Rationality with Preference Discovery Costs

Matthew S. Wilson*

June 10, 2017

Abstract

Economic theory assumes that preferences are rational. However, experiments have found small violations of transitivity. This paper develops a model of rationality with preference discovery costs. Introspection is costly. Thus, agents may find it optimal to use less than full effort, even though this raises the risk of making a poor choice. This model could potentially explain the intransitivities observed in the data while retaining rationality and optimization.

Key words: rationality, behavioral economics, experimental economics

JEL classification: D11, D12

* Economics Department, Truman State University, 100 East Normal Avenue, Kirksville, MO 63501 USA. Email: mwilson@truman.edu. Phone: (660) 785-4828. ORCID: 0000-0003-0490-590X.

Acknowledgements: William Harbaugh, James Andreoni, George Evans, Bruce McGough, and two anonymous referees provided many helpful comments. Any remaining errors are my own.

1 Introduction

"Preferences are complete and transitive." This assumption is one of the cornerstones of microeconomic theory. However, experiments routinely find that human behavior does not conform to this axiom perfectly (Battalio (1973), Andreoni and Miller (2002), Harbaugh et al. (2001)). These results have been known for a while, but standard theory continues to rely on rationality. Perhaps this is because rationality seems so reasonable. How can Bundle A be both better and worse than Bundle B? It is similarly difficult to believe that someone can deliberately prefer X over Y and Y over Z, and yet find Z better than X. Alternatively, preferences could be rational but people might not pick the option that is most preferred. However, this also appears nonsensical. Why would anyone intentionally choose an inferior combination of goods, knowing that superior alternatives are available?

Here is one response that economists frequently give: "Yes, there certainly are rationality violations in laboratory experiments with small stakes. However, if there was repetition or if the stakes were higher, people would learn from their mistakes and the anomalies would go away." A number of studies have tested this claim. However, the evidence is not as clear cut as many economists commonly believe. A heavily studied example is the preference reversal phenomenon. It is usually discussed in the context of risk theory, but it can also be interpreted as a violation of completeness and transitivity. A subject chooses between a "P-bet" (safe) and a "$-bet" (risky). Then the experimenter elicits the subject's valuation of the bets. In many cases, the subject places a higher value on the risky bet, but chooses the safe bet. This reveals the following preferences: *P-bet* > *$-bet* > *P-bet*. Cox and Grether (1996) and Loomes et al. (2003) both found that preference reversals occurred less frequently as the experiment was repeated. However, repetition does not solve everything. Braga et al. (2009) reported their results in the title of their paper: "Market Experience Eliminates Some Anomalies – And Creates New Ones." As subjects gained experience, a different kind of preference reversal occurred. Now people would often choose the risky bet, but they placed a higher value on the safe bet. Completeness and transitivity are still violated. Plott (1996) discussed the effect of repetition on rational behavior, surveying the literature on a wide variety of experiments. He found that violations tended to decrease with repetition. Nevertheless, some violations still remained at the end of the experiments.

Camerer et al. (1999) surveyed 74 papers on the effect of raising the stakes. One of their results (emphasis in the original): “We also note that *no* replicated study has made rationality violations disappear purely by raising incentives.” However, they did note many cases where higher stakes lead to lower variance. Some of the outliers vanish once subjects have a larger financial motivation, but this is not the same as eliminating all the anomalies. Smith and Walker (1993) reviewed 31 studies and found that higher stakes were associated with fewer rationality violations.

Another response to anomalies is to downplay their significance. Many indices have been developed to measure the severity of a violation (Varian (1991), Afriat (1972)). Intransitive choices waste money; the indices measure the fraction of the wealth that is wasted. According to these metrics, the magnitude of the violations found in experiments is quite small (Battalio (1973), Varian (1991), Harbaugh et al. (2001)). For instance, Varian (1991) demonstrates that in Battalio’s (1973) dataset, people seldom waste more than 5% of their endowment. Andreoni and Miller (2002) and Harbaugh et al. (2001) reach the same conclusion in their experiments.

I draw three main conclusions from the literature on rational choice:

1. Rationality violations do exist and have been replicated under many conditions.
2. Giving the subjects higher stakes or the opportunity to learn from experience might reduce the number of violations, but it will not eliminate them entirely.
3. The violations are small in magnitude.

Because departures from rationality are so small, perhaps we only need a slight modification of the theory. I propose adding *preference discovery costs* to the standard model of rational choice. In this new model, people do have rational preferences. However, accessing those preferences is costly. It takes effort to figure out what you want, and you dislike this work. However, if you put in more effort, you will likely receive a bundle that is better, i.e., one that gives higher utility according to your true (but imperfectly known) preferences. Effort is chosen optimally. Agents weigh the disutility of effort against its expected benefits. The cost of effort is known, but the expected benefits may be uncertain. Subjective beliefs are formed about the rewards of increased effort. These beliefs are updated optimally after consumption.

Since picking the most preferred bundle requires costly introspection, using less than full effort could be optimal. Such behavior could lead to apparently intransitive choices, but perhaps the individual is actually rational once preference discovery costs are accounted for. This could explain the small deviations from standard theory that we observe in the data while still retaining rational preferences and optimization. Since these assumptions are crucial to economics, it is important to show that they are compatible with the data.

Other efforts have been made to incorporate the costs of decision making. Their existence is widely acknowledged in many disciplines (Hogarth and Karelaia (2005), Bröder and Schiffer (2003)). Due to these decision-making costs, many papers outside of economics argue that people do not optimize; instead, they rely on heuristics. However, in my model, the choice between optimization and heuristics is a false one. People use heuristics because they are optimal when preference discovery costs are present. One feature of my model is that it does not require that we impose any functional form on the heuristic.

However, economists have made only a few attempts to model decision-making costs. Conlisk (1988) uses "optimization costs" that are a function of the time required to make a decision. Since it is one of the first papers in the genre, there is no general theory or experiment to test it; only a few applications are included and discussed. In all disciplines, a common objection to these models of decision-making cost is the "infinite recursion problem," which Conlisk (1988) mentions frequently. If there is a cost to making a decision, why is there not also a cost to deciding how much effort to use in making the decision on effort? Similarly, there could be a cost to deciding how much effort to use in deciding how to decide on the original question, *ad infinitum*. Most of these models have no response other to assume it away. One feature of my model is that it can accommodate any finite number of recursions.

Smith and Walker (1993) introduce a model of decision-making cost and find indirect evidence for it. They review 31 experiments in which different groups of subjects faced different levels of cash payoffs. As the potential payouts rose, the number of deviations from rationality consistently fell (Smith and Walker (1993)). This is exactly what we would expect if effort is costly. Raising the stakes increased the benefits of making good decisions, so subjects were willing to invest more effort. Plott (1996) takes a very similar approach, but unlike Smith and Walker (1993), he does not include a theoretical model to explain the violations. In Smith and Walker's (1993) model, preferences remain rational, but effort is not chosen optimally. Rather, the unobserved decision cost is a friction that allows choice to deviate from rational preferences. They write, "ε is not 'error' from the point of view of the subject weighing (albeit unconsciously) benefit against decision cost. It is the experimentalist who interprets ε as a prediction error of the theory" (Smith and Walker (1993)). However, economists are certainly inclined to retain optimization if possible.

The assumption that agents optimize is also relaxed in Evans and Ramey (1998). The model is about firms setting their expectations when calculation costs are present, but the same ideas can be applied to consumers making choices when there are decision costs. As calculation rises, firms make better estimates. Rational expectations is a limiting case when calculation approaches infinity. Firms use a heuristic; when the gain from a previous level of calculation exceeds a target level, the firm continues to calculate, but when it falls below the target, further calculation stops. One of their earlier papers kept optimization (Evans and Ramey (1992)). A difficult issue in this model and in other models of decision cost is how to determine effort/calculation optimally. It is easy to speak of agents weighing the benefits of higher quality decisions against the cost of effort, but the mathematics reveal a stubborn problem. How can agents be aware of the utility gain from more effort and yet be ignorant of their utility function? However, if the utility function is known, then the solution to the consumer's problem is already given by standard theory. Choosing anything else cannot be optimal. Hence, any such theory cannot explain the intransitivities we observe in the real world. Evans and Ramey (1992) had to assume the problem away. Firms costlessly knew the benefits of reoptimizing without knowing what reoptimization would require them to do. This makes the math work and also yields some implications for policy that the authors discuss. In several classical models with rational expectations, policy is neutral; firms anticipate and offset the government's actions. However, if calculation is costly, then policy may have limited effectiveness (Evans and Ramey (1992)). If the policy changes are small, then the gain from reoptimizing will not be enough to recover the calculation cost. Firms prefer sticking to their previous plan rather than offsetting the policy change (Evans and Ramey (1992)). Perhaps similar results can come from a theory of consumers with decision costs. Firms and the government may extract some benefits from consumer irrationality, but if they keep trying to increase those benefits, consumers may increase their decision-making effort and the intransitivities will vanish.

There is a related literature in risk theory on intransitive choices over lotteries. These non-expected utility models take a different approach to explaining anomalies. To see the difference, break up revealed preference axioms into two components:

1. People have complete, transitive, and monotonic preferences
2. In every budget set, people pick the bundle that they prefer the most. I.e., choices reveal preferences

Standard theory starts out by distinguishing between preferences and choice, but then the second assumption unifies the two. Most of the behavioral theories keep the second assumption but relax the first one. A key element is probability weighting, introduced by Tversky and Kahneman (1979). The agent does not use the actual probabilities in the lottery; instead, the probabilities are transformed using some function. Viscusi (1989) demonstrated that a Bayesian approach to probability weighting can resolve many of the challenges to the Expected Utility Theorem. However, intransitivity is not addressed. Bordley (1992) generalizes the model to accommodate intransitivity and several other paradoxes. Both models can explain violations of monotonicity, but that is a weakness rather than a strength. Consider the two lotteries in Table 1.

Table 1 *Two lotteries*

Lottery 1	
Outcome	Probability
\$10	30%
\$0	70%
Lottery 2	
\$10	20%
\$10	10%
\$0	70%

Clearly, the two lotteries are the same. However, if agents weight the probabilities, then they will almost surely fail to realize this. To see this, let $u(x)$ be the agent's von Neumann-Morgenstern utility function; $w_f(q)$ is the probability weighting function. Then the expected weighted utility from the first lottery is $w_f(0.3)\ u(10) + w_f(0.7)\ u(0)$. For the second lottery, it is $\big(w_f(0.2) +$

$w_f(0.1)\big)\ u(10) + w_f(0.7)\ u(0)$. These will always be equal if $w_f(0.2) + w_f(0.1) = w_f(0.3)$. But this is not true in either model, except in the special case where they reduce to the Expected Utility Theorem. The same problem arises in other probability weighting models. The reason they allow possibilities such as $w_f(0.2) + w_f(0.1) \neq w_f(0.3)$ is that they want to explain why people overweight low probabilities and other paradoxes. The drawback is that in the model, people will make obvious and unrealistic violations of monotonicity. Suppose that Lottery 2 is strictly preferred to Lottery 1 and preferences are continuous. Then there exists $\varepsilon > 0$ such that Lottery 1's payoff of $10 can be changed to $\$(10 + \varepsilon)$ and the agent will still prefer Lottery 2.

People do commit violations of montonicity and transitivity – but not when it's that obvious. Models that rely on the second assumption (choices reveal preference) cannot explain this. If choices reveal preference and preferences are intransitive, then violations will occur even in the most obvious cases. Tversky and Kahneman (1986) demonstrated the effect of transparency on violations. Table 2 shows some of the lotteries they used.

Table 2 *Lotteries in Tversky and Kahneman (1986)*

Lottery A			**Lottery C**	
Outcome	Probability		Outcome	Probability
$0	90%		$0	90%
$45	6%		$45	6%
$30	1%		$30	1%
-$15	1%		-$15	3%
-$15	2%			
Lottery B			**Lottery D**	
$0	90%		$0	90%
$45	6%		$45	7%
$45	1%		-$10	1%
-$10	1%		-$15	2%
-$15	2%			

In one part of the experiment, subjects chose between A and B. Without exception, they picked B. In that case, it is easy to compare the outcomes and probabilities across lotteries, so no one committed any violations. Later, they had to pick between C and D. Here it is not quite so clear if one lottery dominates the other. The majority chose C, but a more careful look reveals that D dominates C. Thus, violations do occur, but only when people fail to exert full effort.

This suggests that the problem with standard theory is *not* that preferences are irrational. People avoid violations as long as they are obvious or if enough effort is applied. The problem is that people do not always choose according to their underlying rational preferences. This is due to preference discovery costs. Thus, the model in this paper is more realistic than ones with intransitive preferences.

There is another attempt to retain rationality: change of tastes models. It is easy to see how this might explain intransitivities. For example, after consuming Brand B, you find it better than expected, leading you to consume more of it and make choices inconsistent with your previous decisions. Conlisk (1996) observes that these models are very similar to ones with imperfect information about stable preferences, such as DeGroot (1983). “I changed my behavior because my unstable preferences changed” is difficult to distinguish from “I changed my behavior because I learned more about my underlying stable preferences.”

However, this does not fully address the violations that occur in the laboratory. In a common procedure, used in Harbaugh et al. (2001) and others, subjects do not consume their payoff until the end of the experiment. None of the mechanisms that are used to explain why tastes shift – social preferences, culture, emotions (Jacobs (2016)) – have greatly changed within the experiment. In addition, the participant has not gained more information about their preferences for the good, since they have not consumed it yet. In short, within such an experiment, preferences should be stable and choices should be consistent. However, intransitivities remain. This suggests that changing tastes are not the source of all intransitivities. Ballinger and Wilcox (1997) discuss a similar issue in their paper on risk: “[M]any subjects’ choices are still inconsistent across identical trials of choice problems...This suggests that, for whatever reason, discrete choice is inherently probabilistic.” The reason is preference discovery cost. When subjects do not use maximum effort, they make random errors in their choices. This is why my model can explain such behavior while change of tastes models cannot.

Standard theory is incompatible with the intransitive choices observed in the data.

However, many of the models that try to resolve these anomalies have issues of their own. If probability weights are used, then highly implausible violations of monotonicity can occur. Change of tastes models do not explain all of the intransitivities in the laboratory. Preference discovery cost models avoid these flaws. However, they have challenges of their own: the recursion problem and the difficulty of selecting optimal effort with an unknown utility function. The model in this paper solves both of those problems. Since rationality and optimization are retained, it is a minimal modification of standard theory, and yet it still reconciles rational preferences with the data.

2 Theory

The utility function takes the form

(1) $U(x,d) = u(x) - v(d)$.

Clearly the individual's preferences are rational; otherwise no utility function would exist. Here $u(x)$ is the utility from consuming the vector of goods x. The vector of prices associated with these goods is p; the consumer has wealth w to spend. Assume that $u(x)$ is continuous and locally nonsatiated. The effort level is d and the preference discovery cost is $v(d)$. Negative effort is impossible, so $d \geq 0$. There can be no disutility from effort when no effort is exerted, which suggests the normalization

(2) $v(0) = 0$.

All else equal, people prefer less effort to more, so the preference discovery cost $v(d)$ is monotonically increasing in d.

Standard theory provides a useful benchmark. Let $x^*(p,w)$ be the unboundedly rational choice of x, i.e., the solution to

(3) $\max u(x)$ *subject to* $p \cdot x \leq w$.

This is simply the Walrasian demand correspondence. As discussed in the introduction, if agents optimize and know their utility function, then we are back to standard theory and cannot explain intransitivities. Therefore, the utility function $u(x)$ is unknown. However, $v(d)$ is part of the agent's information set. The agent's task is to maximize $E[U(x,d)]$.

Because people rarely use full effort in making their choices, there is an apparent imprecision to our preferences. One speaks of cups of coffee rather than molecules of coffee. It is almost surely true that one cup and another cup will not contain the exact same number of molecules. However, the difference is small and the effort required to find the perfect number of molecules is great; the imprecision is quite tolerable. Thus, the quantity actually chosen will probably be a random number of molecules close to the optimal amount. Standard theory implies that there is a unique solution if the budget constraint is linear and preferences are strictly convex. Deviating by even a molecule would be a violation.

Thus, it is more realistic to treat consumption as a random variable. It will be a function of effort, so actual consumption will be $x(p,w,d)$. Mathematically, we model the agent as choosing effort optimally rather than consumption. After selecting effort, the agent receives a random quantity of consumption. It may be useful to think of consumption as being chosen via a heuristic. d can be thought of as the effort required to implement the heuristic. When choosing between heuristics, the agent considers how much effort each one would require and how good each one is likely to be at making the best choice. However, the agent does not yet know the resulting consumption choice. That is only resolved after the heuristic is chosen. Thus, when deciding on effort d, consumption is a random variable from the consumer's viewpoint. A strength of this model is that no functional form is imposed upon the heuristic. This makes the model more general and allows it to be applied in a wide variety of situations.

How does the agent find the optimal level of effort? The preference discovery cost $v(d)$ is known but the utility function is not. This brings us back to the problem mentioned earlier: how can agents be aware of the utility gain from more effort and yet be ignorant of their utility function? Here is the solution. Agents have beliefs about the distribution of $u(x^*(p,w)) - u(x(p,w,d))$. To simplify notation, call this variable y.

(4) $y(p,w,d) = u(x^*(p,w)) - u(x(p,w,d))$

$y(p,w,d)$ is the difference between a perfectly rational agent's utility and the utility from consumption of our boundedly rational individual. Call this the "utility gap." Since the boundedly rational agent can never outperform a perfectly rational agent, clearly $y(p, w, d) \geq 0$.

The true pdf of y is not known – that is why agents form beliefs about it (Section 3 describes how these beliefs are formed). It is believed to have pdf $f(y|p, w, d)$[1]. The agent would like to solve the problem

$$(5)\ \max_{\{d\}} E[u(x(p, w, d))] - v(d).$$

Unfortunately, $u(x)$ and the density of x are unknown. Note that the agent could solve a similar problem:

$$(6)\ \max_{\{d\}}\ -E[y|p, w, d] - v(d)$$

This is because the agent knows $v(d)$ and has beliefs about the distribution of $y(p, w, d)$. The next two propositions show how the agent overcomes the challenges of finding the optimal level of effort.

PROPOSITION. *The level of effort* d *that solves (6) also solves (5).*

PROOF. By definition of $y(p, w, d)$, we know that (6) can be rewritten as follows.

$$(7)\ \max_{\{d\}} E[u(x(p, w, d)) - u(x^*(p, w))] - v(d).$$

Recall that $x^*(p, w)$ is the quantity that a perfectly rational agent would choose. Standard theory does not consider the cost of effort, so $x^*(p, w)$ and $u(x^*(p, w))$ are independent of d. Thus, the level of effort that solves (7) also solves (5). Since (7) is equivalent to (6), the proposition is proven. □

1 – Note that even if this belief about the distribution was correct and the agent was aware of this, it would still be impossible to recover the utility function. Knowledge of $u(x^*(p, w)) - u(x(p, w, d))$ does not imply knowledge of $u(x)$.

As a result of the proposition, we can rewrite the agent's problem in a form that can be solved. The next proposition shows the solution.

PROPOSITION. *If all the functions are differentiable, then the first order condition for level of effort* d *that solves the agent's problem is as follows.*

$$(8)\ d\left(\frac{\partial}{\partial d}\left(-\int_0^\infty y\, f(y|p,w,d)\,\mathrm{d}y\right)-v'(d)\right)=0.$$

PROOF. Assuming that there is an interior solution, the first order condition is

$$(9)\ \frac{\partial}{\partial d}\left(-\int_0^\infty y\, f(y|p,w,d)\,\mathrm{d}y\right)-v'(d)=0$$

However, it is possible that the solution is not on the interior; exerting zero effort may be optimal. Thus, either $d=0$ or d is given by (9). Both of these possibilities are expressed in (8), proving the proposition. □

This is the basic model. Here is an illustration before we deal with issues such as the infinite recursion problem.

<u>Example</u>

For simplicity, suppose there are two goods, *x* and *z*, that are perfect substitutes. The individual's utility function is $u(x,z)=x+z$ and $v(d)=\frac{7}{2}\log(d+1)$. There are two budget lines, B_0 and B_1. On budget B_0, prices are $p_x^0=4$, $p_z^0=5$, and he has wealth $w^0=160$. He correctly believes that

$$(10)\qquad y_0(p^0,w^0,d_0)=$$

$$u(x_0^*,z_0^*)-u[x(d_0),z(d_0)]\sim Uniform\left(0,\frac{\max\{w^0/p_x^0,\ w^0/p_z^0\}-\min\{w^0/p_x^0,\ w^0/p_z^0\}}{1+d}\right).$$

Since the goods are perfect substitutes, a fully rational person would just buy whichever one is cheaper. In other words, $x_0^*=40,\ z_0^*=0$. Thus, $u(x_0^*,z_0^*)=40$. More generally, $u(x_0^*,z_0^*)=$

$\max\{w^0/p_x^0,\ \ w^0/p_z^0\}$. Note that the worst choice on the frontier of the budget line is picking the opposite corner solution: $x_0^* = 0,\ z_0^* = 30$. This would lead to $u(x,z) = 30$. More generally, picking the wrong corner gives utility $u(x,z) = \min\{w^0/p_x^0,\ \ w^0/p_z^0\}$. The point: as d approaches infinity, the agent is perfectly rational; when $d = 0$, the agent picks from the budget line completely at random.

Note that the optimal level of effort is $d_0^* = 1/7$. This implies that the agent could select any bundle on the budget line between (5, 28) and (40, 0). This wide range of outcomes should not be surprising; when the slope of the budget line is very close to the slope of the indifference curve, the payoff to additional effort is low. Thus, effort will be low and the dispersion of choices will be high.

Now consider another budget line, B_1. Here $p_x^1 = 4$, $p_z^1 = 3$, and $w^1 = 120$. Now

(11) $$y_1(p^1, w^1, d_1) =$$

$$u(x_1^*, z_1^*) - u[x(d_1), z(d_1)] \sim Uniform\left(0, \frac{\max\{w^1/p_x^1,\ \ w^1/p_z^1\} - \min\{w^1/p_x^1,\ \ w^1/p_z^1\}}{1+d}\right).$$

Then $d_1^* = 3/7$ and we know that $x_1^* = 0,\ z_1^* = 40$. The choice could range from (0, 40) to (21, 12). Figure 1 displays the range of possible choices.

Fig. 1 *The agent will pick from the dotted region of the budget lines.*

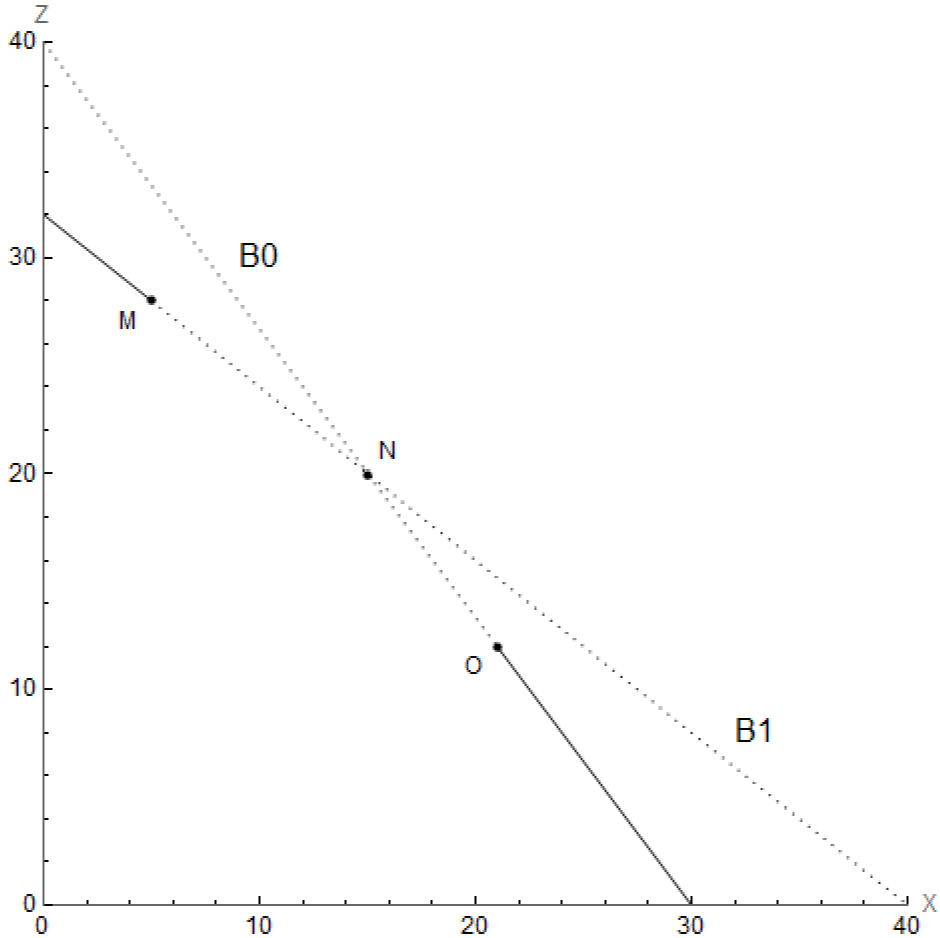

In standard theory, choosing a point from both $\overline{NO}$ on B_0 and $\overline{MN}$ on B_1 would reveal intransitive preferences[2]. However, our boundedly rational agent makes such choices 8.16% of the time. Thus, apparently irrational preferences can be reconciled with rationality if we account for preference discovery costs.

3 Extensions

To pick effort optimally, agents need to have beliefs about the benefits of effort. This is then weighed against the costs in order to find the solution. In this section, we will see where these beliefs come from. In addition, the recursion problem will be solved.

2 – Exception: picking *N* on both budget lines.

The Recursion Problem

In the model above, effort is costly since discovering one's preferences is work. However, it is implicitly assumed that deciding on the optimal amount of effort is cost-free. A natural objection is that it also takes effort to decide on how much effort to use. Further effort is required to determine how much effort to apply to the problem of deciding how much effort to use in picking the effort level for the original consumption problem, etc. There could be any number of recursions in the decision problem. The model can be extended to accommodate N recursions, where N is any positive integer.

Originally, the consumer chose d and received a random draw of consumption $x(p, w, d)$. Consumption was random because the consumer did not use enough effort to perfectly determine his preferences[3]. Now there are N recursions. The agent uses effort d_0 to decide on consumption. Effort d_1 is applied to finding d_0; this is how much effort you use to decide on the effort used on the consumption problem. Similarly, d_2 is the effort used to decide on the effort level d_1 that is used to decide on effort d_0 that is used to decide on consumption, etc. The individual will initially pick d_N optimally. Instead of getting a random draw of consumption, the agent will get a random draw of effort d_{N-1} in the next level of the recursion problem. This will lead to another random draw of effort d_{N-2} and so on. Eventually the agent reaches the last level of recursions and uses effort d_0 in the consumption problem. Mathematically, the agent's problem is

$$(12) \qquad \max_{\{d_N\}} E[u(x(p, w, d_0)) - \sum_{i=0}^{N-1} v_i(d_i(p, w, d_{i+1}))] - v_N(d_N).$$

In the previous model, the agent knew the functional form of $v(d)$ but not $u(x)$. Similarly, in this case $v_k(d_k)$ is known for $k = 0, 1, \dots, N$, but $u(x)$ is unknown. Since the preference discovery cost is $v_k(d_k)$ instead of $v(d_k)$ at recursion k, the disutility of effort can differ across recursion levels.

Now to specify the agent's beliefs about the benefits of effort. Let d_k^* be the optimal

3 – Even if the consumer did use enough effort to perfectly determine his preferences, consumption could still be modeled as a random variable. It would simply have a degenerate distribution.

amount of effort at recursion level k. In the earlier model, the agent thought that $y(p, w, d) = u(x^*(p, w)) - u(x(p, w, d))$ had pdf $f(y|p, w, d)$. Now extend the same idea to allow for N recursions:

(13) $$y_{M+1}(p, w, \boldsymbol{d}) = \left[u\left(x^*(p, w)\right) - \sum_{i=0}^{M} v_i(d_i^*)\right] - \left[u\left(x(p, w, d_0)\right) - \sum_{i=0}^{M} v_i\left(d_i(p, w, d_{i+1})\right)\right]$$
$$for\ M = 0, 1, 2, \dots, N-1.$$

(14) $$y_0(p, w, d_0) = u(x^*(p, w)) - u(x(p, w, d_0))$$

Here $\boldsymbol{d}$ is a vector containing $d_0, d_1, \dots, d_{M+1}$. Note that $y_{M+1}(p, w, \boldsymbol{d}) = y_M(p, w, \boldsymbol{d}) - \left[v_M(d_M^*) - v_M\left(d_M(p, w, d_{M+1})\right)\right]$. The beliefs are that $\left[v_M(d_M^*) - v_M\left(d_M(p, w, d_{M+1})\right)\right]$ has pdf $g_M\left(\left[v_M(d_M^*) - v_M\left(d_M(p, w, d_{M+1})\right)\right] | p, w, \boldsymbol{d}\right)$. Given this and the belief that $y_0(p, w, d_0)$ has pdf $f_0(y_0|p, w, \boldsymbol{d})$, the agent can derive the distribution $f_{M+1}(y_{M+1}|p, w, \boldsymbol{d})$. The agent's problem can be rewritten as

(15) $$\max_{\{d_N\}} E\left[-y_N(p, w, \boldsymbol{d}) + \left[u\left(x^*(p, w)\right) - \sum_{i=0}^{N-1} v_i(d_i^*)\right]\right] - v_N(d_N).$$

Note that d_i^* is independent of d_N for all $i < N$, just as $u(x^*(p, w))$ was independent of d in the original model with no recursions. Thus, the first order condition is

(16) $$\frac{\partial}{\partial d_N}\left(-\int_0^{\infty} y_N(p, w, \boldsymbol{d})\, f(y_N|p, w, \boldsymbol{d})\, \mathrm{d}d_N\right) - v_N'(d_N) \leq 0.$$

This shows that the recursion problem can be solved by the agent for any finite number of recursions; it is not an issue for this model. To simplify the math, in the rest of the paper I will assume that there are no recursions, though the results can be extended to any whole number of them.

Optimal Beliefs

Recall that $y(p, w, d) = u(x^*(p, w)) - u(x(p, w, d)) \geq 0$ is the "utility gap." It is the difference between a perfectly rational agent's utility and the utility from consumption of our boundedly rational individual. To find the optimal level of effort, the agent used his beliefs about the pdf $f(y|p, w, d)$. Here I show how beliefs evolve in response to consumption. It is natural to think that if a particular decision went badly, then the agent will want to avoid making the same mistake. He may realize that he underestimated the benefits of using more effort, and thus will revise his beliefs about $f(y|p, w, d)$ accordingly. This can be modeled more formally.

To update beliefs, the agent needs to gain some information from consumption. However, it must not be enough to allow him to recover the true preferences $u(x)$. If it were, then we are back to standard theory and cannot explain anomalies. This challenge can be met if the agent receives a noisy signal s of the utility gap after consuming $x(p, w, d)$.

(17) $$s = y + \varepsilon$$

Assume that ε is white noise and has known pdf $f(\varepsilon)$. Since a boundedly rational agent can never achieve more utility than a perfectly rational one, y is always nonnegative, so it makes sense to impose $s \geq 0$. If the signal were negative, then the individual would know that this was driven by the random shock ε.

Some notation is needed before proceeding. Let θ be an estimated vector of parameters that is updated in response to the signal. The dimensions of θ are $n \times 1$. The true θ is θ^*. The agent initially believes that $y, p, w, d,$ and θ have pdf $f(y, p, w, d, \theta)$. The optimal level of effort given beliefs is $\hat{d}$. In other words, $\hat{d} = argmax\ E[\left(u\left(x^*(p, w)\right) - y(p, w, d) - v(d)\right)|\theta, p, w] \equiv D(\theta, p, w)$. Equation (8) showed how to solve for $\hat{d}$. Assume that $D(\theta, p, w)$ is continuously differentiable. The information set is I. Prices, wealth, and effort are all observed, and after consumption, the signal s is also observed. Hence, $p, w, d, s, f(\varepsilon)\ \epsilon\ I$. Lastly, the dimensions of p are $L \times 1$.

The initial beliefs are exogenous. The task is to update the beliefs optimally given the signal. This will help the agent in the next consumption problem he faces.

(18) $$\max_{\{\theta\}} E[U(x,\hat{d})|I]$$

Recall that $x^*(p,w)$ is what a perfectly rational agent would choose. That is independent of our boundedly rational agent's beliefs about effort, so $x^*(p,w)$ is not a function of θ. The preference discovery cost $v(d)$ is known, so it can be pulled out of the expectation operator. Now the problem can be rewritten as

(19) $$\max_{\{\theta\}} E\left[u(x^*(p,w)) - y(\hat{d})\middle|I\right] - v(\hat{d}).$$

(20) $$\min_{\{\theta\}} E\left[y(\hat{d})\middle|I\right] + v(\hat{d})$$

From Bayes' Rule, we know

(21) $$f(y|s,p,w,d,\theta) = \frac{f(s,p,w,d,\theta|y)\, f(y)}{f(s,p,w,d,\theta)}.$$

By definition, the cumulative distribution function is

(22) $$F(s,p,w,d,\theta|y) = \Pr[(y+\varepsilon,P,W,D,\Theta) \leq (s,p,w,d,\theta)|y] = \Pr[(\varepsilon,P,W,D,\Theta) \leq (s-y,p,w,d,\theta)|y].$$

Since ε is white noise,

(23) $$f(\varepsilon,p,w,d,\theta) = f(\varepsilon)f(p,w,d,\theta) = f(\varepsilon)\int_0^\infty f(y,p,w,d,\theta)\, dy.$$

Therefore,

(24) $$F(s,p,w,d,\theta|y) = \int_{-\infty}^{\theta_n} \dots \int_{-\infty}^{\theta_1} \int_0^{p_L} \dots \int_0^{p_1} \int_0^{d} \int_0^{w} \int_{-y}^{s-y} \left(\int_0^{\infty} f(Y,P,W,D,\Theta)dY \right) dF(\varepsilon)\, dW\, dD\, dP_1 \dots dP_L d\Theta_1 \dots d\Theta_n$$

(25) $$f(y|s,p,w,d,\theta) = \frac{\frac{\partial^{L+n+3}}{\partial s \partial p \partial w \partial d \partial \theta}\left(f(y) \int_{-\infty}^{\theta_n} \dots \int_{-\infty}^{\theta_1} \int_0^{p_L} \dots \int_0^{p_1} \int_0^{d} \int_0^{w} \int_{-y}^{s-y} \left(\int_0^{\infty} f(Y,P,W,D,\Theta)dY \right) dF(\varepsilon)\, dW\, dD\, dP_1 \dots dP_L d\Theta_1 \dots d\Theta_n \right)}{\int_0^{\infty} \frac{\partial^{L+n+3}}{\partial s \partial p \partial w \partial d \partial \theta}\left(\int_{-\infty}^{\theta_n} \dots \int_{-\infty}^{\theta_1} \int_0^{p_L} \dots \int_0^{p_1} \int_0^{d} \int_0^{w} \int_{-y}^{s-y} \left(\int_0^{\infty} f(Y,P,W,D,\Theta)dY \right) dF(\varepsilon)\, dW\, dD\, dP_1 \dots dP_L\, d\Theta_1 \dots d\Theta_n \right) dF(y)}$$

This can be substituted into $E\left[y(\hat{d})\middle|I\right] = \int_0^{\infty} y\ f(y|s,p,w,d,\theta)\ dy \equiv h(\hat{d},\theta)$. Now the agent's belief updating problem can be solved.

(26) $$\min_{\{\theta\}} E\left[y(\hat{d})\middle|I\right] + v(\hat{d})$$

Recall $\hat{d} = D(\theta,p,w)$. If $D(\theta,p,w)$ and $h(\hat{d},\theta)$ are continuously differentiable, then the first order condition is

(27) $$\frac{\partial\left(E\left[y(\hat{d})\middle|I\right]+v(\hat{d})\right)}{\partial\theta} = v'\left(\hat{d}\right)\frac{\partial\hat{d}}{\partial\theta} + h_1\frac{\partial\hat{d}}{\partial\theta} + h_2 = 0.$$

As the number of signals (N_s) approaches infinity, the effort level $D(\theta,p,w)$ will converge to $D(\theta^*,p,w)$, i.e., the effort the agent would have chosen if the true θ were known. This result holds under very general conditions. The proof is in Appendix A. To clarify, this is *not* a claim that the agent's consumption choices converge to full rationality. However, the model can accommodate learning, as demonstrated in the next section.

Learning from Experience

There is evidence of learning from experience, but as discussed in the introduction, this does not always suffice to eliminate intransitivities. How can this be modeled? In DeGroot (1983), agents start with prior beliefs about the parameters in their utility function. After consumption, they update their beliefs using Bayes rule. This works for certain functional forms in DeGroot (1983); my model can be extended to explain learning from experience for more general utility functions. The utility function remains unknown to the agent. However, after gaining experience, the agent can make better choices with less effort than before. As in the previous section, the agent receives a noisy signal after each choice and N_s is the number of signals. By "better choices," I mean that the utility gap decreases; i.e., the agent's utility from consumption gets closer to the utility of the perfectly rational agent. In other words, the expected value of y is a decreasing function of N_s.

(28) $$\frac{\partial E[y|p,w,d,\theta^*,N_s]}{\partial N_s} < 0$$

The literature on the multi-armed bandit problem can help us refine the learning from experience condition. In this literature, the usual story is about finding the optimal strategy for playing the slot machine at a casino. The slot machine has two or more arms. At each point in time, the agent pulls one of the arms. After each pull, the agent receives a payoff of zero or one. For each arm, the probability of receiving a payoff of one is unknown. Interest in this problem extends well beyond casino gamblers; Jun (2004) notes that it has been applied to "clinical trials, optimal experiments, new product development, job search, oil exploration, research & development, technology choice, and resource allocation."

Suppose the agent is playing on a three-armed slot machine. To learn the probabilities, the agent might experiment by pulling the first two arms a couple of times. However, this experience does not provide any information about the third arm. Here is the insight for my model: experience is more helpful if it was obtained in a situation similar to the problem at hand. As before, the utility gap will shrink as the number of signals increases. This effect will be stronger if prices and wealth in the previous choices are similar to prices and wealth in the current scenario.

To distinguish between current and previous budgets, more notation is needed. Let w^k be wealth in the *k*-th choice; W is a vector of length N_s containing $w^1, w^2, \ldots, w^{N_s}$. Similarly, p_j^i is the price of good *j* in the *i*-th choice. All of the p_j^i's are in vector P. Without loss of generality, assume that $\theta_i^* > 0$ for all *i*. There are many ways to measure the similarity between two budget sets. Due to this, I will only give examples rather than a general result.

In the simplest case, consider an agent who has made one choice already and is currently facing a second choice.

(29) $$E[y|P,W,d,\ \theta^*] = \frac{\theta_1^*}{(\theta_2^*+d)} m_1 m_2$$

(30) $$m_1 = \frac{1}{\theta_3^*+\theta_4^*}\left(\theta_3^* + \theta_4^* - \frac{\theta_4^*}{1+\left(\frac{w^2}{p_1^2}-\frac{w^1}{p_1^1}\right)^2}\right)$$

(31) $$m_2 = \frac{1}{\theta_5^*+\theta_6^*}\left(\theta_5^* + \theta_6^* - \frac{\theta_6^*}{1+\left(\frac{w^2}{p_2^2}-\frac{w^1}{p_2^1}\right)^2}\right)$$

When the two budget sets are more similar, $\frac{w^2}{p_1^2} - \frac{w^1}{p_1^1}$ and $\frac{w^2}{p_2^2} - \frac{w^1}{p_2^1}$ approach zero. This shrinks the m_1 and m_2 terms. As a result, $E[y|P,W,d,\ \theta^*]$ falls. This means that when the budget sets are more similar, your expected utility gap decreases. You make better choices because your experience is more relevant; you approach full rationality.

Equation (32) reveals another property.

(32) $$\frac{\partial E[y|P,W,d,\theta^*]}{\partial d} = \frac{-\theta_1^*}{(\theta_2^*+d)^2}\ m_1\, m_2 < 0$$

This says that the expected utility gap decreases as you put in more effort. This makes sense, since we expect that applying more effort will bring you closer to full rationality.

The next set of equations form a more general example. In this case, the agent has made N_s choices already and is considering what to do in the next choice.

(33) $$E[y|P,W,d,\ \theta^*] = \frac{\theta_1^*}{(\theta_2^*+d)} \prod_{b=1}^{N_s} \prod_{l=1}^{L} m_{lb}$$

(34) $$m_{lb} = \frac{1}{\tilde{\theta}_{lb}^*+\theta_{lb}^*}\left(\tilde{\theta}_{lb}^* + \theta_{lb}^* - \frac{\theta_{lb}^*}{1+\left(\frac{w^{N_s+1}}{p_l^{N_s+1}}-\frac{w^b}{p_l^b}\right)^2}\right)$$

When $L = 2$ and $N_s = 1$, it reduces to the earlier example. Note that m_{lb} is bounded between

$\frac{\tilde{\theta}_{lb}^*}{\tilde{\theta}_{lb}^* + \theta_{lb}^*} > 0$ and one. Thus, for each additional choice, $E[y|P, W, d,\ \theta^*]$ is multiplied by a number between zero and one. In other words, the utility gap cannot increase when you gain more experience. Even if the new situation is entirely different from the previous ones you have faced (i.e., $m_{lb} = 1$), your experience does not harm you; it merely fails to help you. If the choices you have faced in the past are similar to the one you have now, then m_{lb} declines and the expected utility gap falls. Your experience causes you to approach to full rationality.

4 Conclusion

Much of economic theory relies upon everyone being fully rational. However, experiments frequently claim to find violations of this assumption. If the claims are true, then the implications are quite troubling. For the bulk of microeconomic theory, it is not known if the results are robust to violations of transitivity.

However, the model in this paper supports the assumption of rationality. Apparent violations of revealed preference axioms might not be contradictions at all. Instead, people could be – very rationally – weighing the costs and benefits of effort and deciding to use less than 100% of it. People can rationally dislike effort and find that the cost of working harder outweighs the risk of making a bad choice.

The model has many desirable properties. Optimization is retained throughout. Agents choose effort optimally given their beliefs and update their beliefs optimally. The recursion problem which has troubled many decision-cost models is resolved here.

In further research, it may be fruitful to investigate computational costs. Some errors in human decisions do not appear to be the result of preference discovery costs. For instance, in games of strategy, mistakes do not arise due to one player failing to discover that he had a preference for winning. Rather, the player wanted to win but did not manage to calculate the right strategy. The costs of calculation will most likely depend on the agent's cognitive ability; these costs may be substantially different from the costs of discovering one's preferences.

The rationality assumption remains defensible. All that is required is a very minimal addition – preference discovery costs – and the anomalies can be explained.

APPENDIX A

CONVERGENCE OF EFFORT

Let N_s be the number of signals the agent receives. Only a few harmless assumptions are required to show convergence: as N_s approaches infinity, $D(\theta, p, w)$ approaches $D(\theta^*, p, w)$.

Assumption 1: θ^* is unique and dwells in a compact subset of R^n.
Assumption 2: For all $p, w, d \geq 0$, the agent's initial beliefs $f(y, p, w, d, \theta) > 0$ when $\theta = \theta^*$.
Assumption 3: There exists a valid instrumental variable for *d* that is in the agent's information set.

Before going through the mathematical proof, some intuition is in order. An alternative way to update beliefs is by using nonlinear least squares with instrumental variables (NLLS-IV). We instrument for *d* because it is endogenous; the agent selects *d* optimally given his beliefs θ. So if there is an instrument for *d*, the agent can obtain consistent estimates for θ using NLLS-IV. Thus, $\lim_{N_s \to \infty} \theta_{NLLS\_IV} = \theta^*$. This implies that $\lim_{N_s \to \infty} D\left(\theta_{NLLS\_IV}, p, w\right) = D(\theta^*, p, w)$.

By definition of optimality, the optimal beliefs are no worse than any other beliefs. Therefore, if effort under the non-optimal beliefs θ_{NLLS_IV} converge, then effort under the optimal beliefs will also converge. Now for the math.

Let $\hat{\theta}$ denote the optimal beliefs. The beliefs under NLLS-IV are θ_{NL}. The information set I augmented with N_s signals is I_{N_s}. Since θ^* is the true θ, if you knew θ^*, then other beliefs θ would be irrelevant for assessing the expected value of the utility gap. Hence, $\int_0^\infty y\, f(y|p, w, d, \theta^*, \theta_{NL}, \hat{\theta}, I_{N_s})\, dy = \int_0^\infty y\, f(y|p, w, d, \theta^*)\, dy$. First, consider the problem faced by an agent who knows θ^*.

$$(35) \qquad \max_{\{d\}} E[u\left(x(p, w, d)\right)|p, w, d, \theta^*, \theta_{NL}, \hat{\theta}, I_{N_s}] - v(d)$$

This is equivalent to the following.

$$(36) \qquad \min_{\{d\}} E\left[y(p, w, d)\middle|p, w, d, \theta^*, \theta_{NL}, \hat{\theta}, I_{N_s}\right] + v(d)$$

Maximizing utility is equivalent to minimizing the utility gap plus the disutility of effort. The problem above is the same as

(37) $$\min_{\{d\}} E[y(p,w,d)|p,w,d,\theta^*] + v(d).$$

The solution is $d = D(\theta^*, p, w)$. Since $D(\theta^*, p, w)$ is optimal given $p, w,$ and θ^*, we know that for any θ_0,

(38) $$E[y(p,w,D(\theta^*,p,w))|p,w,d,\theta^*] + v(D(\theta^*,p,w)) \leq E[y(p,w,D(\theta_0,p,w))|p,w,d,\theta^*] + v(D(\theta_0,p,w)).$$

One such θ_0 is θ_{NL}. Since the NLLS-IV estimator is consistent, $\lim_{N_s\to\infty} \theta_{NL} = \theta^*$. Thus, $\lim_{N_s\to\infty} D(\theta_{NL}, p, w) = D(\theta^*, p, w)$ and $\lim_{N_s\to\infty} E\left[y(p,w,D(\theta_{NL},p,w))\middle|p,w,d,I_{N_s}\right] + v(D(\theta_{NL},p,w)) = E[y(p,w,D(\theta^*,p,w))|p,w,d,\theta^*] + v(D(\theta^*,p,w))$. Another possible θ_0 is the optimal beliefs $\hat{\theta}$. Therefore, we know

(39) $$\lim_{N_s\to\infty} E\left[y(p,w,D(\theta_{NL},p,w))\middle|p,w,d,I_{N_s}\right] + v(D(\theta_{NL},p,w)) \leq \lim_{N_s\to\infty} E\left[y\left(p,w,D(\hat{\theta},p,w)\right)\middle|p,w,d,I_{N_s}\right] + v\left(D(\hat{\theta},p,w)\right).$$

However, because $\hat{\theta}$ is the optimal beliefs, we know that for all N_s,

(40) $$E\left[y(p,w,D(\theta_{NL},p,w))\middle|p,w,d,I_{N_s}\right] + v(D(\theta_{NL},p,w)) \geq E\left[y\left(p,w,D(\hat{\theta},p,w)\right)\middle|p,w,d,I_{N_s}\right] + v\left(D(\hat{\theta},p,w)\right).$$

Together, the two equations above imply

(41) $$\lim_{N_s\to\infty} E\left[y(p,w,D(\theta_{NL},p,w))\middle|p,w,d,I_{N_s}\right] + v(D(\theta_{NL},p,w)) =$$

$$\lim_{N_s \to \infty} E\left[y\left(p, w, D(\hat{\theta}, p, w)\right) \middle| p, w, d, I_{N_s}\right] + v\left(D(\hat{\theta}, p, w)\right).$$

Since $\lim_{N_s \to \infty} E\left[y(p, w, D(\theta_{NL}, p, w)) \middle| p, w, d, I_{N_s}\right] + v(D(\theta_{NL}, p, w)) =$
$E[y(p, w, D(\theta^*, p, w))|p, w, d, \theta^*] + v(D(\theta^*, p, w))$,

$$\lim_{N_s \to \infty} E\left[y\left(p, w, D(\hat{\theta}, p, w)\right) \middle| p, w, d, I_{N_s}\right] + v\left(D(\hat{\theta}, p, w)\right) = \lim_{N_s \to \infty} E\left[y\left(p, w, D(\hat{\theta}, p, w)\right) \middle| p, w, d, \theta^*\right] + v\left(D(\hat{\theta}, p, w)\right) = E[y(p, w, D(\theta^*, p, w))|p, w, d, \theta^*] + v(D(\theta^*, p, w)). \quad (42)$$

Therefore, $\lim_{N_s \to \infty} D\left(\hat{\theta}, p, w\right) \epsilon D(\theta^*, p, w)$. Note that there is no requirement that $D(\theta^*, p, w)$ has to be a function; more generally, it can be a nonempty-valued correspondence.

REFERENCES CITED

Afriat, S. (1972). Efficiency Estimates of Production Functions. *International Economic Review*, 8, 568-598.

Andreoni, J. and Miller, J. (2002). Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism. *Econometrica*, Vol. 70, No. 2 (Mar., 2002), pp. 737-753.

Ballinger, T. P. and Wilcox, N. T. (1997). Decisions, Error and Heterogeneity. *The Economic Journal*, 107(443), 1090-1105.

Battalio, R.C., Kagel, J.H., Winkler, R.C., Fisher, E.B., Basmann, R.L. and Krasner, L. (1973). A Test of Consumer Demand Theory Using Observations of Individual Consumer Purchases. *Economic Inquiry* 11, no. 4: 411-428.

Braga, J., Humphrey, S. J., and Starmer, C. (2009). Market Experience Eliminates Some Anomalies—And Creates New Ones. *European Economic Review*, 53(4), 401-416.

Bordley, R. F. (1992). An Intransitive Expectations-Based Bayesian Variant of Prospect Theory. *Journal of Risk and Uncertainty*, 5(2), 127-144.

Bröder, A., and Schiffer, S. (2003). Take The Best Versus Simultaneous Feature Matching: Probabilistic Inferences From Memory and Effects of Representation Format. *Journal of Experimental Psychology: General* 132, no. 2: 277.

Camerer, C. F., Hogarth, R. M., Budescu, D. V., and Eckel, C. (1999). The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework. In *Elicitation of Preferences* (pp. 7-48). Springer Netherlands.

Conlisk, J. (1988). Optimization Cost. *Journal of Economic Behavior & Organization* 9, no. 3:

213-228.

Conlisk, J. (1996). Why Bounded Rationality?. *Journal of economic literature*, 34(2), 669-700.

Cox, J. C., and Grether, D. M. (1996). The Preference Reversal Phenomenon: Response Mode, Markets and Incentives. *Economic Theory*, 7(3), 381-405.

DeGroot, M. H. (1983). Decision Making with an Uncertain Utility Function. In *Foundations of Utility and Risk Theory with Applications* (pp. 371-384). Springer Netherlands.

Evans, G. and Ramey, G. (1998). Calculation, Adaptation and Rational Expectations. *Macroeconomic Dynamics* 2, no. 02: 156-182.

Evans, G. and Ramey, G. (1992). Expectation Calculation and Macroeconomic Dynamics. *American Economic Review*, Vol. 82, no. 1: 207-224.

Harbaugh, W.T., Krause, K. and Berry, T.R. (2001). GARP for Kids: On the Development of Rational Choice Behavior. *American Economic Review*, American Economic Association, vol. 91(5), pages 1539-1545, December 2001.

Hogarth, R.M. and Karelaia, N. (2005). Simple Models for Multiattribute Choice with Many Alternatives: When It Does and Does Not Pay to Face Trade-offs With Binary Attributes. *Management Science* 51, no. 12 (2005): 1860-1872.

Jacobs, M. (2016). Accounting for Changing Tastes: Approaches to Explaining Unstable Individual Preferences. *Review of Economics*, 67(2), 121-183.

Jun, T. (2004). A Survey on the Bandit Problem with Switching Costs. *de Economist*, 152(4), 513-541.

Loomes, G., Starmer, C., and Sugden, R. (2003). Do Anomalies Disappear in Repeated

Markets?. *The Economic Journal*, 113(486), C153-C166.

Loomes, G., Starmer, C. and Sugden, R. (2010). Preference Reversals and Disparities Between Willingness to Pay and Willingness to Accept in Repeated Markets. *Journal of Economic Psychology* 31, no. 3: 374-387.

Plott, C. R. (1996). Rational Individual Behavior in Markets and Social Choice Processes: The Discovered Preference Hypothesis. In *The Foundations of Economic Behaviour* (pp. 225-250). McMillian.

Smith, V.L. and Walker, J.M. (1993). Monetary Rewards and Decision Cost in Experimental Economics. *Economic Inquiry* 31, no. 2: 245-261.

Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 263-291.

Varian, H.R. (1991). *Goodness-of-fit for Revealed Preference Tests*. Department of Economics, University of Michigan.

Viscusi, W. K. (1989). Prospective Reference Theory: Toward an Explanation of the Paradoxes. *Journal of risk and uncertainty*, 2(3), 235-263.

Tversky, A. and Kahneman, D. (1986). Rational Choice and the Framing of Decisions. In *The Limits of Rationality* (pp. 251-278). University of Chicago Press.