

Rational Markets, Irrational Actors

Thesis: How can markets be efficient if economic actors are irrational? Irrationality is the result of high subjective costs of transaction, which can be reduced through a variety of social and cognitive mechanisms. These mechanisms allow for market efficiency even with actors prone to highly irrational decision-making.

Neoclassical decision theory has been able to explain various economic phenomena by making use of a psychologically implausible decision theory. How is this possible? Several authors have suggested that the structure of social and economic institutions is responsible for allowing real-world economic actors to approximate the traditional homo economicus model. Institutions effectively limit actors' choices, leaving them with only rational options to select among. Yet the mechanism by which choice limitation is achieved remains obscure.

The goal of this paper is to expand on and modify existing models of institutional cognition. Many models suggest that institutions produce effective economic decisions by limiting choices, a success-through-elimination model. But this explanation is too simple. Instead, effective social and economic institutions reduce *costs of cognition* through providing agents with the information they need to make rational decisions. The information provided is usually inadequate but the effects of this meager information can lead a market to converge on an economically rational outcome. Institutions that effectively reduce cognition costs allow economic actors to coordinate effectively, which then through imitation and adaptation enables markets to clear as they would were they populated with rational actors.

This essays proceeds in four parts. In Section I, I review the structure of neoclassical decision theory and the psychological challenges posed against it. In Section II, I discuss some of the proposed remedies to the conflict between economic explanations and the false decision theory underlying those explanations and critique them. In Section III, I will advance my own "costs of cognition" approach and show how my model can accommodate the benefits of the models discussed in Section II. In Section IV, I conclude.

Section I: Neoclassical Decision Theory and Its Critics

(I.i) Most economists employ traditional neoclassical decision theory in their economic modeling¹; traditional neoclassical decision theory is an *ordinal utility theory*. Economists

¹ This said, it is important not to overgeneralize about the ubiquity of this form of the use of neoclassical decision theory in economics; not only is neoclassical decision theory employed in different ways by

agree that cardinal measures of utility are infeasible and that interpersonal comparisons of utility are impossible or at least impracticable. Instead, economics focus on understanding rationality in terms of consistent preference orderings, where consistency is defined by an appeal to certain rules of rationality. Second, traditional neoclassical decision theory is an account of *instrumental* rationality. It makes no rational assessment of ends, only of means and specifies what the implications of an individual's beliefs and preferences are. The traditional model suggests that an individual is rational when he has an ordinal utility function, or a consistent ranking and weighting of all his preferences. In order to generate a consistent utility function, an individual's preferences must be *weakly ordered*. There are four standard conditions for a weak ordering.²

First, an ordering must be *complete*. If Arthur has a complete ordering, he can rank any two preferences vis-à-vis one another, either where $A > B$, $B > A$, or even that he is indifferent between them. Second, the preference relation must be *asymmetric*: if Arthur prefers A to B, then he cannot prefer B to A. The indifference relation is *symmetric*. If Arthur is indifferent between A and B, he must also be indifferent between B and A. Finally, preferences must be *transitive*. If Arthur prefers A to B and B to C, he must prefer A to C.³ If an individual's preference ordering is complete and transitive, the preference relation is symmetric and the indifference relation asymmetric, then the individual can be assigned an ordinal utility function. When he acts in accordance with this function, he is said to be rational.

(I.ii) Economics call a more specific version of this rational individual, *homo economicus*.⁴ Homo economics is understood to have an ordinal utility function with several features. First, *she prefers more to less*. This is a *ceteris paribus* condition. If Arthur wants a pizza, he prefers more pizza to less *all things equal*. Suppose that for \$10 Arthur can have two pizzas or one pizza. If he chooses one pizza we would call him irrational. Even if he does not want to

economists in different subfields, other subfields use modified versions of it. For more on the plural understanding of rationality used by economists, see Tyler Cowen, "How Do Economists Think About Rationality?," *Satisficing and Maximizing*, 213-236. The debate over the rationality of markets, however, has usually focused on the use of neoclassical decision theory, so our scope is restricted to those areas of economics which employ the theory in its most unmodified form and to the extent that other sub-fields of economics employ traditional neoclassical decision theory, to that extent will what I say here apply.

² I draw on the standard list provided by Gerald Gaus, *On Politics, Philosophy and Economics*, 36-37.

³ This requirement is often defended by the 'money-pump' argument, or the claim that if an individual had intransitive preferences. She could be swindled out of money continuously.

⁴ The following list of features of homo economicus is given by Gerald Gaus, *ibid*, 1.3, 19-27.

eat the second pizza, he may be able to find a use for it, perhaps to be eaten later, given to a friend or the homeless, or so on. One way of putting it formally is to say, following Russell Hardin, that “the simplest definition of rationality ... is that one should choose more rather than less value.”⁵ Next, homo economicus makes choices *at the margin*. Thus, if she is consuming units of pizza, she will prefer the first piece more than the second piece, the second more than the third, and so on. This is what economists know as the *law of diminishing marginal utility*. The value we place on goods diminishes for each additional good that we acquire. The rationale for the law is easily illustrated. Suppose that I am a farmer and I need grain to meet a variety of my ends. The first bag of grain that I acquire, I will turn to my most vital end, perhaps eating. The second bag of grain will go to sustaining my livestock, so that they do not die. And from there I meet further goals, saving grain for future eating, saving grain for feeding my livestock in the future and eventually saving grain in order to sell at the market.⁶ My behavior would be irrational if I did not explicitly or implicitly act in this way. I would effectively act to satisfy a preference that I rank lower than another preference I could satisfy just as easily. My action would thereby reveal a preference ordering inconsistent with my actual ordering. A third rule homo economicus follows is that she has *downward sloping demand curves*. In other words, the more that a good costs, the less that it is demanded. As the price of X increases, the demand for X decreases. This is true not only on the aggregate level but of the individual. This is because of the familiar economic concept of *opportunity costs*. The cost of any action A is the benefit of the next most beneficial feasible action B. B is the *opportunity cost* of A. As the opportunity cost of A increases, A will be demanded less by *homo economicus*.

All three of these rules are essential to the *homo economicus* model and are widely agreed upon as such by the economics profession. But there are two more assumptions that are built-in. First, there is the notion of *wealth maximization* or *self-interested action*. Homo economicus acts to maximize the satisfaction of her preferences, to be sure. This is an essential feature of a rational actor. But typically homo economicus is taken to have *self-interested preferences*. In other words, she seeks to maximize her psychological satisfaction, often through the accumulation of wealth. Economists make simplifying assumption as the accumulation of wealth is the primary purpose of economic transactions.⁷ Finally, and this

⁵ Russell Hardin, *Indeterminacy and Society*, 16.

⁶ The grain example is well-known. See Eugen von Bohm-Bawerk, *The Positive Theory of Capital*, Bk III, Ch. IV.11.

⁷ It is not essential to homo economicus that her actions be self-interested, however, for she may have altruistic preferences.

will be important later, homo economicus may sometimes decide not to maximize the satisfaction of her preferences. Herbert Simon, in his work on bounded rationality, argued that agents are often rational to ‘satisfice’ rather than maximize.⁸ Satisficing is choosing a solution that is ‘good enough’ rather than perfect in order to economize on costs of cognition. Gerald Gaus notes that “it is not always clear whether satisficing is really an alternative to maximizing.”⁹ Sometimes those who maximize do best to satisfice as they will waste time calculating.

(I.iii) Expected utility theory applies the homo economicus model to decision-making under risk. Expected utility theories multiply expected utility of an outcome by the probability of the outcome and hold that rational individuals will select the greater multiple than the lesser. Thus, even if option B produces more utility than option A, if option A has a sufficiently high probability, and B sufficiently low, a rational individual will select option A over option B. Since homo economicus maximizes wealth, in choices under risk, homo economicus will maximize her expected value, choosing the highest probability-utility multiple available to her. Expected utility theory is a straightforward implication of the homo economicus model when applied to decision-making under risk.¹⁰ All of the choice axioms that apply above apply to standard expected utility theory.

(I.iv) The homo economicus model has come under a variety of important criticisms¹¹, and I cannot hope to cover them all. But prominent examples often derive from cognitive

⁸ See Herbert Simon, 1959, “Theories of Decision Making in Economics and Behavioral Science.” *American Economic Review*, 49, 258-283. Some hold that satisficing is rational as an *alternative* to maximizing not just as a *method* of maximizing. See David Schmitz, *Rational Choice and Moral Agency*, Chapter 2, “When Satisficing Is Rational,” 31-37.

⁹ Gaus, *Politics, Philosophy and Economics*, 26.

¹⁰ This brief discussion will suffice for our purposes, but for a clear, concise explanation of expected utility theory, see Michael Resnik, *Choices: An Introduction to Decision Theory*, 45 – 118.

¹¹ Many anthropologists have argued that in traditional societies, economic decisions are governed by norms of reciprocity. See Karl Polanyi, 1944, *The Great Transformation*, Beacon Press, 1991 and Marshall Sahlins, 2004, *Stone Age Economics*, Chapter 1, “The Original Affluent Society,” 1-40. John Maynard Keynes famously criticized the homo economicus model for assuming that economic actors always calculate based on risk rather than uncertainty. See John Maynard Keynes, 2006, *A Treatise on Probability*, Cosimo, Inc. New York, NY, Chapter XXVI, “The Application of Probability to Conduct,” 313. A similar criticism is made in Israel Kirzner, 1973, *Competition and Entrepreneurship*, University of Chicago Press, Chicago, IL, Chapter 2, “The Entrepreneur,” 30-87. Kirzner argues that entrepreneurship cannot be modeled with homo economicus as entrepreneurs make decisions under uncertainty, i.e. without known probabilities. Finally, Herbert Simon’s classical work on bounded rationality argues that humans are rational but subject to certain cognitive

psychology, especially the criticisms advanced by psychologists, such as Daniel Kahneman, Richard Thaler and Amos Tversky.¹² There are a voluminous number of cases of poor human decision-making that we cannot review in the detail they deserve. Instead, we will cover some of the more prominent cases. To give one example, humans often commit the *conjunction* fallacy. The conjunction fallacy concerns a systematic error we make in probabilistic judgments. Kahneman and Tversky performed a famous experiment demonstrating the fallacy by giving subjects the following question¹³:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Please rank the following statements by their probability, using 1 for the most probable and 8 for the least probable.

- (a) Linda is a teacher in elementary school.
- (b) Linda works in a bookstore and takes Yoga classes.
- (c) Linda is active in the feminist movement.
- (d) Linda is a psychiatric social worker.
- (e) Linda is a member of the League of Women Voters.
- (f) Linda is a bank teller.
- (g) Linda is an insurance sales person.
- (h) Linda is a bank teller and is active in the feminist movement.

For ordinary undergraduate students with no statistics coursework, 89% ranked (h) as more likely than (f). The problem with this ranking, of course, is the probability of being an F cannot be lower than the probability of being an F and an H. The probability of F and H can only be as large as the probability of F and could be a great deal lower.

limitations. See Herbert Simon, 1984, *Models of Bounded Rationality, V. I, Economic Analysis and Public Policy*, The MIT Press, Cambridge, MA.

¹² Some of the classical studies can be found in Amos Tversky and Daniel Kahneman, "Rational Choice and the Framing of Decisions," in Hogarth, Robin and Reder, Melvin, eds. *Rational Choice: The Contrast Between Economics and Psychology*, Chicago: University of Chicago Press, 1987. An original version of the critique appears in "Judgment under Uncertainty: Heuristics and Biases." *Science*, September 1974, 185(4157), 1124-31. See also Richard Thaler, *Quasi Rational Economics*, New York: Russell Sage Foundation, 1991, and Thaler, Richard, "Mental Accounting Matters," *Journal of Behavioral Decision Making*, July 1999, 12(3), 183-206.

¹³ Kahneman, D. and Tversky, A. 1982, *Judgment Under Uncertainty: Heuristics and Biases*, New York, NY, Cambridge University Press, 92.

Other examples abound.¹⁴ One of the most important phenomena is known as *preference reversal*. In cases where the same decision is presented under distinct descriptions, preferences over multiple events often reverse. In the cognitive psychology literature, these are known as ‘framing effects’.¹⁵ Preference reversal is a strong challenge to standard neoclassical decision theory because, as we saw above, neoclassical decision theory requires that preferences be complete and intransitive. In many cases, experiments appear to show that our preferences are not complete or intransitive. The above case is an example of a framing problem, but when framing is applied to decision-making under risk, the same problems arise. Framing effects are not the only sources of criticism for neoclassical decision theory but they are core examples of how cognitive psychologists have challenged the foundations of many prominent, mainstream sub-fields of economics. Other comment effects derived from cognitive psychology include the *status-quo* bias, where subjects prefer what they currently have to alternatives that would make them better off.¹⁶

Experiments like these have generated the ‘heuristics and biases’ tradition in cognitive science. It holds that humans have evolved to process information according to crude cognitive heuristics that produce systematically biased behavior. The heuristics have low cognitive cost and promote survival and so they were selected for but they do not track proper modes of reasoning. These heuristics often get the answer right¹⁷, but they often get the answers wrong.¹⁸ What we can learn from the tradition is that under certain experimental conditions, individual human reasoning is poor and far removed from the rational capacities and practices of homo economicus. In the following section, we will

¹⁴ A 1966 study by Peter Wason has also generated a great deal of interest. Cf. Wason, P. C. "Reasoning." *New Horizons in Psychology*. Ed. Brian M. Foss and P. C. Dodwell. Harmondsworth: Penguin Books, 1966.

¹⁵ Amos Tversky and Daniel Kahneman, "Rational Choice and the Framing of Decisions," in Hogarth, Robin and Reder, Melvin, eds. *Rational Choice: The Contrast Between Economics and Psychology*, Chicago: University of Chicago Press, 1987. The original case of preference reversal can be found in Maurice, Allais, 1953, "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine," *Econometrica* 21, 503-546. The Allais paper has given rise to the preference reversal phenomenon known as the "Allais Paradox."

¹⁶ See D. Kahneman, J.L. Knetsch, and R.H. Thaler, "The endowment effect, loss aversion, and status quo bias," *Journal of Economic Perspectives*, 5, 1, 193-206, 1991 for more on the status-quo bias. We will discuss the endowment effect below.

¹⁷ Cf. Kahneman, Daniel, and Amos Tversky. "On the Psychology of Prediction." *Psychological Review* 80 (1973): 237-51.

¹⁸ Cf. Bower, B. "Rational Mind Design: Research Into the Ecology of Thought Treads on Contested Terrain." *Science News* 150 (1996): 24-25. Also see: Gould, Stephen Jay. *Bully for Brontosaurus: Reflections in Natural History*. Boston: W. W. Norton & Company, Incorporated, 1992.

review some proposals for reconciliation between the efficiency of markets and the irrationality of individual decision-makers.

Section II: Attempts at Reconciliation

Sometimes markets work as if economic agents are rational, but our best psychological evidence suggests that people aren't as rational as one might assume from observing the efficiency of markets. In this section, four attempts to reconciling this tension are assessed and rejected.

(II.i) One common attitude amongst economists is that the irrationality of individual actors doesn't matter so long as models have predictive power. This prevalent view amounts to a crude sort of positivism. This methodological disposition derives from Milton Friedman's early 1953 article "The Methodology of Positive Economics," where Friedman famously wrote:

A hypothesis is important if it "explains" much by little, that is, if it abstracts the common and crucial elements from the mass of complex and detailed circumstances surrounding the phenomena to be explained and permits valid predictions on the basis of them alone. To be important, therefore, a hypothesis must be descriptively false in its assumptions; it takes account of, and accounts for, none of the many other attendant circumstances, since its very success shows them to be irrelevant for the phenomena to be explained. ... Truly important and significant hypotheses will be found to have "assumptions" that are *wildly inaccurate* descriptive representations of reality, and, in general, *the more significant the theory, the more unrealistic the assumptions* ... [emphasis mine].¹⁹

Many economists are known to create models with highly implausible psychological assumptions. For instance, the permanent income hypothesis (developed by Milton Friedman) holds, to oversimplify, that patterns of consumer choices are not fixed by their current income. Instead, consumers consume according to their income expectations over their lifetime. The theory predicts that short-term changes in income will not affect

¹⁹ Milton Friedman, *Essays in Positive Economics*, Chicago: Chicago University Press, 1953, 3-43, 14-15.

consumer behavior significantly.²⁰ The theory predicts the empirical fact that consumption is relatively smooth over short periods, but while the hypothesis was once popular, other views have superseded it.²¹ Nonetheless, the hypothesis was once considered to have explanatory power despite its seemingly implausible key assumption: that ordinary humans make consumption decisions bearing in mind a coherent and determinate expectation of their incomes over their entire lives (or at least for a sufficient long period of time). For Friedman and those economists that follow him, however, the psychological implausibility of the view *did not matter*. In fact, for Friedman, a *virtue* of a good theory is that it has implausible assumptions!

The positivist strategy is inadequate. Rather than solving the problem, it tries to show that the problem is insignificant to economic theory. Three replies are in order. First, even if Friedman is right, it is important for psychologists and economists to ask how individual psychology and economic reality mesh. It is a fascinating question in its own right. Second, one thing that we often want science to do is to illuminate *the truth* about its subject matter. The permanent income hypothesis is probably false. But what true theory then explains how it fits with data on consumer choices? This is a question appropriate for economists to ask. But perhaps the biggest problem is as follows. Philosophers of science and many scientists knew long before Quine made it commonplace that a vast number of theories are compatible with most data.²² And this is true across the scientists. In economic life, there are many more variables than there are in many models in the natural sciences, so the indeterminacy of theories is even worse.²³ One advantage of trying to develop a psychologically plausible economic model is the hope of developing a theory that not only predicts the data but gives a psychologically plausible explanation of it. This can limit the number of theories under consideration for explaining particular empirical phenomena and make the process of scientific progress more effective. For these reasons, positivism seems like a poor way to solve our problem.

²⁰ For a helpful discussion of the permanent income hypothesis, see:

<http://ingrimayne.com/econ/FiscalDead/PermIncome.html>.

²¹ A. S. Deaton, "Life Cycle Models of Consumption: Is the Evidence Consistent with the Theory?" In *Advances in Econometrics*, ed. Amsterdam: North Holland, 1987. H. Molana, "Consumption Function" in *The New Palgrave Dictionary of Money and Finance*, eds. P. Newman, M. Milgate, and J. Eatwell, London: Macmillan, 1992.

²² For more on the underdetermination of theory by data, see W.V. Quine, "On the Reasons for Indeterminacy of Translation," *Journal of Philosophy*, 67 (6), 178-183.

²³ Nobel Laureate economist, F. A. Hayek, discusses his theory of complex phenomena in his Nobel Lecture, see http://nobelprize.org/nobel_prizes/economics/laureates/1974/hayek-lecture.html.

(II.ii) Another proposed solution to the problem of conflict between individual irrationality and market efficiency is to claim that despite the economic evidence to the contrary, psychological evidence shows market failure to be extensive and serious. Many behavioral economists and cognitive scientists pursue this line, arguing that economic evidence for market successes is ambiguous and that economists' theories are highly abstract and psychologically implausible. As a result, the appearance of market efficiency is an illusion. This is not to say that behavioral economists believe that markets always fail, but that they appear to work much more often than they do and that close psychological examination of individual choice in economic environments will reveal high degrees of irrationality at work.

A recent popular book in behavioral economics makes such an argument. In *Predictably Irrational*, Dan Ariely argues that humans are often fooled into thinking that we make rational decisions, when we don't. Ariely asks his reader: "If I asked you for the last two digits of your social security number (mine are 79), then asked you whether you would pay this number in dollars (for me this would be \$79) for a particular bottle of [wine], would the mere suggestion of that number influence how much you would be willing to spend online?" It turns out that the answer is yes, even for a group of MIT MBA students.²⁴ Humans irrationally anchor themselves to initially suggested prices and are resistant to move off of them. This is known as the anchoring effect. Ariely draws a moral from this. Markets assume that prices are determined by supply and demand, but this depends on "the assumption that the two forces are independent and that together they produce the market price."²⁵ Ariely, et al.'s experiments are evidence to the contrary: "what consumers are willing to pay can easily be manipulated."²⁶ Ariely goes on to suggest that, for instance, increasing the price on gasoline won't permanently reduce demand because people will adjust their anchors over time; as a result, free market exchanges may not always be mutually beneficial due to irrational anchoring. These are both significant challenges to orthodox views in economics. Ariely's argument is radical. Since we can be mistaken about what will bring us utility, then why should economists and policy makers assume that there is a presumption that individuals will act in their best interests? Ariely suggests that policy makers should "take this into account" but if he is right the case for the free-market is destroyed. The classic defense of markets is that they utilize the expertise of individuals on their own preferences

²⁴ Dan Ariely, *Predictably Irrational*, 26-28.

²⁵ Ibid, 45.

²⁶ Ibid.

and local information in order to increase economic efficiency through the extension of the division of labor and the accumulation of capital.²⁷ If individuals do not know what will satisfy them, and routinely make rational errors in deciding what to buy and for how much, then the classic defense is threatened.

An odd feature of Ariely's book is that he neglects to mention a fifty-year old field of economics that has already struggled with some of these questions: experimental economics. Founded by Nobel Laureate Vernon Smith, this branch of economics attempts to verify the details of economic models by running economic experiments with test subjects. This contrasts with the work of behavioral economists like Ariely, who primarily study individual economic behavior. Smith and his colleagues have conducted hundreds of experiments and have replicated the efficiency of competitive markets in many cases. While Smith notes that his experiments do not always confirm standard expected utility theory, "experimental tests of market theories, which explicitly assume expected utility (or value) maximization, have *not* falsified many of these theories."²⁸ And while many initial tests show marked irrationalities, over time behavior comes to approximate economic models. Smith has long thought "that markets may induce greater 'rationality' in behavior because they force or promote a response to, or discovery of, opportunity cost conditions, that need not be readily forthcoming when agents merely think about the choices they make."²⁹ He has argued for decades that the Kahneman-Tversky framing effects "do relatively low level damage to [the expected utility hypothesis]" because when at least some cost is a stake, economic actors can learn to avoid these effects.³⁰ This is not to say that economic actors become expert statistical reasoners; rather, they learn to avoid options that lead to net cost. To illustrate, imagine how an experimental economist could address the anchoring effects that Ariely describes; certainly, the experimental economist will admit, one can produce anchoring effects in individual persons, even in groups initially; the question is whether these effects will persist over time as individuals become more familiar with how to avoid net costs to themselves. The problem with Ariely's discussion of anchoring effects on the experimental economist's view would be (i) that it does not include a comparison class of objects that can be chosen in the context of subjection to a budget constraint, and (ii) it does not tell the researcher

²⁷ One of the classic articulations of this view in the 20th century can be found in F.A. Hayek, "The Use of Knowledge in Society," *American Economic Review*, XXXV, 4, 519-530. A classic 19th century defense can be found in J.S. Mill's *Principles of Political Economy* and a classic 18th century exposition is Adam Smith's *The Wealth of Nations*.

²⁸ Vernon Smith, "Experimental Economics: Reply," 267.

²⁹ Ibid.

³⁰ Ibid, 268-269.

whether in repeat conditions the learned rationality of subjects will ‘wash-out’ or reduce anchoring effects. In fact, experimental economists have argued just this with respect to the endowment effect, which is the observed behavior that individuals prefer what they possess to what they do not possess even when a clearly superior choice presents itself. Experimental economists have argued that the endowment effect is reduced in repeat interactions.³¹

Smith argues that economists have learned from experiments that institutions matter.³² The formal and informal rules that govern institutional environments are required to enable economic learning. On Smith’s view, economic rationality is an *achievement* not a given that only occurs under the right sets of rules.³³ In many cases, market optimization is an *unconscious* process: Smith reports that “hundreds of double auction experiments ... would spotlight the crucial importance of not ruling out the rationality of unconscious decision in rule-governed repeat interaction settings.”³⁴ What is even more odd is that market actors often become more irrational when given more information. A bargainer with better information, “... knowing that the other player knew only his own payoff, is more forgiving when his opponent makes large demands.”³⁵ Surprisingly, “This concessionary posture works to the disadvantage of the completely informed player.”³⁶ Finally, and of interest to game theorists, experimental economists have learned that it is crucial for players to play dominated strategies.³⁷ Players should cooperate even when it will lead to losses in

³¹ Some may be surprised, but the empirical literature on this matter is quite large; depending on how experiments are set-up different experimenters have found that endowment effects are reduced (cites) or even eliminated (cites) by repetition; other experimenters find the endowment effects are not reduced by repetition.

³² This may seem like a trivial insight, but many periods in the history of economic thought have been characterized as engaging in the development of abstract theorizing at the expense of studying institutional context. I partly document the criticism as it was developed in the mid-19th century in “Production, Distribution and J.S. Mill,” *Utilitas*, forthcoming.

³³ Vernon Smith, “Economics in the Laboratory,” *Journal of Economic Perspectives*, 8, 1, Winter 1994, 113-131, 116.

³⁴ *Ibid*, 118.

³⁵ *Ibid*, 119. This is often called the “curse of knowledge”, a term coined in Colin Camerer, George Loewenstein, and Martin Weber, “The Curse of Knowledge in Economic Settings: An Experimental Analysis,” *Journal of Political Economy*, October 1989, 97:5, 1232-54.

³⁶ *Ibid*.

³⁷ A dominated strategy is one that can be systematically beaten by a dominant strategy. A dominant strategy is defined as follows: “A strategy is dominant if, regardless of what any other players do, the strategy earns a player a larger payoff than any other. Hence, a strategy is dominant if it is always better than any other strategy, for any profile of other players’ actions. Depending on whether “better” is defined with weak or strict inequalities, the strategy is termed strictly dominant or weakly dominant. If one strategy is dominant, then all others are dominated. For example, in the prisoner’s dilemma, each player has a dominant strategy.” See gametheory.net for further definitions: Shor, Mikhael, “Dominant Strategy,” Dictionary of Game Theory Terms,

the short-term is because the process of cooperation helps to form common expectations. Common knowledge, Smith has found, is not sufficient to produce efficient outcomes. Common expectations must be established before efficiency can be fully achieved. These common expectations are formed through repeated cooperation. Thus Smith says, “Dominated Strategies Are for Playing, Not Eliminating.”³⁸

Experimental economics generates other important insights into economic behavior. But the above discussion should prove sufficient to show that the Ariely critique fails. Ariely’s experiments no doubt show that individuals can be irrational. Further, other behavioral economists have challenged experimental economists’ account of market learning, arguing that some irrationality is not reduced through learning.³⁹ Some recent work even purports to show that some irrationality even increases through repetition.⁴⁰ But it is widely established that repetition reduces irrational economics behavior in a wide range of cases; experimental economists have shown that at least in very many cases markets can regularly achieve efficient outcomes despite individual irrationality. Market efficiency is possible if individuals are irrational by themselves, *so long as individuals are rational together and interact over time*.⁴¹

Most behavioral economists make more modest claims than Ariely, never denying that market repetition sometimes reduces market irrationalities. Thus the appearance of at least widespread market success need not be denied. The problem with the ‘extensive market failure’ solution to the problem is similar to the problem with positivism. Instead of struggling to understand an important tension, it simply discounts one side of the issue, either by arguing that the other side is irrelevant or by ignoring it. If the critique fails, then

Game Theory.net, < <http://www.gametheory.net/dictionary/DominantStrategy.html> > Web accessed: June 15, 2009.

³⁸ Smith, “Economics in the Laboratory,” 122.

³⁹ A nice behavioral economist critique of experimental economics that raises some of this work is George Loewenstein, “Experimental Economics from the Vantage-Point of Behavioral Economics,” *The Economic Journal*, 109 (February), F25-F34. Some irrationalities do not go away after repetition in social settings or after extensive learning, but irrationalities are often substantially reduced. Other work suggests that repeated experience significantly reduces the endowment effect, John A. List, “Does Market Experience Eliminate Market Anomalies?,” *The Quarterly Journal of Economics*, February 2003, 41-71.

⁴⁰ Cf. Jacinto Braga, Steven J. Humphrey and Chris Starmer, “Market Experience Eliminates Some anomalies – And Creates New Ones,” *CeDex Discussion Paper No. 2006-19*, October 2006, forthcoming, *European Economic Review*. However, another study contradicts this, arguing that preference reversal in the Allais paradox can be overcome through repetition, Gijs van de Kuilen and Peter P. Wakker, “Learning in the Allais Paradox,” working paper, August, 2006.

⁴¹ This point is emphasized in David K. Levine, “Is Behavioral Economics Doomed? The Ordinary versus the Extraordinary,” Max Weber Lecture, June 8th, 2009.

the presumption of market success in at least many cases is retained. The apparent tension between individual irrationalities and widespread market success cannot be denied; it must be resolved.

(II.iii) Perhaps the most sophisticated attempts to reconcile irrational psychologies with market rationality lies in the field of alternative expected utility theories. These theories depart, in one way or another, from the standard assumptions of expected utility theory outlined in I.i and I.iii. The theories vary substantially, so I cannot cover them in detail here. I will instead describe two general classes of alternative expected utility theory, and attempt to show that both of them, while accounting for the psychological evidence within a decision theoretic framework, still fail to generate a satisfying reconciliation between the effectiveness of markets and individual irrationality.

The two types of alternative expected utility theories [henceforth, EUT/AEUT] are “conventional” EUTs and “nonconventional” EUTs.⁴² Conventional EUTs try to preserve the general form of expected utility analysis, that is, indifference curve analysis. But they do so by relaxing some of the standard assumptions of expected utility theory. They assume, however, that agents behave as though they are optimizing a preference function. Nonconventional EUTs, in contrast, reject this view. They are often called ‘procedural’ theories, which assume merely that economic actors use decision heuristics or rules rather than maximizing a preference function.⁴³ Conventional EUTs come in many varieties but we will focus on one – weighted expected utility theory.⁴⁴ Weighted EUT treats standard EUT as a specific case of a general theory. EUT holds when all consequences are equally weighted.

⁴² These terms derive from Chris Starmer, “Developments in Non-Expected Utility Theory,” *Journal of Economic Literature*, 38, 332-82, 2000.

⁴³ Incidentally, I do not think this distinction is well formed. One can both use heuristics and be understood as optimizing a preference function. The preference function is simply a model for representing choice. It does not need to have explicit, conscious psychological underpinnings. See Gerald Gaus, *Politics, Philosophy and Economics*, introduction.

⁴⁴ Among them are generalized expected utility analysis, see Mark Machina, “‘Expected Utility’ Analysis Without the Independence Axiom,” *Econometrica* 50, 277-323, Disappointment Aversion, see D. E. Bell, “Disappointment in Decision Making under Uncertainty,” *Operations Research*, 30, 961-981, 1985 and G. Loomes and R. Sugden, “Disappointment and Dynamic Consistency in Choice under Uncertainty,” *Review of Economic Studies*, 53, 271-282, 1986, Rank-Dependent EUT, see J. Quiggin, “A Theory of Anticipated Utility,” *Journal of Economic Behavior and Organization*, 3, 323-343, 1982, and the theory we will focus on, Weighted Utility Theory, S. H. Chew and K. R. MacCrimmon, “Alpha-Nu Choice Theory: A Generalization of Expected Utility Theory,” Working Paper No. 669, University of British Columbia, Faculty of Commerce and Business Administration, 1979.

But decision weights can differ and advocates of weighted EUT argue that cases that diverge from standard EUT are best explained by distinguishing between differently weighted consequences. For instance, in preference reversal cases, sometimes two bets are presented, one with a low probability of a high payout and the other with a moderate probability or a moderate payout. The expected utility of both bets are identical, yet bettors consistently prefer moderate probability. Weighted expected utility theorists argue that innate decision-making heuristics produce a tendency to assign more weight to the less risky bet.

The most prominent nonconventional EUT is prospect theory.⁴⁵ Prospect theory argues that decision-making under risk has two separate stages, an 'editing' phase and an 'evaluation' phase. In the editing phase, the realized options are structured and formed to make choice easy. During this cognitive process, heuristics are used to assign perceived outcomes as positive or negative judged against some baseline. Common factors are left out, probabilities of the same outcomes are combined, and dominated options are removed if they are detected. However, many mistakes occur during this process. Next comes the evaluation phase. The evaluation phase contains both a value function and a weighting function, as weighted EUT suggests. The decision weighting function usually involves assigning irrationally high risk to small probability events. Some economists have worried about the relevance of the editing phase, so Kahneman and Tversky developed cumulative prospect theory, which avoids relying too heavily on the editing phase.⁴⁶ In cumulative prospect theory, only cumulative probabilities are transformed rather than individual probabilities themselves. Consequently, expected utility theory ranges over the probabilities of the outcomes as a set. This produces the irrationally high weights we assign to low probability events that result from the editing phase.

Both conventional and nonconventional EUTs purport to explain why markets work in some cases and not in others. In some cases, their explanations are illuminating. Both theories can account for assigning irrational probabilities to events of low probability, and both can accommodate preference reversal phenomena, among others. To this extent, alternative expected utility theories help to reconcile how markets work with the fact that humans are sometimes irrational economic actors. In those cases where preference reversal, framing, etc. are significant, markets will not work as well as they could; but in cases where

⁴⁵ The original theory premiered in Kahneman, Daniel, and Amos Tversky (1979) "Prospect Theory: An Analysis of Decision under Risk", *Econometrica*, XLVII (1979), 263-291.

⁴⁶ Tversky, Amos; Daniel Kahneman (1992). "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and Uncertainty* 5: 297-323.

irrationalities do not show up, markets may be efficient. However, there is one experimental phenomenon that these theories do not directly explain. This is the phenomenon of anomaly reduction through repetition. Often deviations from rational decisions can be reduced significantly through repeating the experiment. In other words, economic actors adapt in the lab. We saw in II.ii that repetition allows actors to converge on value maximizing behavior; over time weights and preferences change so as to allow markets to approximate the behavior that standard expected utility theory assumes they have. The feedback from repeated social interactions helps isolated actors or actors in single-shot bets to act rationally according to their self-interest.⁴⁷ So while alternative expected utility theories may be part of the story, social interaction, adaptation and common expectations are required to give a fully satisfying resolution to our problem.

(II.iv) The large literatures in behavioral and experimental economics converge on a few points of agreement. The one most significant for our purposes is that rational outcomes can be achieved through *limiting options*. Experimental economists argue that market discipline limits choices in market contexts that remain open in single-shot and individual decision-making. Behavioral economists argue that the artificial conditions of the experimental economist's lab artificially limit options, so as to make rational behavior a trivial accomplishment. But both agree that limiting options is one way to generate seemingly rational behavior.⁴⁸

Limiting options can be given a more concrete form, however, when we connect it to what philosophers sometimes call 'situated' or 'embedded' cognition.⁴⁹ The view, in short, is that human cognition is not merely 'in the head' or entirely contained within the brain.

⁴⁷ See ft. 26 for relevant citations. Note that I claim that weights *and* preferences change. Experimental economists tend to explain adaptation as individual's adjusting weights as they discover their preferences and how to satisfy them. But behavioral economists sometimes argue that individuals are simply constructing new preferences in response to losing out on potential gains. This is sometimes called the discovered-constructed preference debate. A recent paper on the constructed preference side is Graham Loomes, Chris Starmer and Robert Sugden, "Do Anomalies Disappear in Repeated Markets?," *The Economic Journal*, 113 (March), C153-C166, 2003. For our purposes, the victor in this dispute is not relevant. In any event, over time individuals come to have their preferences increasingly satisfied whether they are constructed during the process of interaction or not. In some ways, as a result, the constructed-discovered preference distinction seems to be a bit of a bugbear depending on how 'constructed' or 'discovered' is defined. 'Discovered' might be understood dispositionally with context-sensitive activation for the disposition, whereas 'constructed' might be understood as the preference coming to drive action directly. Under these definitions, discovered and constructed preferences would not be distinguishable.

⁴⁸ See ft. 36.

⁴⁹ One early statement of this view can be found in Andy Clark, *Being There ...*

Instead, it holds that human cognition arises out of a complex interplay between brain, body and environment. Philosopher Andy Clark has argued that what makes human cognition unique is not that humans are particularly good reasoners in isolation, but instead possess “amazing capacities to create and maintain a variety of special external structures.”⁵⁰ He continues: “these external structures function so as to complement our individual cognition profiles and to diffuse human reason across wider and wider social and physical networks whose collective computations exhibit their own special dynamics and properties.”⁵¹ Cognition, in other words, is extended into the environment, making complex forms of cognition easier for the human brain. Clark argues that advanced cognition requires the ability to “dissipate” some forms of reasoning. We have to reduce the stress of cognition on our brains by situating ourselves in complex social structures. In this way, we are not much different than robots, save in our ability to structure our environment to our advantage.

This movement in cognitive science has extended into philosophical interpretations of decision theory. Debra Satz and John Ferejohn have defended a position in the philosophy of social science known as “moderate externalism” about social scientific explanations (including economics).⁵² The internalist holds that social scientific events can only be explained by intentional states interior to the individual human. Moderate externalism does not deny that individual intentional states are relevant to social scientific explanations but their connection to social scientific explanations is ‘remote’. Moderate externalism “does not explain behavior in terms of these mental entities; it merely shows that behavior can be interpreted as consistent with them.”⁵³ It explains social scientific events in terms of the “structure in which [individual action] is embedded.”⁵⁴ On this view, firms are relations of individuals structured such that they [firms] can be understood *as if* they were intentionally maximizing profits. The basic unit of analysis for the moderate externalist is not the individual but a set of individuals and institutional conditions.

The New Institutional Economics shares much with the embedded cognition paradigm. Standard neoclassical economics relies on what new institutional economists Arthur Denzau and Douglass North call “substantive rationality” or rationality that is

⁵⁰ Clark, 179.

⁵¹ Ibid.

⁵² Debra Satz and John Ferejohn, “Rational Choice and Social Theory,” *Journal of Philosophy*, 91, 2, 71-87, 1994.

⁵³ Ibid, 77.

⁵⁴ Ibid, 78.

contained entirely “in the head.”⁵⁵ But Denzau and North argue that we cannot make sense of economic and political history if individuals simply know their self-interest and act in accord with it. Denzau and North argue that institutions reduce transactions costs, thereby making market efficiency possible. Their focus on transactions costs follows economist Ronald Coase. Coase was one of the first economists to focus on the role transactions costs played in economic models. Standard neoclassical modeling assumes zero transactions costs, but clearly this assumption is not realistic. Sometimes transactions costs are monumental or at least have a substantial impact on market functioning. One problem that Coase saw with ignoring transactions costs is that it became impossible to see why firms would arise on the free-market. Firms are simply centrally directed exchanges – command economies on an extremely small scale. But if markets are efficient, why should firms exist at all? Coase’s answer: to reduce transactions costs.⁵⁶ Discovering the information necessary to make exchanges, coordinating exchanges and producing transaction assurance are costly; firms arise to reduce those costs. However, firms are limited in their expansion, as the inefficiencies of central direction eventually outweigh the reduction in transactions costs. Denzau and North apply this framework to economic history, arguing that firms are not the only institutions that arise to reduce transactions costs. Ideologies and cultural institutions arise for the same reason. Just like the embedded cognition tradition, they see economic and social institutions as mechanisms that ‘offload’ cognition into the environment, or in their terms, reduce transactions costs.⁵⁷

How do these two traditions solve the conflict between psychology and economics? They argue that individual psychology *just isn’t that important* when explaining how markets work. Instead, most of the efficiency of markets is due to the institutional structure in which economic actors are embedded. Humans create these institutions, to be sure. But once transactions costs are sufficiently reduced and economic cognition is sufficiently embedded, our options for choice are limited. Cognition then becomes easy enough for ordinary humans to make reliably rational decisions. One nice feature of these accounts is that they can explain why markets work in some environments and not others. North’s view was influenced by the implementation of markets in the former Soviet Union, which proved

⁵⁵ Arthur T. Denzau and Douglass C. North, “Shared Mental Models: Ideologies and Institutions,” working paper, 1995.

⁵⁶ The classic article here is Ronald Coase, “The Nature of the Firm,” *Economica*, New Series, 4, 16, 386-405, Nov., 1937.

⁵⁷ This view may smack of economic determinism, but one need not be an economic determinist to appreciate this point.

less effective than the implementation of markets in Western countries. North's theory is that markets failed in former Soviet Union because there were no cultural and moral institutions that reduced transactions costs sufficiently to make markets work. Individuals could not trust one another and did not know how to run businesses; as a result, they faced enormous transactions costs.⁵⁸

North and Clark cite experimental evidence that individual rationality isn't that important in making markets efficient. Their primary evidence comes from market experiments where human traders are replaced with "zero-intelligence" programs. These programs submit random bids and offers. Dhananjay K. Gode and Shyam Sunder find that "imposing a budget constraint (i.e., not permitting traders to sell below their costs or buy above their values) is sufficient to raise the ... efficiency of these auctions close to 100 percent."⁵⁹ They go on to claim that the efficiency of markets appears to be due to "its structure independent of traders' motivation, intelligence, or learning. Adam Smith's invisible hand may be more powerful than some have thought; it can generate aggregate rationality not only from individual rationality *but also from individual irrationality*."⁶⁰ The argument then is that it is the shape of institutions that produce rational action. Like Satz and Ferejohn, Gode and Sunder explain the rationality of markets partly in terms of institutional structure. Like North and Denzau, they find that institutions are required to make markets work. And like Clark, they argue that rational cognition occurs when cognition is embedded in the world. Vernon Smith has argued in favor of what he calls the Hayek Hypothesis, the thesis "that strict privacy together with the trading rules of a market institution are sufficient to produce competitive market outcomes at or near 100% efficiency."⁶¹ Smith finds that markets can often produce efficient outcomes with far less information than standard neoclassical analysis requires. But Gode and Sunder go farther: Economic actors do not need much information *or rationality* in order to produce efficient outcomes under the right rules.

⁵⁸ For an elaboration of this story, see Douglass C. North, *Understanding the Process of Economic Change*, Princeton, Princeton University Press, 2005, Chapter 11, The Rise and Fall of the Soviet Union, 146-170.

⁵⁹ Dhananjay K. Gode and Shyam Sunder, "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *The Journal of Political Economy*, 101, 1, 119-137, 1993, 119. Their full explanation of the allocational efficiency of markets can be found in Dhananjay Gode and Shyam Sunder, "What Makes Markets Allocationally Efficient?" *The Quarterly Journal of Economics*, May 1997, 603 – 630.

⁶⁰ Ibid.

⁶¹ Smith's testing of the hypothesis comes out positive. See Vernon Smith, "Markets as Economizers of Information," *Economic Inquiry*, Vol. 20, April 1982, 165-179.

The embedded cognition/transactions cost approach is compelling. It explains how individuals can be irrational and ignorant and yet produce efficient outcomes under certain economic conditions. But it has two drawbacks. First, it has the counterintuitive implication that individual rationality has practically no causal efficacy.⁶² None of these experimenters deny that rationality has *any* causal power in bringing about efficient outcomes. Instead, they argue that economic rationality is so embedded in the world and subject to external constraints that the common sense conception of individual rationality is largely eviscerated. Some will welcome this counterintuitive consequence, but it has theoretical costs. The causal inefficacy of our rational faculties seems to undermine our common sense understanding of the world. For instance, we have the experience of making rational decisions in many economic environments. Buying groceries at the grocery store, selecting a house and the like are all activities that some participate in irrationally but many work through with great care and thought. Surely these calculations make a substantive causal contribution to our decisions.

A stronger worry is that the extended cognition model also makes it hard to understand how economic actors in the lab substantially improve efficiency in repeat interactions. Economic actors learn to interact more effectively *under the same set of rules*. Perhaps zero-intelligence traders are nearly as efficient as humans are in one-shot auctions where the rules do not change. But they do not update their behaviors and it is this feedback process that may be crucial to producing efficient markets. A theory that can explain both how irrational humans produce efficient markets and how they can improve over time is a better theory. While Clark and North point towards historical stories about how humans alter their institutions, these stories remain vague; the concrete local-level mechanisms involved are not well understood.

Section III: Subjective Costs of Transacting as a Solution

(III.i) In this section, I defend an approach to our problem I call the *subjective cost of transacting* (SCT) theory.⁶³ The theory is not meant primarily as a *prescriptive* theory or a

⁶² The data is also compatible with the view that rationality has causal efficacy and that choices are causally overdetermined by the environment. This hypothesis has not been pursued.

⁶³ The first mention of this theory is in Jacob Marschak, "Economics of Inquiring, Communicating, Deciding," *American Economic Review Proceedings*, May 1968, 58, 1-18. Vernon Smith develops this view in Vernon Smith, "Microeconomic Systems as an Experimental Science," *American Economic Review*, December 1982, 72, 923-55, esp. 934. Smith's view is also sometimes referred to as 'the labor theory of cognition' but to adopt this

theory about what humans should do; rather I advance the SCT theory as an *explanation* of the compatibility of irrational actors and rational markets. In this section, my aim is to show that the SCT theory can explain how repetition and institutions help to produce rational market behavior in otherwise often irrational actors; it can show how rational market behavior is compatible with extensive heuristic use in human cognition. If humans generally are more rational when subjective transactions costs are minimized, then we will expect to see irrational behavior when these costs are high and how market conditions reduce subjective transactions costs sufficiently to generate rational market behavior. The SCT theory can also explain how rationality has causal efficacy despite often appearing not to.

The subjective costs of transacting are the costs of “thinking, calculating, deciding and acting.”⁶⁴ I understand subjective transactions costs as those costs of transacting that are *cognitive*, or internal to an agent’s cognitive systems. Further, these are costs that are – to some extent – under the agent’s control. To illustrate, imagine that you go to the grocery store to buy a box of Wheat Thins. There are boxes of two sizes. The smaller box is less expensive than the larger box, but contains fewer Wheat Thins. Typically, the large box is less expensive per Wheat Thin, but today the small box is on sale. You now face a calculation: you must estimate the number of Wheat Thins in each box and then divide the price by the number of Wheat Thins to figure out the price per Wheat Thin.⁶⁵ This calculation can be tedious. For many people, division with decimals and multiple digits is a chore and few people carry around calculators. Even if they have one on their mobile phone, they would still have to slow down, focus and not move onto the next item. All of these frustrations and calculations are subjective costs of transacting. They are (a) cognitive – internal to our cognitive systems, and (b) under our control to some extent. We can choose to forgo the costs but we can also exert the attentional resources and do the calculation.

External structure in the economic environment can reduce the subjective costs of transacting. For instance, common expectations of behavior can be implicit and learned at an early age. These expectations reduce the stress involved in trusting others by assuring economic actors that if they cooperate, the other actor will reciprocate. Clear institutional rules like contract laws and informal moral norms can reduce not only the stress of worrying

label would imply that it is only *labor* costs that matter to cognition; capital costs might matter as well. See below.

⁶⁴ Smith, 268.

⁶⁵ Suppose that your rate of eating Wheat Thins is an exogenous variable (you won’t eat them faster if you have more) and that you will continue to eat Wheat Thins for the foreseeable future, so the sheer quantity of Wheat Thins is irrelevant.

about whether to trust others, but the costs of deciding what to do. Many of our economic decisions become unconscious habit as we practice them over time. Humans become habituated to these practices because formal and informal rules remain stable. If humans are rational when subjective transactions costs are low but often become irrational when they are high, then we would expect for rational market behavior to increase within institutions that reduce transactions costs. Thus, the SCT theory is consistent with the insights of the embedded cognition/new institutional approach.

Repetition can also reduce subjective transactions costs so long as the payoffs and rules remain stable. Learning simply *is* a form of reducing the subjective transactions costs of acquiring, storing and accessing information. We learn how to make good decisions through repeating the decision, observing the outcome and adjusting our behavior accordingly. Consider many of the experiments in experimental economics; as trials are repeated, individuals approach efficient market outcomes. If repetition lowers subjective transactions costs, then it is no surprise that repetition increases rational market behavior. In fact, Vernon Smith defends the subjective transactions costs theory to explain how repetition produces market efficiency.⁶⁶

The SCT approach also helps to explain how market rationality is consistent with the widespread use of heuristics observed by cognitive psychologists. Thinking is costly even when we have all the relevant information. But gathering information is costly as well. In many cases, then, it would save time and cognitive exertion to on a less cognitively expensive but less reliable cognitive heuristic. On this view, the ‘irrationalities’ discovered by behavioral economists are rational responses to high subjective costs of transacting. Relying on heuristics makes sense when higher-level cognition is cognitively expensive. Additionally, there is good neurological evidence to suggest that calculation in more cognitive, less emotional parts of the brain are slower than other processes.⁶⁷ Richard Samuels and Stephen Stich defend a *dual processing* theory of reasoning that has been gaining prominence over the last decade.⁶⁸ These theories argue that we make decisions with two distinct systems. Samuels and Stich argue that one system is “fast, holistic, automatic, largely unconscious, and requires relatively little cognitive capacity.”⁶⁹ But the other system is “relatively slow, rule based, more readily controlled and requires significantly more

⁶⁶ See ft. 63.

⁶⁷ See Richard Samuels and Stephen Stich, “Rationality and Psychology”.

⁶⁸ For some of the recent literature, see Evans, J. and Over, D., “Dual Processes in Thinking and Reasoning.” (forthcoming), Sloman, P. “The empirical case for two systems of reasoning.”

⁶⁹ Samuels and Stich, 24.

cognitive capacity.”⁷⁰ Many psychologists who adhere to this theory think that the former system is evolutionarily *older* than the latter system and that the latter system can be deeply affected by cultural and institutional factors. Some psychologists think that differences in aptitude will affect the degree to which the latter, newer system functions adequately, yet they hold that the former, older system is ubiquitous in its error-ladenness and hard to alter. Some recent work in moral decision-making has been built off of dual processing models in the non-moral decision-making literature. Joshua Greene, for instance, has tried to correlate deontological judgments with the older, more instinctive system and utilitarian judgments with the latter, more ‘rational’ system. And he has marshaled some neurological evidence which shows that different parts of the brain simultaneously react to the same moral problems.⁷¹ Greene has shown that in standard moral dilemmas, the parts of the brain that generate deontological judgments exhibit some of the same putatively erroneous behavior as those in the non-moral decision-making tasks we discussed above. On these views, agents may often have a *choice* as to which system to rely on. We could use our “fast, holistic, automatic” system that requires “little cognitive capacity.” Or we could spend the resources to use our “relatively slow, rule-based, more readily controlled” system of reasoning that has higher subjective costs of transacting. But when the information required to do the relevant calculation is unavailable or too costly to acquire or the calculations required are too cognitively expensive, it is rational to rely on systems of reasoning that are less accurate given their cost-benefit ratio. I do not want to suggest that the SCT theory commits us to the view that we always choose which system to rely on. Most of our decisions are governed by the faster, less conscious system. But in economic interactions we have a choice as to whether the faster, less conscious system has the final say in how we act. We can choose to allow our slower, more rule-based system of reasoning to override the outputs of the faster, quicker system. In many cases, we cannot prevent the older system from operating, but we can often prevent its outputs from determining our actions.

In the same way, when a decision is relatively cognitively inexpensive, i.e., when the subjective transactions costs of calculating are low, we can expect for heuristic use to diminish or for its application to the environment to be modified such that it is consistent with market efficiency. In fact, experimental economists have discovered just this phenomenon in the laboratory; as we have already discussed, repetition reduces the endowment effect sometimes entirely and almost always to a significant degree. Many of the

⁷⁰ Ibid.

⁷¹ Greene, “The Secret Joke of Kant’s Soul.”

experiments where cognitive psychologists have found individual irrationalities are ones that rely on probability calculations that for many people would require extra cognition, such as with framing effects. Individuals may not realize ahead of time that they would do better to calculate rather than let their natural heuristic cognition dictate their decisions. However, the SCT does run aground when faced with numerous studies which find that statistics courses often have no significant reduction in probability miscalculations and actually seem to increase some probability biases.⁷² One would expect for knowledge of statistics to reduce subjective transactions costs, and therefore increase the level of rational decision-making. An experimental economist could counter that even those with statistics coursework have no incentive to make the correct judgment – even for those with statistics courses, utilizing that information may be deemed too costly for a psychological experiment. A variety of experiments have found that significant financial payoffs substantially improve rational risk-taking, although errors are not eliminated.⁷³ Smith and Walker argue that “increased financial rewards [may] shift the central tendency of the data toward the predictions of rational models ... [and] in virtually all cases rewards reduce the variance of the data around the predicted outcome.”⁷⁴ However, some have argued that the ‘labor cost’ of cognition is not the only cost of cognition that matters. Camerer and Hogarth emphasize the ‘capital’ aspect, which include cognitive abilities that, while largely invariant in the short run, can be improved through repetition.⁷⁵ In any event, the SCT theory does not specify which costs are relevant, only that cognition costs, when reduced, lead to rational behavior and when increased, lead to less rational, heuristic-based, less effective behavior.

⁷² Many of these studies are documented in Kahneman, Daniel, and Amos Tversky (eds.) 2000. *Choice, Values, and Frames*. Also see Levy, Jack S. “Daniel Kahneman: Judgment, Decision, and Rationality.” *PS: Political Science and Politics* 35: 271-273, 273. For a very recent study which comes to the same conclusion, see Kinga Morsanyi, Caterina Primi, Francesca Chiesi, Simon Handley, “The effects and side-effects of statistics education: Psychology students’ (mis-)conceptions of probability,” *Contemporary Educational Psychology*, 34, 3, 210-220 (2009).

⁷³ Smith and Walker, “Monetary Rewards and Decision Cost in Experimental Economics,” *Economic Inquiry*, 31, 245-261. Although, these effects do not increase monotonically, see Gneezy and Rustichini, “Pay Enough or Don’t Pay at All,” *Quarterly Journal of Economics*, 115, 791-811. Experimenters found that at low levels of compensation, errors *increase*. The authors postulate that this might result from subjects feeling insulted.

⁷⁴ *Ibid*, 245.

⁷⁵ C. Camerer and R. Hogarth, “The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework,” *Journal of Risk and Uncertainty*, 19, 7-42. For an attempt to generate a ratio between these two costs, see Ondrej Rydval and Andreas Ortmann, “How Financial Incentives and Cognitive Abilities Affect Task Performance in Laboratory Settings: An Illustration,” *Economics Letters*, 85, 3, 315-320. The authors find around a 2:1 ratio between ‘capital’ and ‘labor’. I am not sure whether the capital/labor contrast is the most illuminating here; it appears to be a distinction between short-term fixed labor costs/long-term variable labor costs and short-term variable labor costs.

It is easier to see how the SCT theory saves a place for individual rationality to have causal efficacy. Subjective transactions costs make rationality more expensive; as a result, individuals are less rational when these costs are high. In other words, subjective transactions costs are *barriers* to the causal efficacy of rationality; when subjective transactions costs are reduced, individual rationality can more readily exercise its causal power. On the SCT theory, institutions function as enablers for rationality's causal power rather than replacing it or showing that it never had any causal power to begin with. But what about the zero-intelligence traders? One might argue that their success rate shows that rationality either lacks causal power or that the results of economic decision-making are overdetermined by institutions and individual reasoning. The advocate of the SCT theory could reply that the zero-intelligence traders do so well because the rules are clearly defined and all relevant learning has taken place; as a result, the traders mimic human rational decision-making. The SCT theory shows that there is no need to adopt the counterintuitive position that individual rationality has no causal power; it shows that individual rationality's causal power is consistent with the data.

(III.ii) The SCT theory has advantages over all four of our previously discussed proposals. First, it avoids the error of positivism by refusing to ignore individual psychology. But it also avoids claiming that markets often do not work because humans are irrational in some settings. It can accommodate the findings of cognitive psychologists, explaining why we rely on prospects and heuristics when we could spend the resources to do more expensive calculations. Further, it has a comfortable place for views that rely heavily on the causal impact of the environment and explains how learning and repetition reduce the subjective transactions costs of future actions. The great advantage the SCT theory has over more externalist theories is that it does not ascribe primary causal power to the environment. Rationality matters and has causal force. The SCT theory preserves the common sense intuition that rational deliberation has causal efficacy while acknowledging a vital role for the environment. But the SCT theory also has another important implication. Note that the SCT theory makes clear recommendations on how to produce efficient market outcomes: the best methods for producing efficient market outcomes involve reducing subjective transactions costs. This can occur through generating stable institutional rules, and the feedback-based learning that results from economic practice within these systems of rules. While these recommendations are abstract, there is no in principle reason why one could not

attempt to analyze successful and unsuccessful economic institutions based on these recommendations.

(III.iii) The SCT theory has a significant drawback, however. Smith comments that “considerations of SCT are hard to formalize within a framework as general as that attempted in [expected utility theory] and [weighted expected utility theory], that allow the latter to be deduced as limiting cases when SCT goes to zero, or when outcome values get large relative to a fixed SCT.”⁷⁶ The best way to accommodate these worries is to suggest some methods of quantifying subjective transactions costs in the lab and in real-world data collection. Yet formalization proves difficult. As a result, the SCT theory is hard to apply and test. There is no space to speculate on methods of quantification, however, so this will have to be left for future work. Despite this difficulty, the SCT theory has clear advantages over the four alternatives covered in Section II and thus merits an attempt to mitigate its deficiencies.

Section IV: Conclusion

This essay began with a simple problem: economists think the data show that markets work but their models assume that humans are rational economic actors, while psychologists think the data show that humans are not rational economic actors. There is an apparent tension that needs to be resolved. Four strategies for resolving the conflict were considered: (a) accepting methodological positivism in economics, (b) rejecting the effectiveness of markets, (c) building an alternative utility theory to accommodate the problem, or (d) resorting to an embedded cognition model of economic behavior. Strategies (a) and (b) were rejected as inadequate and while strategies (c) and (d) have merits, they need supplementation. In response, I advanced the subjective transactions costs (SCT) model, which embraces and explains the economic and psychological data, resolves the tension, and integrates the advantages of (c) and (d). The SCT theory assigns individual rationality a causal role, while explaining the importance of heuristics and the environment in facilitating rational choice; the theory predicts that economic actors will behave rationally when the subjective costs of transaction are low. The SCT theory may be the key to showing how human irrationality can result in efficient market outcomes.

⁷⁶ Smith, 268. Smith cites one experiment that tries to build in SCT elements. See Sydney Siegel, “Decision Making and Learning Under Varying Conditions of Reinforcement,” *Annals of the New York Academy of Science*, 1961, 89, 766-83.