover, presuppose that somebody—the teacher creating the incentive to learn and assigning the things to be learned—*does* know ''the answer.'' But in a capitalist economy (like any economy), nobody reliably knows ''the answer,'' even though many people may think that they do, and even though many people really may know particular answers at a given time. Consumers, for example, may think they know what they want to buy, but they may encounter an unpleasant surprise after making the purchase. Entrepreneurs may think they know what consumers want to buy, yet they would (as a group) never lose money if this ''knowledge'' were reliable—a point about which we shall have much more to say.

Thus, in a more realistic example, one may ''know'' that if one learns how to trade options or repair a car, one will be able to use these skills. Thus, one may be able to guess the value of learning them. But even stipulating that one can *accurately* guess the cost of learning them, one cannot know ex ante how beneficial such skills will be once they are acquired. The going rate for auto mechanics or options traders may justify a certain level of effort in acquiring the necessary skills, but one cannot know in advance if one will be able to command the going rate, so one may err in devoting the effort to learning the skills. Once again, it is not irrational to make one's best guess at the costs and benefits of learning a discrete body of knowledge, and then to proceed or not depending on the outcome of the calculation (or guesstimate). It is quite rational to do so; but one may err in one's rational calculations because one is ex ante ignorant of the true costs or benefits—which, in the case of learning, can be known only ex post.

The Triviality of Search Theory

Stigler (1961, 213) defines ''search'' too narrowly to encompass such examples of deciding whether to learn unknown knowns: ''A buyer (or seller) who wishes to ascertain the most favorable price must canvass various sellers (or buyers).'' ''Search'' is used here as a synonym for the wider concept of ''canvass,'' and thus could be interpreted in a broad

manner. However when Stigler (1961, 219) provides more detail about the concept, its limitations are striking:

> If a buyer enters a wholly new market, he will have no idea of the dispersion of prices and hence no idea of the rational amount of search he should make. In such cases the dispersion will presumably be estimated by some sort of sequential process, and this approach would open up a set of problems I must leave for others to explore.

Stigler brackets issues of *genuine* ignorance (and thus the possibility of error) so that he can assume that buyers *do* have knowledge of the dispersion of prices; if they do not, and thus might be mistaken, we are in the territory where we would need something more than the theory of search. (With ''some sort of sequential process,'' searching begins to blend into browsing. One encounters information not in its order of importance—due to the fact that one's search is somehow pre-informed by that order—but in hit or miss fashion.)

Similarly, in George A. Akerlof's classic contribution to the literature, ''The Market for 'Lemons,''' the possibility of ignorance is radically shrunk by assumption. ''Suppose,'' Akerlof (1970, 492) writes, ''that the quality of all cars is uniformly distributed.'' Thus, the buyer's problem is not so much to discover whether or not a *particular* car is a lemon, but to compute the probability of this for *any* used car, based on inputs derived from the known distribution of lemons as opposed to non-lemons. If not for the fact that this knowledge is hidden by the seller, then the buyer could no more make a genuine ''error'' than could a gambler who knows the odds against winning but decides to risk them anyway just because he or she feels lucky. Such a gamble *would* be irrational—assuming that the purpose of taking this risk is to win, rather than to experience the thrill of the activity in itself. However, search theory is supposed to be describing a process of learning that is instrumental to pecuniary gain.

Of course we expect simplifying assumptions to be made in any model. For example, Akerlof (1970, 489) assumes that there are just four kinds of car in the market, which is unrealistic but unobjectionable, because it is merely an effort to apply the theory to a small sample for the sake of exposition. But the assumption of a uniform distribution is unrealistic in a manner that relates to the suitability of the model itself. Indeed, the model rests on the assumption that the distribution is not only uniform, but also knowable, in principle, to the prospective car buyer, given that it *is* known to the seller. Thus, the purchase of a lemon is due to ''an asymmetry in *available* information'' (Akerlof 1970, 489),

suggesting that "unavailable" information, which is to say information aside from the known distribution—not only unknown unknowns, but known unknowns—is irrelevant.

An important literature on "signaling" in labor markets displays similar features. The 1973 *locus classicus* of this literature, by Michael Spence, begins by explaining the problem with hiring employees: since employers do not know the productivity of an applicant, this "makes the decision one under uncertainty" (Spence 1973, 356). Employees, knowing this fact, must signal their productivity by using heuristics such as educational degrees.

This would be a case of genuine, albeit not radical, ignorance—the worker's productivity is known in advance to be unknown both by the employer and the employee—except for the fact that Spence assumes the existence of a "wage schedule": an objective statistical distribution that tells prospective employees the value of signaling. Thus, a prospective employee "will invest in education if there is sufficient return as defined by the offered wage schedule" (ibid., 358). A case in which genuine errors on both sides may occur is transformed into an instance of accurate signaling that accomplishes the objectives of both parties, albeit creating rent-seeking type inefficiencies.[9]

Similarly, Joseph Stiglitz and Andrew Weiss (1981) use the term *search* to refer to a bank's efforts to discover profitable business plans. Their model assumes "that the bank has identified a group of projects; for each project κ there is a probability distribution of (gross) returns R." Again, the striking point is that the probability distribution is *given, i.e., known*. If so, "search" becomes little more than probabilistic modeling. But it is obvious why this assumption is made. If the probability distribution is not known, at least to the economist, behavior cannot be modeled; there is simply no way to generate predictive results without determining behavioral parameters.[10]

This is only an indicative overview of a vast literature, but the seminal works are telling. Whether the particular market in question is for cars, education, or banking services, all of the authors claim that their findings are applicable more generally. But as we can see, their findings rest on assumptions of knowledge, whether symmetrical or asymmetrical, which are achieved by radically shrinking the types of ignorance under consideration. Yet these are models that are supposed to explain learning, where one moves from a position of ignorance to a position of knowledge. The models, in fact, never explain learning, because learning is obviated by the "givenness" of the probability distributions.

Capitalism, Socialism, and Knowledge

In his encyclopedia entry on "Information," Stiglitz (2008) cites Hayek's 1945 *American Economic Review* essay, "The Use of Knowledge in Society," as a pioneering text in the economics of information. There Hayek argued that under capitalism, freely fluctuating prices *convey information* that would otherwise be unavailable because it is dispersed and local. All the sales of various consumer and capital goods are aggregated to produce a constantly changing array of prices that *accurately* reflect supply and demand.

Hayek reached this conclusion in response to the "socialist calculation debate" of the 1920s and 1930s, which pitted him and another Austrian economist, Ludwig von Mises, against an array of economists who defended the feasibility of central economic planning. Stiglitz, however, cleverly combines Hayek's 1945 argument with the concept of information asymmetry to criticize the adequacy of market prices. Consumers will not volunteer to pay a higher price than the one on offer, even if they would be willing to pay more. Likewise, producers will not volunteer to reduce prices to the point where they just cover costs if they can avoid it.

"If individuals were willing or able to reveal their information," Stiglitz writes, "everybody could be made better off" by having the information reflected in accurate prices (Rothschild and Stiglitz 1976, 629). A capitalist price system *could* convey accurate information about actual supply and demand, but the strategic behavior of consumers and producers gets in the way. Stiglitz (2008) quotes Hayek against Hayek's own idealization of the price system: "A much better approach is to assume the world we have, one in which everyone has only a little information."[11] Thus, the economics of information "substantiates Hayek's contention that central planning faces problems because it requires an impossible agglomeration of information," but it also shows that these problems are at least as bad under capitalism due to market inefficiencies.[11]

Stiglitz is assuming, however, that consumers and producers *know* the price they would be willing to pay or receive in an imaginary auction, not in the concrete situation of a transaction where a single price is offered. Stiglitz is begging the question against Mises and Hayek[12] by assuming that objective supply and demand information could "exist" in the form of curves that are not merely blackboard heuristics, but are carried around in the heads of consumers and producers—information that is both knowable and known absent actual transactions.[13] Again, there is no real

ignorance here except of the asymmetrical variety. On each side of the supply and demand equation, market participants are omniscient, but they have no incentive to share their wisdom with each other.

In the Stiglitz model, then, the informational problem is merely one of aggregation—of transporting existing and known information from dispersed locations (for instance, by better incentivizing the holders of information to reveal it). This, indeed, is the received view among economists. The responsibility for this trivialization of the problem of knowledge rests on Hayek's shoulders.

It was Hayek, after all, who insisted that prices aggregate producers' (and consumers') *local knowledge* of *actual* supply and demand conditions. (If it were not ''knowledge'' of actual conditions, it would be mere speculation.) Thus, only the central planning board is ignorant, because it does not have access to ''knowledge of people, of local conditions, and of special circumstances'' (Hayek 1945, 80). In that case, however, if we had some way to transmit this knowledge to the central planning board—as with portable telecommunications devices—the problem would be solved.

Henceforth the ''knowledge problem'' of both socialism and capitalism was seen as being the ''transmission'' of known information, not the discovery of unknown information or of information whose value was unknown. As Hayek (1945, 78-79, emph. added) himself put it, ''The various ways in which the *knowledge* on which people base their plans is *communicated* to them is the crucial problem for any theory explaining the economic process, and the problem of what is the best way of utilizing *knowledge* initially dispersed among all the people is at least one of the main problems of economic policy—or of designing an efficient economic system.''[14]

Israel Kirzner (1973), a student of Mises, took a step in the right direction by replacing Hayek's knowledgeable men on the local scene with ex ante ignorant ''entrepreneurs.'' Yet his explanation of how ignorant entrepreneurs learn what they need to know (let alone interpret it correctly), so as to make a profit, is, at best, a non-explanation: they are ''alert'' to profit opportunities, hence they notice them, hence they know what to do about them.

It is true in a trivial sense that entrepreneurs can be defined as those who are ''alert to profit opportunities,'' but we wonder why agents of the central planning board could not be equally alert. The real issue is not alertness, but the magical property that Kirzner attributes to those who

are alert: the property of *thereby* finding what they are looking for (a profit opportunity) and knowing what to do about it. If mere alertness—activated by "the profit motive" (Kirzner 1978, 74)—were all it took to produce the requisite knowledge, one could incentivize central planners with the same motive or an even stronger one, such as the death penalty (Friedman 2006, 486-97).

Kirzner's answer represents a headlong retreat from the purely *epistemological* challenge to socialism that Mises ([1920] 1937) had posed in opening the socialist-calculation debate in 1920. Most previous critiques of the feasibility of socialism had focused on incentives problems: How to motivate workers or managers without the lure of differential pay or profits? Such critiques begged the question against the likes of Marx, who held that the self-interestedness that would make such incentives necessary was an historically contingent outcome of capitalism; it would disappear once the socialist revolution liberated the altruistic human "species essence." Rather than argue over the motivational "nature of man," as political philosophers had done for millennia, Mises made a truly epochal advance, bracketing motives and thus incentives so he could focus solely on the question of how the central planners could know what they would need to know.

Kirzner reverses this advance by conflating epistemological and motivational issues: the trigger for the "alertness" that leads the entrepreneur (but not the central planner) from ignorance to knowledge is, as it turns out, selfishness after all. Thus, Kirzner (1978, 74) concludes his paper on "Economics and Error" by suggesting that "only . . . the lure of pure entrepreneurial profits" can "harness entrepreneurial activity toward the systematic elimination of error." Kirzner fails to explain, however, why some entrepreneurs are "alert" to profit opportunities while others, who err, are not—despite the presumably equal pull on all of them that is exerted by the profit motive. Why does the alertness of some of them move them from a position of ignorance to one of knowledge, while the alertness of others does not? Kirzner does not answer this question, nor could he within the incentives-based framework he is using, since the question—ever since Mises in 1920—had been, and remains, epistemological, not motivational.

The epistemologist asks, "How can one learn the truth?" The economist answers, "If one is sufficiently motivated to learn it, one will find a way." We submit that this is literally religious thinking, based on a groundless faith in incentives to produce knowledge.[15] We say this in full

awareness of the fact that Kirzner (1997) provides the best extant critique of the ''economics of information,'' without which our own would be impossible. Kirzner challenged the relevance to capitalist realities of economic models that simply assumed an equilibrium in which all supply and demand curves intersected at the price per marginal unit that precisely covered the costs of production. Kirzner noted that in such a world there would be no profits. We are simply noting that in that world, as well as in Kirzner's alternative model of capitalism, there would be no losses.

Information vs. Interpretation

Epistemologically, Kirzner did exactly what the market socialists had done when ''responding to Hayek's argument that the relevant data cannot be considered objectively 'given' to the central planner'': they ''proposed a solution that assumes such data to be objectively given to the decentralized plant managers instead'' (Lavoie 1985, 160). This is also what Hayek had done in ''The Use of Knowledge in Society.'' In Kirzner's version of this recurrent mistake, however, the ''alert'' entrepreneurs are, as it were, engaged in a sort of instantaneous information ''search'' that provides them the objectively given knowledge.

We say this because alert entrepreneurs somehow know *where to look* for profit opportunities. One cannot simply be alert to them in general, or in the abstract, since in the abstract, *everything* could conceivably be interpreted as a profit opportunity. There is far too much information to allow a searcher who looks everywhere to succeed. It is literally impossible for any non-omniscient being to be uniformly ''alert'' to the infinite amount of information that *might* indicate profit opportunities, rather than being alert to particular information due to particular stimuli *that this particular entrepreneur, but not another one, interprets as significant*. As in the search literature, however, for Kirzner's entrepreneurs, it is as if the ''data'' are given and require no interpretation, including the meta-data about where to search for the relevant and accurate information about profit opportunities.

Kirzner's 1997 critique of mainstream search theory justly criticizes the assumption that the needed knowledge—and thus its location—are ''given,'' such that all entrepreneurs know what to do and therefore none can make a profit in competition with the others. Kirzner's own schema also assumes that the needed knowledge is given once entrepreneurial

"alertness" kicks in. Hayek failed to acknowledge that the perceptions and interpretations of the man on the spot can be wrong, and Kirzner fails to acknowledge that entrepreneurial "discoveries" prompted by alertness can be errors that produce losses.

Knowledge vs. Opinion

There is no genuine uncertainty among the Kirznerian entrepreneurs, or the Hayekian "men on the spot," or the market-socialist plant managers. We mean uncertainty as Keynes and Shackle meant it: not a psychological (or motivational) wavering, but radical ignorance. When one is radically ignorant, one may well misperceive one's situation as requiring a mere information search—*as if the information has already been discovered by someone else*. At the conclusion of the search, one may then be as confident that one has found "the answer" as is a student who has looked something up on Wikipedia. But, like that student, *one may be wrong*. If one is wrong and does not know it—which really is the only way one can be wrong, unless one is irrational—then one is radically ignorant, no matter how confident one feels that one is actually knowledgeable. Hence the commission of error by people who turn out to be "overconfident," ex post.

The problem with treating economic information as given, whether by one's location, one's alertness, or by an array of prices, is that all of these factors leave out the rightful starting point of modern epistemology, as noticed by Descartes: that one's opinions about the truth may be in error. We submit that this is the problem with capitalism—but also with socialism: Whoever is in charge, whether competing entrepreneurs or central planners, may be in error. Error is possible even in the face of the "hard data" of a given array of prices, or in the face of what the man on the spot observes with his own eyes, or what the entrepreneur knows in his gut, because data never speak for themselves. Data must always be interpreted as relevant and as instructive about what one should conclude (and therefore do); and interpretations, being the products of fallible human knowledge, reasoning, and imagination, can be mistaken.

It is not true that if one is located in a particular place in the world, one "know[s] directly of the relevant changes and of the resources immediately available to meet them" (Hayek [1945] 1948, 84). The only thing that one "knows directly" is the *cogito*. The rest of the data to

which one has access must be interpreted. Two different people in exactly the same ''spot'' could have quite different interpretations of which circumstances *are* relevant and which resources *can* be more profitably employed. Judgments about such matters entail theorizing about ''the data,'' not merely observing, accumulating, and thus, in principle, communicating them to other market partcipants.

Hayek (ibid.) is, however, backed into the corner of relying on constant *changes* in circumstances to explain the unfeasibility of communicating local data to a central location, given his non-interpretive conception of what might be communicated. An interpretive conception of what might be communicated would recognize that what might be communicated is *mere opinion*: any interpretation created by non-omniscient beings may not be *worth* communicating because it may well be wrong.

We receive sense data that persuade us that there is an external world, but as Descartes pointed out, there are other possible interpretations of the same data. We receive sense data that persuade us that the world is flat, which is also an interpretation, and one that seems to be wrong. It is not happenstance that Kirzner's theory compels him to posit that entrepreneurs can ''smell'' profit opportunities (Kirzner 2006, 414), as if infallible sense data, not fallible interpretations of data, are at issue.[16] We doubt that human evolution, 99 percent of which took place before capitalist economies came into being, would have equipped anyone with a ''nose for profit opportunities'' (except in the tautological sense that *successful* entrepreneurs might be said, ex post and figuratively, to have such a faculty); we doubt, therefore, that *unsuccessful* entrepreneurs do *not* feel ex ante that they ''smell'' a profit opportunity that later turns out to have been nonexistent. Hiding under the false rubric of sensory perceptions are interpretations—some of which turn out to be less accurate than others. Similarly, consider the prices that, according to Hayek in ''The Use of Knowledge in Society,'' convey information: the only information that they *really* convey is that someone is offering something in exchange for a given quantity of money.[17] That this quantity reflects an intersection of ''supply and demand'' is an interpretation—and a somewhat inaccurate one, since there has been no infinite bargaining between every consumer in the world and every producer so as to produce the *correct* price point. There have been only recent sales to specific consumers at specific prices, and each sale might have taken place at a lower price or a higher one—as Stiglitz noted.

Moreover, the task of the socialist plant manager or the capitalist entrepreneur is to interpret the *implications for action* of the current array of prices: What do strong sales of iPhones at a given price mean that we should now do?[18] It could mean that there is strong demand for portable phones, portable Web browsers, browsers with many applications, browsers with particular applications, electronic fashion statements, or any of a number of other qualities that different consumers prize in an iPhone—and that, if asked, they might not be able accurately to pin down. One's decision about whether to produce particular models of iPhones in response to strong sales will depend on one's interpretation of what consumers are "really" looking for when they buy the product.

Assume that under either competitive capitalism or market socialism there is a shortage of iPhones at a given price; in response, Apple or the plant manager could simply raise the price. However, this mechanical response to "the data" would harm consumers unable or unwilling to pay the higher price of whatever qualities they would be willing to buy at a lower price. If as many consumers' desires as possible are truly to be met,[19] then the plant manager, or alternatively Apple and its would-be competitors, must try to interpret the *reason* for this shortage and then venture a hypothesis about the best way to satisfy the excess consumer demand. The hypothesis will involve inferences not only about the specific consumer demand that stands behind a product purchase, but about how best to meet this demand, possibly through different (say, unbundled) products or even *less* expensive ones. No matter how confident they may be that their interpretations of such complex matters are correct, either the entrepreneurs or the socialist plant manager may be wrong.

The market socialists were therefore right to point out against Mises and Hayek that *whatever* it is that allows capitalism to work could, *in principle*, be duplicated under socialism. There is a basic symmetry that reasserted itself at each stage of the socialist-calculation debate: whether in a socialist or a capitalist economy, decisions about how best to satisfy consumers are being made. But neither the market socialists nor the critics of socialism appreciated that at bottom, this is a problem of interpretation, not a problem of gathering data or performing calculations based on data. As Don Lavoie (1985, 123, emph. added) put it in his blow-by-blow account of the debate, capitalism "depends on the competition of separate private owners who *disagree* about which techniques are better" at satisfying consumer demand. They could

only disagree if the questions of which data are relevant and accurate, and how best to interpret them, were open questions. We note that entrepreneurs can and do use many different types of data—not merely prices but news reports, marketing research, single-store sales reports, aggregate statistics, reports from agents "on the spot," theories learned in business school, personal experience, hunches, and so on. The relevance of this or that datum depends on one's interpretive framework.

Disagreement as the Heart of Capitalism

When the market socialists conceded that the central planning board would have to allow surpluses and shortages to develop as a way of zeroing in on the "correct price," it was an opportunity for Mises or Hayek to point out that a mechanical price-raising or -lowering response is an abdication of the fundamentally interpretive function of any economic system. Yet even if they had made this argument, the market socialists could have conceded it, too, and included interpretation of *what to do about* shortages or surpluses (not merely raise or lower prices, but consider the production of alternative products) as part of the mandate of the central planning board or of those to whom the board delegated authority. At this juncture in the debate, perhaps Mises and Hayek would have recognized that *given human ignorance,* it is far better to have many fallible interpretations of what to do competing against each other than to risk everything on one centrally decided interpretation. The downside of competition among entrepreneurs who have different and fallible interpretations of "the data" is that some of them will necessarily be wrong, and will waste resources on their mistakes. The upside is that all the resources of an entire economy are not bet on the infallibility of the central planners' interpretation of what to do, or on the interpretation of a *single* manager to whom this authority has been delegated. (Capitalism, of course, simply *is* the delegation of this authority to anyone and everyone who can get ahold of some capital.)

Hayek ([1945] 1948, 77, emph. added) had dropped a hint in this direction on the first page of "The Use of Knowledge in Society," where he characterized "the knowledge of the circumstances of which we must make use" as existing "as the dispersed bits of incomplete *and frequently contradictory* knowledge which all the separate individuals possess." Only decades later, however, did he emphasize the implication of contradictory

"knowledge": that it is *not* knowledge, which, after all, could only *appear* to be contradictory. What *can* be contradictory are different interpretations of what *is* knowledge. Thus, Hayek maintained in 1968 that the reason for such an inherently wasteful method of decision making as competitive capitalism is our ignorance of which decision will succeed. If anyone could know that ex ante, "knowledge" would indeed be given and socialist central planners or their subordinates would know what to do:

> It is salutary to remember that, *whenever* the use of competition can be rationally justified, it is on the ground that we do *not* know in advance the facts that determine the actions of competitors. . . . It would clearly be pointless to arrange for competition, if we were certain beforehand who would do best. (Hayek 1968, 69)

By 1968 we think it is fair to say that Hayek was envisioning competitors who, whatever their location, do *not* necessarily possess knowledge, as opposed to opinions, or hypotheses, such that their conflicting versions of "knowledge" may have contradictory implications. Had Hayek pointed this out earlier, he might have placed the *prima facie* case for competitive capitalism on an epistemological foundation and perhaps turned economics in an epistemological direction, too. For the primacy of economic interpretation, and the desirability of competing interpretations, presupposes Descartes' fundamental acknowledgement that our "knowledge" may be wrong.

It is at least imaginable that if this had happened, then by 1961, most economists would not have retained such a simplistic view of the interpretation-free "givenness" of "information" that they could construct models of information search that conferred effective omniscience on (at least some of) the participants in capitalist economies. Since such models include no role for genuine, unforced error, however, the economists using them are ill suited to acknowledging the possibility that an initially ignorant market participant (or central planner, or bureaucratic regulator)[20] will *not* know what information he or she needs, or where to look for it, or how to interpret it even if he or she should find it.

The Irrationalist Theory of Economic Error

Because of the naïve epistemology, or nonepistemology, of mainstream neoclassical economics, economists who *do* recognize that error may be

due to factors other than asymmetrical data-hiding in the rational pursuit of self-interest have little alternative in *explaining* error but to make appeals to ''irrationality.'' Ignorance simply isn't part of their theoretical vocabulary.

We first distinguish between vulgar and sophisticated versions of psychological—as opposed to epistemological—explanations of error.

The vulgar usage of ''irrationality'' does not actually explain anything. In this usage, ''irrational'' is a synonym for *inexplicable*: when we casually say that people are behaving irrationally, we really mean that we don't understand why they are doing what they are doing.

Specifically, the vulgar usage of ''irrationality'' treats people's *mistakes* as inexplicable. For instance, one may retrospectively ponder the many investors and bankers who bought into the housing bubble and were devastated by its collapse, and if one is insufficiently attuned to the reality that ''the bubble'' was *not* an ''obvious fact,'' but rather was one possible interpretation of rising house prices, then one may end up asking why anyone would buy into a bubble and conclude that they must have been ''irrationally exuberant.''

In fact, though, the increase in house prices was interpreted differently by perfectly sane, rational observers who simply disagreed with each other about what this increase meant. The dominant, mistaken interpretation was that what we now retrospectively see as a ''bubble'' was in reality merely a ''boom'' that was driven by the growth of America's (and other countries') population and wealth. But the vulgar notion of irrationality treats objective, ex post reality as so easily accessible ex ante that it cannot acknowledge the possibility that behavior betokening a failure to understand a true, relevant, and important reality might be due simply to inadvertent ignorance, despite the incentives to know what is true, relevant, and important. Hence the non-explanation of such behavior as irrational, i.e., willfully ignorant of the allegedly transparent reality.[21] Like motivation (incentives)-obsessed economists, vulgar-irrationality theorists treat what would otherwise be considered inadvertent, *involuntary* mistakes as if they were deliberately (but derangedly) *willed*.

One way of summarizing our point is to say that the vulgar irrationality hypothesis is more normative (''they *should* have known it was a bubble'') than it is scientific (causally explanatory). Another way of summarizing it is to say that the attribution of error to irrationality is itself an (inadvertent) error—an exercise in hindsight bias:

> In hindsight, people consistently exaggerate what could have been anticipated in foresight. They not only tend to view what has happened as having been inevitable but also to view it as having appeared "relatively inevitable" before it happened. People believe that others *should have* been able to anticipate events much better than was actually the case. They even misremember their own predictions so as to exaggerate in hindsight what they knew in foresight. (Fischhoff 1982, 341, emph. added)

However, these marks of hindsight bias can also be seen in the non-vulgar, scholarly wing of the irrationality school. For instance, the leader of this school, Robert J. Shiller, argues with the benefit of hindsight (Shiller 2008, 39) that a certain Figure 2.1 in the 2005 edition of his book, *Irrational Exuberance*, demonstrated, at the time, that the increase in house prices since the late 1990s could not have been due to secular factors such as rising construction costs, population growth, or rising interest rates. He concludes that anyone who in 2005 relied on a continuation of the housing boom was being irrational. However, Shiller omits one of the secular factors that influenced many people to interpret the "boom" as being fundamentally sound, instead of being a "bubble": the rise in disposable income since World War II, a period during which house prices had, until recently, barely moved up at all. As Ben Bernanke (2005, 7) put the conventional wisdom,

> House prices have risen nearly 25 percent over the past two years. Although speculative activity has increased in some areas, at a national level these price increases largely reflect strong economic fundamentals, including robust growth in jobs and income, low mortgage rates, steady rates of household formation, and factors that limit the expansion of housing supply in some areas.

Given this interpretation of what was happening in the housing industry, someone who read Shiller's 2005 book when it was published (and it is important to note that, of course, most people did *not* read it) might reasonably have interpreted his Figure 2.1 to mean that long-pent-up housing demand was being expressed in higher house prices.

Indeed, despite his 2008 retrospective interpretation of his own Figure 2.1, back in 2005 Shiller (2005, 18) himself did not recognize that there was a nationwide housing bubble, as opposed to bubbles in a few "big glamour cities"—bubbles that were widely recognized as bubbles at the time, and that Shiller devoted a large part of the book to causally explaining with great insight.

Shiller characterizes precisely the same error, when made by others, in the language of ''psychological'' or ''emotional'' failure caused by ''*social contagion*'' (Shiller 2008, 41, emph. original), all of which suggests irrationality. This is the barren conceptual framework, but when Shiller actually describes the causal mechanisms at work, the behavior turns out to be *errors* that are not caused by any particular psychological pathology, and certainly not by the emotional processes suggested by such terms as ''exuberance.'' Rather, Shiller describes purely epistemological lapses that have *cultural*, not psychological, roots.

Shiller's ''irrationality'' is actually market participants' ignorance of important information, or market participants' misinterpretation of information. The ''social contagion'' is in fact a ''contagion of ideas'' (Shiller 2008, 41). Without the pseudo-medical terminology, a social ''contagion'' of ideas translates as the widespread transmission and acceptance of incorrect beliefs. When people convey to each other, or hear from the media, *sound* beliefs, we assume that Shiller would not call it a ''social contagion,'' since the very point of such terminology is to suggest pathology.

We agree with Shiller that people often get their beliefs socially, meaning from other people (although we prefer the term *culturally*, since the beliefs must be mediated by symbols). But when a socially transmitted idea is unsound, it bespeaks ignorance of some aspect of reality or logic on the part of the transmitter or the recipient of the idea, not irrationality on the part of anyone. To make an error in logic, let alone to be unaware of a fact, is not the same thing as to be irrational rather than merely to be mistaken. However, in the world of the economist, where all action is voluntary and thus susceptible to being influenced by incentives, it *would* be irrational to make such errors, because one would be doing so deliberately.

We find it highly suggestive that Shiller teamed up with Akerlof, the author of ''The Market for Lemons,'' to coauthor *Animal Spirits* (2009). *Animal Spirits* is, in effect, an answer to Stigler's call for a theory of error. Economists such as Shiller and Akerlof, who rightly reject the assumption that economic agents are *infallible*, have unfortunately equated this with the assumption that economic agents are *rational*. Accordingly, Akerlof and Shiller's concrete attempt to create an economics of *ignorant* and therefore error-prone agents is transformed, verbally, into a theory of economic agents' *irrational* ''animal spirits,'' leaving intact the assumption that the agents are effectively omniscient. Yet one's use of reason has nothing to do with how well informed or confused one is. Only if we

conflate rationality with the possession of perfect information, and then conflate the possession of perfect information with perfect interpretive abilities, does a rejection of the perfect-information assumption appear to entail ''irrationality'' as the alternative.

If economics is to avoid collapsing into a rather implausible version of psychology, an explicit embrace of (radical) ignorance would seem to be the only alternative that would account for widespread *mistakes*, such as occurred in the years preceding 2008.[22]

''Search'' versus ''Browse''

How would one undertake a neoclassical-economic search for an *unknown* unknown? Perhaps one would ''browse.''

The online journal repository J-STOR has two functions for researchers: ''search'' and ''browse.'' ''Search'' will check whichever term has been submitted against the directory, to see if any pages match. The directory is a pool of journal articles; every search goes through every article, touching upon all items within the data set. By contrast, the ''browse'' option permits the user to select a particular publication and meander through an edition of it. The browsing client sets off on a path between different items that is completely detached from the other articles contained within the database. By opting to wander through a particular journal, the user's research will necessarily neglect whatever may be contained within *other* journals. His or her opinion formation is no longer underpinned by the entire data set; of necessity, his or her opinions will now have blind spots.

This might seem to be an indictment of browsing, but only if we assume that the alternative, searching, does *not* produce blind spots. This assumption is false, however. A comprehensive ''search'' of an electronic database is possible only because the infinite universe of potential information has been severely curtailed by the database administrators.

A comprehensive search of an infinite universe being impossible, the illusion of comprehensiveness is created, in the case of J-STOR, by administrators who choose which journals to include, by journal editors and referees who choose which articles to publish, and by article authors who have chosen which topics are interesting due to *their* previous searching or browsing. Similarly, a Google search is limited by the choices of those who find this or that information interesting enough to put online, and then by those whose previous searching or browsing induces

them to click on this or that page, processed through Google's interpretive algorithm—and, finally, by the willingness of the user of Google to draw the line, somehow, after the first 10 or 100 or 1000 hits, meaning that he or she is trusting that what would be subjectively interesting enough to other Google users to have produced an early hit is roughly equivalent to the objective relevance (and importance) of the information in those hits for the present user's needs. Unless all the human beings involved in radically delimiting searches are omniscient, then we have no reason to believe that a J-STOR or Google search will discover the truth, as opposed to the truth about which terms are contained in the tiny subset of "information" about the world that is being searched.

Thus, searching is not necessarily any less free of error than browsing, even bracketing the reliability of the information contained in the database being searched. On the other hand, "browse" is by far a more suitable description than "search" of the ultimate determinants of people's opinions. The data sources among which they might search in shaping those beliefs are limited to the particular search terms they think would be relevant to pursue in a particular case. But what determines their understanding of "relevance"? If they do not already realize that housing is booming only due (for example) to the Fed's expansionary monetary policy, it would not occur to them to search for this information. Only if they had previously acquired a theory of business cycles that raised such suspicions would they go to the time and trouble to start their search. How, then, would they have acquired such a theory? In effect, by "browsing."

Perhaps they happened to have been interested in Austrian or Keynesian economics, and in reading this literature had acquired suspicions about booms that eventually go bust. Within the body of literature in which they were browsing because of their antecedent interest in Austrian or Keynesian economics, they *inadvertently* encountered an interpretation of business cycles that, years later, they remember and think might prove useful. The contrast with search is illustrated when we ask how one might have acquired the antecedent interest. In our example, it certainly would not be because one knew that years later, it would equip one to invest wisely in the face of an unprecedented housing bubble. For one could—at best—make such predictions about the usefulness of the literature only after having already read it.

No, one would first have to have had some "random" encounter with Austrian or Keynesian economics—some friend was interested in it, or it was taught in a course that was required or that sounded interesting (leading

back to previous "random" factors that might have made Economics 101 seem, *ex ante*, more interesting than History 101 at that particular moment in one's life). Only after having stumbled onto this or that knowledge might one be in a position to recognize, later on, the need for information that would confirm or falsify the applicability of this knowledge to a particular economic situation—that is, information relevant to a certain interpretation of the situation. Then one might "search" for these unknown knowns, such as data about current monetary policy or the financial condition of mortgagors. But the potential importance of the antecedent interest was, when it was acquired, an unknown unknown.

Thus, not only is the "search" model of learning dependent on previous "browsing," but a *rational* browse (hence a truly rational search) is impossible, because rationality connotes deliberate choice. The accidents that lead to a browsed discovery are *involuntary*, as are all accidents. Therefore, no incentives can affect them. However, if the discoveries—or the failure the to make them, due to equally accidental factors—are not the result of a rational process, that does not make them the results of an irrational process. *Rationality is no explanation of serendipity*, regardless of whether one was engaged in a rational or irrational activity when one accidentally made a discovery.

Browsing is only a weak metaphor for the actual process of opinion formation—regardless of whether the result is true or false opinions. The actual process is *life itself,* with all its pleasant and unpleasant surprises. We bump into ideas and "data" (whether true or false) that determine our future interpretations of what would be "interesting" to browse or "valuable" to search. The accidental process that shapes our lives determines our future ideas about which information is worth "searching" for, where to find it, and at what cost to get it. Nothing in this process guarantees that we will end up with good interpretations: Fallibility is part of the human condition.

This is the most powerful reason we know for pluralizing economic decision making among *competing* human interpretations.

An Economics Beyond Choice

From an epistemological perspective, what is the "life itself" that we have just described? It is a process of cultural and genetic conditioning—not a process of deliberate choice (whether "rational" or "irrational").

Information reaches us either from the precognitive organization of our neural apparatus, which is based ultimately on our genes; or from the conscious organization of our neural apparatus by means of culturally (or as Akerlof and Shiller say, "socially") mediated messages.

We are then left with the task of processing and ordering "the data": interpreting it. We do this by using our notions of what is true and important (i.e., what is relevant), and of how it fits together into a coherent whole. These notions, like "the data" themselves, have been transmitted to us either genetically, as intuitions, or explicitly by other human beings (contemporary or long dead). In this process of cultural and genetic influence, there is path dependence, there is structural amnesia (Douglas 1987), and for these and other reasons, there are gaps in the knowledge that passes, for us, as the truth, the whole truth, and nothing but the truth.

The progress from ignorance to knowledge is therefore anything but inevitable. Our ignorance as scholars may leave us convinced that we are making social-scientific progress when, in fact, we are simply reiterating earlier mistakes. Likewise, an economic actor may think he or she is pursuing a brilliant business opportunity that turns out to be a disaster. In the process, he or she may assiduously acquire all the new data that seem to be worth acquiring, yet still commit a grievous error.

To err is human. We should not build entire social-science disciplines on a denial of this fundamental truth, or when we turn from blackboard models to the decidedly imperfect markets—and polities—that these models are designed to illuminate, we will find ourselves unable to explain the manifold departures from perfection except by appealing to insanity, rather than to unchosen ignorance.

NOTES

1. It is a mistake because voters vote, so they must *not* realize the steep odds against their vote being decisive (which the theory of irrational voting counts as a decisive reason not to vote) (Downs 1957). The theory of rational ignorance may explain the political ignorance of *nonvoters*, but *voters* must either overestimate the odds that their vote will make a difference or they must feel a civic duty to vote. In the former case, they must feel compelled to inform themselves to their own satisfaction that they are casting an *accurate* vote; otherwise it would not make sense to *want* to cast the decisive vote. In the latter case, one's civic obligation to vote would not be fulfilled by casting an uninformed vote that could lead to undesirable social consequences.

2. Friedman and Kraus 2011, ch. 1, surveys some of the baseless assertions that prominent mainstream economists have made about the causes of the financial crisis; it traces these assertions to economists' discomfort with admitting that economic actors make genuine mistakes. Joseph E. Stiglitz (2010, 153, emph. original), for one, writes without irony that the supposed fact that "the disaster that grew from . . . *flawed incentives* can be, to us economists, somewhat comforting: our models predicted that there would be excessive risk-taking and short-sighted behavior, and what has happened has confirmed these predictions." Incentives, as we argue below, can affect only *deliberate* behavior; genuine errors can no more be made deliberately than can an actor deliberately decide not to learn something that, if known, would alter his or her behavior for the better, as the agent sees it.
3. However, see Friedman 1996.
4. Keynes (1921) and Knight (1921) made a key distinction between risk and uncertainty: risk can only be calculated if the events form a probability distribution. For unique events that are not part of any distribution, their probability of occurring is essentially unquantifiable. Thus, we are radically ignorant of their probability, which is to say genuinely ignorant about how to respond to their possibility. This is surely a realistic point that confronts human actors daily, yet economists refuse to deal with it.

   Stiglitz, whose work is so often thought to represent a bracing challenge to the unrealistic assumptions of mainstream neoclassical economics, actually goes out of his way to downplay the importance of the Keynes-Knight point, at least in the case of the financial crisis. In a review essay on Robert Skidelsky's *Keynes: The Return of the Master*, Stiglitz (2010b, 17) takes Skidelsky to task for emphasizing Keynes's sharp distinction between "situations in which we have good statistical data so that we can talk meaningfully about the probability that a particular event will happen," and *uncertainty*, i.e., ignorance (which Stiglitz does not even try to define). "Much of the behavior that led to the crisis," Stiglitz asserts, "did not depend on this distinction. More important, for instance, were the incentives, which encouraged banks to take on too much *risk*."
5. Following the example of Weber (1968) 1978, 24.
6. Radical ignorance often concerns one's own unstated assumptions. People who assumed that the housing boom was not a bubble probably did not tend to state it to themselves that explicitly, or else they might indeed have tried to search for information about whether it was a bubble.
7. Admittedly, the overabundance of information entails the scarcity of one's time in tracking it all down. So a precondition of radical ignorance can be reworded in the language of scarcity to make it more familiar to economists. But radical ignorance is not *reducible* to "scarcity of time resources" because the problem facing an agent in allocating these scarce resources cannot be solved rationally (or irrationally), due to the special nature of the object one would like to "acquire": *relevant, accurate, and properly interpreted* "information," or "knowledge" for short. Knowledge cannot be known until it is known, so however much one can use rational heuristics to try to guide one's knowledge searches, the strong-sense "rationality" of these heuristics—their success—can be determined only ex post,

while the path to success in a search for an unknown unknown is, by definition, unknown ex ante. Therefore the high incentive to come up with an excellent search heuristic in advance is irrelevant to the success of the search. Put differently, incentives can encourage efficacy only when the incented agent knows where to look, but if there unknown unknowns, one of the things we don't know about them is where to look for them; if we did, they would be known unknowns. If we knew where the next Black Swan was coming from, it would not be a Black Swan.

8. See n7 above.
9. Spence also assumes that the prospective employee knows in advance how much education will send a successful signal.
10. We are referring to whether a probability distribution is known to the economist, as opposed to the agent. For example, Rothschild (1974) provides a model for when the distribution of prices is not known to the agent, and they are forced to make inferences based on previous information. We do not doubt that one can inject some degree of realism in this way, but we question whether this truly solves the problem. Rothschild's result—that ''the qualitative properties of demand functions which arise from optimal search from unknown distributions are the same as those which arise from optimal search from known distributions'' (ibid., 710)—relies on analysis that is based on an underlying distribution (ibid., 694) and an assumption that the agents hold a prior belief that information follows a Dirichlet distribution.
11. Prychitko 1996, Boettke 1996, and Boettke 1997 also argue that Stiglitz conflates information and efficiency.
12. See Lavoie 1985, 160 and *passim*.
13. Although it is true that Stiglitz (1994) explicitly attacks the Arrow-Debreu assumptions of perfect knowledge, he tacitly revives them in defending the ability of the state to know how to correct market failures. See Boettke 1997.
14. Hayek ([1945] 1948, 82) also misperceived the epistemological problem as being brought about only by changing circumstances—as if unchanging circumstances present unmediated ''local knowledge'' to those who are on the scene, without the need for their interpretation of the *meaning* of this information, i.e., their need to interpret what to do about local circumstances. ''The shipper who earns his living from using otherwise empty or half-filled journeys of tramp-steamers'' (ibid.) might well *lose* money if his interpretation of what to put in those empty holds turns out to be unprofitable. See Friedman 2006, 495–96, for a critique of the same mistake in Kirzner's theory of entrepreneurship.
15. Cf. Friedman 2006.
16. Note that if this sixth sense existed, central planners could be selected on the basis of their possession of it, assuming that it could be identified. We could put Warren Buffett or George Soros in charge of the entire economy. However, if it cannot be identified except ex post, by virtue of certain entrepreneurs' profits, then calling it a sensory faculty is, even metaphorically, quite empty. This ''nose for profit opportunities'' is then merely *whatever* factors—genetic, cultural, or accidental—happen to have produced profits in a given case.

17. Indeed, strictly speaking, they do not convey any information at all. They are simply meaningless numbers until they are interpreted.
18. This raises the issue of "signal extraction" problems. "Austrian" business cycle theory points to the problems entrepreneurs face in understanding whether price increases are caused by rising demand or by inflation, and also the difficulty of estimating the permanency of such changes (e.g., Horwitz 2000, 108). These are interpretive problems. However, there is more than a slight tension between Austrian business-cycle theory's acknowledgement that monetary expansion may produce a cluster of interpretive errors by entrepreneurs, on the one hand, and on the other the lack of a microfoundational theory of entrepreneurial ignorance to explain how such errors could occur among businesses run by people who are "alert" to profit opportunities because of their incentive to be knowledgeable about such opportunities.
19. We posit this as desirable in the broadly utilitarian yet antipaternalist tradition of welfare economics, even though in reality consumers, too, may be mistaken in their desires, in principle justifying paternalism. In the real world of capitalism, however, mistaken consumers can "exit" from an error by returning the product that turns out to be unsatisfactory, or refusing to buy it again in the future (Friedman 2006). This remedy does not, of course, address the larger question, raised *inter alia* by Rousseau's *Discourse on Inequality* and Weber's *Protestant Ethic and the Spirit of Capitalism*, of whether the entire modern economy—capitalist *or* socialist—is an irreversible mistake, from a utilitarian perspective.
20. For an application of this argument to regulatory fallibility, see Friedman and Kraus 2011, ch. 4.
21. This is also the approach taken by Bryan Caplan (2007) in explaining popular economic opinions with which he disagrees. Economic truth is so self-evident to him that he "explains" the disagreement as being due to the public's irrationality, for which he provides no evidence except the very fact that the public disagrees with him (and with his fellow professional economists). Why, then, is the public's side of this disagreement called irrational, rather than simply incorrect? Because Caplan imputes to the public (without evidence) intuitive *agreement* with allegedly self-evident economic truths, which, Caplan infers (without evidence), the public then repudiates because it "feels good" to disagree with these truths. See Bennett and Friedman 2008. The reason for Caplan's "burlesque of the relevant psychology" (Quirk 2008, 351) is the lack of conceptual space in economics for divergent interpretations of a complex, opaque reality.
22. This section is adapted from a longer critique of "irrationality" economics in general, and its application to the financial crisis in particular, in Friedman and Kraus 2011, ch. 1.

REFERENCES

Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84: 488–500.

Akerlof, George A., and Robert J. Shiller. 2009. *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism*. Princeton: Princeton University Press.

Bennett, Stephen Earl, and Jeffrey Friedman. 2008. "The Irrelevance of Economic Theory to Understanding Economic Ignorance." *Critical Review* 20(1): 195–258.

Bhidé, Amar. 2010. *A Call for Judgment: Sensible Finance for a Dynamic Economy*. Oxford: Oxford University Press.

Boettke, Peter J. 1996. Review of Joseph Stiglitz, *Whither Socialism*. *Journal of Economic Literature* (March).

Boettke, Peter J. 1997. "Where Did Economics Go Wrong? Modern Economics as a Flight from Reality." *Critical Review* 11(1): 11–64.

Boettke, Peter J., and Christopher J. Coyne. 2009. "Context Matters: Institutions and Entrepreneurship." *Foundations and Trends in Entrepreneurship* 5(3): 135–209.

Caplan, Bryan. 2007. *The Myth of the Rational Voter: Why Voters Choose Bad Policies*. Princeton: Princeton University Press.

Douglas, Mary. 1987. *How Institutions Think*. London: Routledge.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: HarperCollins.

Fischhoff, Baruch. 1982. "For Those Condemned to Study the Past: Heuristics and Biases in Hindsight." In *Judgment Under Uncertainty: Heuristics and Biases*, ed. Daniel Kahneman, Paul Slovic, and Amos Tversky. Cambridge: Cambridge University Press.

Friedman, Jeffrey. 1996. Introduction to *The Rational-Choice Controversy: Economic Models of Politics Reconsidered*. New Haven: Yale University Press.

Friedman, Jeffrey. 2006. "Taking Ignorance Seriously: Rejoinder to Critics." *Critical Review* 18(4): 469–532.

Friedman, Jeffrey, and Wladimir, Kraus. 2011. *Engineering the Financial Crisis: Systemic Risk and the Failure of Regulation*. Philadelphia: University of Pennsylvania Press.

Greenwald, Bruce, and Joseph E. Stiglitz. 1986. "Externalities in Economies with Imperfect Information and Incomplete Markets." *Quarterly Journal of Economics* 101: 229–64.

Hahn, Frank. 1981. "General Equilibrium Theory." In *The Crisis in Economic Theory*, ed. Daniel Bell and Irving Kristol. New York: Basic Books.

Hayek, F. A. [1945] 1948. "The Use of Knowledge in Society." In idem, *Individualism and Economic Order*. Chicago: University of Chicago Press.

Hayek, F. A. 1968. "Competition as a Discovery Procedure." In idem, *New Studies in Philosophy, Politics, Economics, and the History of Ideas*. Chicago: University of Chicago Press.

Horwitz, Steven. 2000. *Microfoundations and Macroeconomics*. London: Routledge.

Ikeda, Sanford. 2003. "How Compatible Are Public Choice and Austrian Political Economy?" *Review of Austrian Economics* 16: 63–75.

Keynes, John Maynard. 1921. *A Treatise on Probability*. London: Macmillan.

Kirzner, Israel M. 1973. *Competition and Entrepreneurship*. University of Chicago Press.

Kirzner, Israel M. 1978. "Economics and Error." In *New Directions in Austrian Economics*, ed. Louis M. Spadaro. Kansas City: Sheed Andrews and McMeel.

Kirzner, Israel M. 1997. "Entrepreneurial Discovery and the Competitive Market Process: An Austrian Approach." *Journal of Economic Literature* 35: 60–85.

Kirzner, Israel M. 2006. "Hayek and Economic Ignorance: Reply to Friedman." *Critical Review* 18(4): 411–16.

Lavoie, Don. 1985. *Rivalry and Central Planning: The Socialist Calculation Debate Reconsidered*. Cambridge: Cambridge University Press.

von Mises, Ludwig Edler. [1920] 1937. "Economic Calculation in the Socialist Commonwealth." *Archiv für Sozialwissenschaften und Sozialpolitik* 47(1). Trans. in *Collectivist Economic Planning: Critical Studies on the Possibilities of Socialism*, ed. F. A. Hayek. London: George Routledge & Sons.

Prychitko, David L. 1996. Review of Joseph Stiglitz, *Whither Socialism? Cato Journal* 16 (Fall).

Quirk, Paul J. 2008. "Putting Experts in Their Place." *Critical Review* 20(1): 333–57.

Rothschild, Michael. 1974. "Searching for the Lowest Price when the Distribution of Prices is Unknown." *Journal of Political Economy* 82(4): 689–711.

Rothschild, Michael, and Joseph E. Stiglitz.. 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics* 90: 629–49.

Shackle, G.L.S. 1972. *Epistemics and Economics: A Critique of Economic Doctrines*. Cambridge: Cambridge University Press.

Shiller, Robert J. 2005. *Irrational Exuberance*, 2nd ed. New York: Broadway.

Shiller, Robert J. 2008. *The Subprime Solution: How Today's Global Financial Crisis Happened, and What to Do about It*. Princeton: Princeton University Press.

Somin, Ilya. 1998. "Voter Ignorance and the Democratic Ideal." *Critical Review* 12(4): 413–58.

Somin, Ilya. 2012. *Democracy and Political Ignorance*. Cambridge, Mass.: Harvard University Press.

Spence, Michael. 1973. "Job Market Signalling." *Quarterly Journal of Economics* 87: 355–74.

Stigler, George J. 1961. "The Economics of Information." *Journal of Political Economy* 69: 213–25.

Stigler, George J. 1976. "The Xistence of X-Efficiency." *American Economic Review* 1976: 216.

Stiglitz, Joseph E. 2008. "Information" *The Concise Encyclopaedia of Economics*. Library of Economics and Liberty. http://www.econlib.org/library/Enc/Information.html

Stiglitz, Joseph E. 2010. *Freefall: America, Free Markets, and the Sinking of the World Economy*. New York: Norton.

Stiglitz, Joseph E., and Andrew Weiss. 1981. "Credit Rationing in Markets with Imperfect Information." *American Economic Review* 71: 393–410.

Weber, Max. [1968] 1978. *Economy and Society*, vol. 1. ed. Guenther Roth and Claus Wittich. Berkeley: University of California Press.