12-1-2013

Soros's Reflexivity Concept in a Complex World: Cauchy Distributions, Rational Expectations, and Rational Addiction

John B. Davis

Marquette University, john.davis@marquette.edu

Soros's Reflexivity Concept in a Complex World: Cauchy distributions, rational expectations, and rational addiction[1]

John B. Davis, Marquette University and University of Amsterdam

Abstract: George Soros makes an important analytical contribution to understanding the concept of reflexivity in social science by explaining reflexivity in terms of how his cognitive and manipulative causal functions are connected to one another by a pair of feedback loops (Soros, 2013). Here I put aside the issue of how the natural sciences and social sciences are related, an issue he discusses, and focus on how his thinking applies in economics. I argue that standard economics assumes a 'classical' view of the world in which knowledge and action are independent, but that we live in a complex reflexive world in which knowledge and action are interdependent. I argue Soros's view provides a reflexivity critique of the efficient market hypothesis seen as depending on untenable claims about the nature of random phenomena and the nature of economic agents. Regarding the former, I develop this critique in terms of Cauchy distributions; regarding the latter I develop it in terms of rational expectations and rational addiction reasoning.

Keywords: Soros, reflexivity, efficient market hypothesis, Cauchy distributions, rational expectations, rational addiction

JEL codes: B41, D84, G14

1   Soros's cognitive and manipulative causal functions

Soros's analysis explains reflexivity in terms of how cognitive and manipulative causal functions are connected to one another by a pair of feedback loops (Soros, 2013). The cognitive function concerns our knowledge of the world and the role that the world itself plays in determining our knowledge of it. The manipulative function concerns action and how people's motivations given their knowledge of the world determine their actions. His method of explanation is to have us first imagine each function *as if* it operated in isolation from the other, and then have us see how that function is modified by a feedback loop from the other function.

(i) In the case of the cognitive or knowledge function, the direction of *causation is from the world to the mind.* Taken in isolation, our knowledge of the world is determined by what the world is independently of us, *as if* what the world is were fully independent of our actions:

$$\text{world} \rightarrow \text{mind} \qquad (1)$$

But in a complex reflexive world the manipulative or action function interferes with the cognitive function because people act on what they know, action changes what the world is, and thus changes knowledge. Action thus has a feedback effect on knowledge so that:

$$\text{world} \leftarrow \text{mind} \qquad (2)$$

So it incorrect to say as in (1) that the world is fully independent of us and independently determines our knowledge of it. The main direction of causation when we think of knowledge is still from the world to the mind, but the feedback loop from action in (2) produces a reverse causation that modifies the world and thus knowledge.[2] The cognitive function modified by manipulative function might thus be represented as:

$$\text{world} \rightarrow \text{mind [world} \leftarrow \text{mind]}$$

$$\text{or} \qquad (3)$$

$$\text{'world} \rightarrow \text{mind'}$$

(ii) In the case of the manipulative function, the direction of *causation is from the mind to the world.* Taken in isolation, people act on what they know *as if* their motivations in doing so were fully independent of their actions.

---

[2] Positive feedback effects of action on knowledge that reinforce knowledge have been termed self-fulfilling prophecies (Merton, 1948).

$$\text{mind} \rightarrow \text{world} \qquad (4)$$

But the cognitive function interferes with this because acting on the world also produces knowledge about one's motivations for acting, so that the basis for action is changed.

$$\text{mind} \leftarrow \text{world} \qquad (5)$$

So action depends not only on our motivations as in (4) but also on how our motivations change as we learn about them in acting upon the world as in (5), and so it is incorrect to say that the basis for action is independent of our actions. The main direction of causation is still from mind to the world – our motivations still basically drive our actions – but the feedback loop involving the reverse causation from knowledge gained in acting modifies motivation.[3] The manipulative function modified by cognitive function might thus be represented as:

$$\text{mind} \rightarrow \text{world } [\text{mind} \leftarrow \text{world}]$$

$$\text{or} \qquad (6)$$

$$\text{'mind} \rightarrow \text{world'}$$

(iii) Putting these two functions together with their respective feedback effects – the cognitive function modified by the manipulative one (3) and the manipulative function is modified by the cognitive one (6) – we have the following schema in a complex reflexive world:

$$\text{'world} \rightarrow \text{mind'} \text{ and } \text{'mind} \rightarrow \text{world'} \quad (7)$$

Thus, Soros explains reflexivity analytically by first isolating the two basic causal relationships concerning knowledge and action, and then showing how they are modified and interconnected by two separate feedback relationships.

2   A 'classical' world vs. a complex reflexive world: The efficient market hypothesis

(i) Soros reasons that what I call a 'classical' view of the world, in which knowledge and action are independent and can be investigated separately from one another, is an untenable simplification in a

---

[3] This second form of reflexivity is less commonly discussed, in no small part because it is associated with endogenous preferences, a subject assiduously left outside standard economics.

complex reflexive world of two-way feedback loops. His analysis offers two linked dimensions through which reflexivity can be investigated.

The first is an epistemological one in which the focus is the modified cognitive/knowledge function as in (3), and where the main subject of investigation is the fallibility of scientific knowledge. Soros points out that there are multiple epistemic sources of fallibility associated with the complexity of the world and the structure of the brain (p. 5). I associate the first with radical uncertainty (in the Keynes-Knight sense) and the latter with bounded rationality (in the Simon sense). But Soros's principle message in his contribution to this issue is that the reflexive feedback principle associated with how the manipulative function modifies the cognitive one is the fundamental source of the fallibility of scientific knowledge.

The second dimension of the analysis is an ontological one where the focus is the modified manipulative/action function as in (6), and the main subject of investigation is how agents act in a world of radical uncertainty. Uncertainty is often seen as an epistemological concept associated with the impossibility of knowing the future.[4] But it also needs to be seen as an ontological concept associated with not having a basis for action.[5] Soros also points out that there are multiple sources of uncertainty, but again his principle message in his paper here emphasizes the special role of reflexivity, or how feedback from the cognitive function interferes with the manipulative one in unsettling the basis for action.

The 'classical' view of the world of standard economics, then, entails an epistemology and ontology in which knowledge and action are independent.[6] Since this is the world in which the efficient market hypothesis (EMH) applies, Soros's reflexivity analysis should show us why the EMH cannot apply in a complex reflexive world.

---

[4] As J. M. Keynes put it: "The outstanding fact is the extreme precariousness of *the basis of knowledge* on which our estimates of prospective yield have to be made. Our knowledge of the factors which will govern the yield of an investment some years hence is usually very slight and often negligible. If we speak frankly, we have to admit that our *basis of knowledge* for estimating the yield ten years hence of a railway, a copper mine, a textile factory, the goodwill of a patent medicine, an Atlantic liner, a building in the City of London amounts to little and sometimes to nothing; or even five years hence" (1936, pp. 149-50; emphasis added).

[5] As Keynes also put it: "Most, probably, of our decisions to do something positive, the full consequences of which will be drawn out over many days to come, can only be taken as a result of animal spirits—of a spontaneous *urge to action rather than inaction*, and not as the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities" (1936, p. 161; emphasis added).

[6] The 'classical' view of the world entails a natural science view of social science. In natural science, human action does not change the laws of nature, so the cognitive function is unmodified by the manipulative one. Thus standard economics employs a natural science view of social science.

The EMH, whether in its strong, semi-strong, or weak versions, basically says that in financial markets agents use 'all available information' about market prices efficiently. Behavioral finance economists generally argue this does not occur because people are subject to various psychological imperfections in reasoning or heuristics and biases that distort their information processing. Soros accepts this but for him the main reason the EMH is wrong is more fundamental: it ignores reflexivity. On the one hand, the EMH assumes traders' cognitive function is not modified by the manipulative function (1), and that accordingly 'all available information' can be reasonably taken to correspond to the way the world is. But, Soros tells us, the operation of financial markets affects what the world is because trading can distort the so-called underlying asset fundamentals which market prices are thought to reflect (2). So it is incorrect to say that 'all available information' corresponds to the way the world is independently of us. That is, traders' cognitive function is modified by their manipulative one (3). On the other hand, the EMH assumes traders' manipulative function is not modified by the cognitive function (4), from which it follows that traders are, as it were, motivationally complete (or have complete preferences). But, Soros tells us, turning points in boom-bust processes or other price cycles provide traders information about the consequences of their actions, which can change attitudes, and thus change traders' motivations. So it is incorrect to say that traders have 'all available information' when they set out to make any given set of trades since that information does not include information about how their motivations will change, or how their manipulative function is modified by their cognitive one (6).

The EMH is wrong, then, because it ignores the two-way feedback loops that operate on the two causal functions (7), whereas proponents of the EMH assume we live in a simple 'classical' world in which knowledge and action are independent of one another so that their effects can be investigated separately. This 'classical' thinking rests on two questionable ontological claims, one pertaining to the nature of random phenomena (the random walk hypothesis) and the other pertaining to the nature of economic agents (the expected utility conception). Consider first the claim about the nature of random phenomena.

i   The random walk hypothesis

The random walk hypothesis says that stock prices evolve in a random manner and thus cannot be predicted. Randomly distributed phenomena which are normally or bell-shaped distributed possess

determinate means and finite variance, and these properties are needed to be able to say that economics agents can use all available information efficiently. So the random walk hypothesis requires normally distributed stock prices. In contrast, if random phenomena have a Cauchy distribution (a family of fat tail distributions), mean and variance are undefined, and it no longer makes sense to say that economics agents use all available information efficiently. The 'classical' view of the world and the EMH rule out Cauchy distributions by assumption, but the recent financial crisis seems to have shown that stock prices are not always normally distributed. That is, one of the clear lessons of the crisis is that fat tail 'extreme' events, that is, phenomena which are rare when we assume normal distributions, are somehow not all that rare after all.

Suppose, then, that stock prices can be Cauchy-distributed with undefined means and variances. Why we might suppose this can be explained in terms of Soros's understanding of reflexivity in terms of the modified cognitive and manipulative functions (3 and 6). Soros's view is that the two functions are dominant causal processes that are only modified by feedback loops. Thus it is can easily *appear* that the pure cognitive function (1) and the pure manipulative function (4) accurately represent the phenomena (particularly when methods of observation are imperfect). Note, then, that normal distributions presuppose the Central Limit Theorem, which says that for a sufficiently large number of observations of a phenomenon the distribution of the phenomena converges on a normal distribution. But when the cognitive and manipulative functions are modified by two-way feedback (7), in reality the phenomena are changed, and so there no longer are a large number of the identical phenomenon. Soros thus tells us that the fundamentals behind stock prices change in trading, but his view is actually more subtle. It is that fundamentals are most of the time only modestly changed by the action of traders so that they may well appear not to have changed at all (at least until boom and bust processes become undeniable). So it may appear that the phenomena are the same and unchanged, the central limit theorem applies, and random phenomena are normally distributed when in fact this is not the case.

Indeed, this deceptive combination of the appearance of normality and the emergence of crisis is reflected in the idea of a fat tail 'extreme' event occurrence. How can an 'extreme' event both be truly not rare (because it occurs more frequently than expected) and nonetheless still be regarded as 'extreme'? The answer is that reflexive worlds mask their rare but nonetheless very possible non-normal behavior with the appearance of normality, because feedback relationships only modify market actors' pure cognitive and pure manipulative functions gradually and often imperceptibly.

Nonetheless these feedback effects can still be real and cumulative, and market phenomena – particularly where there are increasingly strong incentives for agents to act on their knowledge and affect the phenomena – should accordingly not be regarded as always normally distributed. Thus the EMH thus rests on a mistaken claim about claim about random phenomena.

ii   The Homo economicus expected utility conception

Now consider what the EMH requires about the nature of economic agents. There are two sides to the standard conception: the probabilities and knowledge side and the preferences and motivation side. On the former, the *Homo economicus* expected utility individual conception is objectivist about probabilities and knowledge. It assumes that learning occurs in a Bayesian way given common priors.[7] Everyone starts with the same knowledge of the world, and people acquire different posterior beliefs according to their differential experience. Further, since expected utility depends on forecasts of future prices, and since forecasts are subject to error, the standard view crucially also assumes that *errors in forecasting are uncorrelated across agents*, so that one can aggregate individual expectations into a representative agent rational forecast – the rational expectations hypothesis – and say that people acquire objective knowledge about the world because on average they use probabilistic information efficiently (see Hommes, 2013). Bayesianism thus employs a pure cognitive function feedback analysis since the assumption that forecasting errors are uncorrelated leaves this up-dating process fully in the domain of probabilities and knowledge, and the cognitive function is unmodified by the manipulative one (1).

One way in which this view has been contested in reflexivity terms is through the development of bounded rationality agent conceptions employing expectations feedback systems with correlated forecasting errors. In terms of Soros's framework, expectations feedback systems with correlated forecasting errors show how agents' cognitive function – here their forecasting function – is modified by the manipulative one (3). If agents' forecasts are correlated, they do not balance out on average (ruling out the representative agent rational forecast), and what people do (the actions they undertake) changes the overall pattern of future forecasts. This rules out the 'classical' rational expectations view that knowledge and action are independent. It also rules out the important associated assumption that agents are independent with respect to knowledge. Soros brings out this

---

[7] See Morris (1995) for a critical discussion of the common prior probability distribution assumption.

latter assumption in his emphasis on interaction between multiple agents. Brock and Hommes (1997, 1998) have developed this framework in systematic way by introducing heterogeneous expectations into expectations feedback systems with correlated forecasting errors, and by showing that when agents revise their forecasts of future prices, and switch back and forth between different forecasting rules based on their relative performance, excess volatility with bubbles and crashes in asset markets can result.[8]

But bounded rationality analysis strictly concerns knowledge and belief; that is, how the cognitive function is modified by the manipulative one (3). A quite different and often overlooked reflexivity analysis applies on the preferences and motivation side of the expected utility conception where the issue is how the manipulative function is modified by the cognitive one 6). The *Homo economicus* conception, then, is objectivist about probabilities and knowledge, but subjectivist about preferences and motivation. Utilities are fully personal. In contrast to the common priors assumption on the knowledge and probabilities side, people's preferences have no common basis and are individual-specific (though people can accidentally happen to have the same preferences). This has meant for most standard economists that agents have no preference adjustment process, process of preference formation, or preference feedback system, analogous to expectations feedback systems. Generally speaking, saying preferences are subjective means they have an exogenous character, or are given and unchanging.[9]

But looked at more closely, this is an odd conception. It assumes that a person's preferences constitute a finite set, and, according to the completeness axiom, are sufficient to evaluate all choice opportunities. People encounter new choices but since their preferences do not change, they must evaluate these new choices by means of the same preferences they employed to evaluate their past choices. Agents' given preferences, then, are not specific to the circumstance of choice, but have a kind of general, all-purpose nature or versatility for making an infinite number of possible choices. A particular preference *x*, that is, can be exercised to make any number of different choices *A, B, …*

---

[8] See Hommes (2013) for a discussion of the literature on expectations feedback systems as a response to rational expectations modeling. Note that agents change their mix of forecasting rules but do not change the motivational basis of their choices. The analysis is still reflexive in a way that Bayesianism is not because action feeds back onto belief. Hommes stresses the fact that in positive expectations feedback systems non-rational almost self-fulfilling equilibria with correlated forecasts easily arise.

[9] Stigler and Becker (1977) influentially argued that preferences, like the "Rocky Mountains" are just there, and thus essentially do not change. This assumption serves to isolate the effects of changes in prices and incomes on choice, and gives a market-driven analysis of choice. They also feared that if were explained by changing preference, the explanation of choice would become hopelessly confused by competing theories of human psychology.

N, though those choices can differ substantially due to differences in the person's information about the circumstances of choice. We might thus ask: why should *preference x* provide the means for making any particular choice versus any other choice? That is, what explains its versatility? The answer is that since preferences are subjective, that is, essentially unexplainable, this question simply cannot be raised.

One way, then, in which it may seem this view has been contested is in the analysis of fads, fashion, and related consumer phenomena as found in the recent literature on information cascades and 'herding' behavior (Banerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992). But these models are explanations of 'rational' herding behavior, where agents substitute public information for private information, and preferences are unaffected by choice. Thus they only concern a pure cognitive function (1), and tell us nothing about how the manipulative function is modified by the cognitive one (6) – which would seem to be the point of talking about fads and fashion.

Where the manipulative function has actually been the focus is in an older literature on endogenous preference. Preferences are endogenous and an intertemporal preference formation process – a preference feedback system – exists when individuals' preferences somehow depend on their past choices. Here, however, the standard analysis – the 'rational addiction' model (Becker and Murphy, 1988) – assumes individuals' preferences are dynamically consistent across time in that their later preference sets are rationally derived from their earlier ones. This rules out the possibility that a person might have dynamically *in*consistent preferences, and dissolve into a series of disconnected multiple selves each associated with a different set of (present-biased) preferences. It essentially allows the endogenous preference feedback system to operate as if the person's preferences over time were the person's representative individual preferences.[10] This analogue or preference side representative individual conception also isolates the person's preference change process from the influence of other individuals, and thus essentially employs a pure manipulative function (4). Thus it also fails to get at Soros's understanding of reflexivity.

However, recent behavioral economics research provides strong evidence that people commonly employ hyperbolic rather than exponential time rates of discount, and thus exhibit present-bias in decision-making (Frederick, Loewenstein, and O'Donoghue, 2002). This implies that their preferences are dynamically inconsistent, that people effectively function as successions of multiple

---

[10] In effect, it is as if the person's different preference sets over time are 'uncorrelated' in the sense of sustaining dynamic consistency of choice of the person's representative individual.

selves, and that a rational addiction representative individual assumption is not a reasonable way to model preference feedback systems and endogenous preference.[11] Thus, in the multiple selves time-inconsistent behavior literature (e.g., Laibson, 1997) no such assumption is employed, present-bias is interpreted as a lack of self-control (on the part of the representative individual planner self), and people are said to be reliant on external commitment devices, such as mandated 401(k) plans, to overcome present-bias. Recall that on Soros's understanding of how the manipulative function is modified by the cognitive one that the motivational basis for action is not independent of our actions; that is, we learn about ourselves when we act and this subsequently changes our motivations for acting. It might be argued, then, that generally people become aware of their present-bias and its sometimes unhappy consequences, and are then disposed to avail themselves of external commitment devices in forming future preferences. In this way the individual's preference feedback system functions in such a way that the manipulative function can be said to be modified by the cognitive one (6).

More certainly needs to be said about the nature of endogenous preferences to adequately explain preference feedback systems. But in closing this section I simply note that replacing the expected utility individual conception with a reflexive individual conception depends on bringing together the two sides of the agent discussed above. That is, a full reflexivity account of agents would need to integrate expectations feedback systems and preference feedback systems by showing how each acts upon the other, and thus how both the cognitive and manipulative functions modify each another in a two-way feedback system.

3   *Reflexivity and the EMH*

Soros's analysis makes an important analytical contribution by explaining reflexivity as a two-way feedback system between knowledge and action. The basic idea is that people acquire knowledge, act on what they know, this changes the world and what they can know about it, and also changes what they know about the basis for action and how they will subsequently act. This means there are always both 'new' things true about the world and 'new grounds for action, and that time's arrow

---

[11] When individuals exhibit present-bias, their preferences are no longer uncorrelated over time but rather correlated around periods.

has a one-way direction.[12] This is ruled out in the 'classical' world of the EMH in which knowledge and action lack feedback channels. In standard economics, this is reflected in how the random walk hypothesis and the *Homo economicus* expected utility conception – the two pillars on which the EMH depends – assume that the phenomena of the world and economic agents constitute two bracketed-off ontologically independent domains.

$$Homo\ economicus \rightarrow \quad EMH \quad \leftarrow random\ walk$$
$$\downarrow \qquad\qquad\qquad \downarrow$$
$$[agents] \qquad\qquad [phenomena\ of\ the\ world]$$
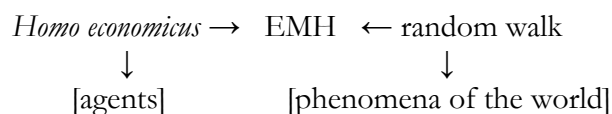
But in a complex economic world with reflexive feedback relationships between knowledge and action, agents and the phenomena of the world are ontologically interdependent and mutually influencing:

$$agents \quad \leftrightarrow \quad phenomena\ of\ the\ world$$

This is the vision Soros employs. Setting out this alternative ontological conception, then, is part of what is involved in explaining what reflexivity in economics entails.

References

Becker, Gary and Kevin Murphy (1988) "A Theory of Rational Addiction," *Journal of Political Economy* 96: 675-700.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy* 100 (5): 992-1026.

Brock, William A. and Cars Hommes (1997) "A Rational Route to Randomness," *Econometrica* 65: 1059-1095.

---

[12] The arrow of time idea originates with Eddington (1928) as a property of entropy. As he put it, "If your theory is found to be against the second law of Thermodynamics I can give you no hope; there is nothing for it but to collapse in deepest humiliation" (1928, p. 74).

Brock, William A. and Cars Hommes (1998) "Heterogeneous beliefs and routes to chaos in a simple asset pricing model," *Journal of Economic Dynamics & Control* 22: 1235-1274.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue (2002) "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40 (2): 351-401.

Eddington, Arthur S. (1928) *The Nature of the Physical World*, New York: Macmillan.

Hommes, Cars (2013) "Reflexivity, Expectations Feedback and Almost Self-Fulfilling Equilibria: Economic Theory, Empirical Evidence and Laboratory Experiments," *Journal of Economic Methodology* 20 (4).

Keynes, J. M. (1936) *The General Theory of Employment, Interest and Money*, London: Macmillan.

Merton, Robert K. (1948), "The Self-Fulfilling Prophecy," *Antioch Review* 8, 193-210.

Morris, Stephen (11995) "The Common Prior Assumption in Economic Theory," *Economics and Philosophy* 11: 227-253.

Soros, George (2013) "Fallibility, Reflexivity and the Human Uncertainty Principle," *Journal of Economic Methodology* 20 (4).

Stigler, George and Gary Becker (1977) "De Gustibus non est disputandum," *American Economic Review* 67 (2): 76-90.