Bayesian Revealed Preferences*

Daniel Martin†

April 17, 2014

Abstract

In this paper I identify a condition on stochastic choices from budget sets, called Bayesian GARP (BGARP), which characterizes the purchases of consumers who base their decisions on noisy signals of price. For a balanced panel of grocery purchases, I show that while most households fail GARP, BGARP is satisfied. In addition, I show that BGARP is restrictive for this data set.

Key Words: Stochastic Choice, Revealed Preferences, GARP, Bayes' Rule

1 Introduction

There is evidence that consumers may not know the exact prices of the goods they select for purchase (Chetty, Looney, and Kroft [2009]). Such failures could be the result of lack of attention to prices, poor memory about prices, or uncertainty about prices due to complex pricing or tax policies.

However, existing revealed preference tests, such as the Generalized Axiom of Revealed Preference (GARP), assume that decision makers are certain about prices when they select bundles (see Cherchye, Crawford, De Rock, and Vermeulen [2008] for a review). This is even true for those tests, following Varian's [1985], that allow for error in the measurement of prices.[1]

*I thank Abigail Adams, Andrew Caplin, Ian Crawford, Mark Dean, Paola Manzini, Marco Mariotti, Hiroki Nishimura, Collin Raymond, and John Quah for valuable comments.

†Paris School of Economics.

[1] The substantive difference between measurement error and imperfect perception is that measurement error does not impact the consumer's optimization problem, whereas imperfect perception does.

In this paper I produce a revealed preference test for a simple model of consumption in which consumers base their purchases on noisy signals of price. In the spirit of Afriat [1967] and Varian [1982], this test is nonparametric. In addition to no functional form assumptions about utility, there are also no functional form assumptions about signals. The only two assumptions I preserve are common to most models of demand with uncertainty: utility maximization and Bayes' rule.

I call the condition that characterizes this behavior Bayesian GARP (BGARP) because of its close connection to traditional GARP. There are two innovations in this test that allow it to account for noisy signals of price. The first is that unlike GARP, BGARP is a revealed preference condition on stochastic choices, as in Hoderlein and Stoye [2013], Caplin and Dean [2014], Caplin and Martin [2014a,b], and Manzini and Mariotti [2014]. The second is that revealed preferences are determined using an expected price that is generated by looking across all budget sets where a bundle is chosen.

I apply GARP and BGARP to a balanced panel of grocery purchases for 977 households in the Denver metropolitan area over 2 years.[2] It is in a similar setting and for a similar set of goods that Chetty, Looney, and Kroft [2009] find that consumers are uncertain about price. For these purchases, I find that almost half of households fail GARP. However, I find that BGARP is satisfied. Moreover, BGARP is restrictive for this data set: for three forms of simulated choice, BGARP is rarely satisfied.

2 Theory

In the model I characterize, consumers make standard consumption choices, but are uncertain about the prices of goods. Formally, in each decision problem consumers face a price vector $p \in P \subset \mathbb{R}^n_{++}$ and choose a bundle $x \in \mathbb{R}^n_+$. Let X be the finite set of observed choices, and P be the finite set of possible price vectors. In most theoretical work, the set P is assumed to be infinite, but it is natural to imagine that this set is large, but finite, because stores may be limited in the prices they can realistically set. In addition, there is a probability distribution over possible prices $\mu \in \Gamma = \Delta(P)$. Let μ_p be the probability of price p.[3]

[2] As in Hoderlein and Stoye [2013], stochastic choice functions are generated by looking across households. However, I attribute variation in choices to variation in perception, not variation in preferences.

[3] A suitably amended version of the characterization theorem also holds when this probability distribution is unobservable or decision makers hold subjective priors.

Because consumers are uncertain about prices when they select a bundle, it is necessary to make assumptions about the budget constraint that consumers face. I assume that they maximize subject to an expected budget constraint in each decision problem.[4] In an era of ubiquitous consumer credit, it seems reasonable to assume that most consumers do not have a hard budget constraint when they make smaller purchases, but instead maximize over a long run constraint. In section 2.1, I present a behaviorally equivalent representation with a fixed budget and an unobservable residual good.

The data set I consider is a set of stochastic choices from these decision problems, given by the function q, which is a map from possible prices P into probability distributions over observed bundles X. Let q_p^x be the probability of choosing bundle x when the price is p.

The revealed preference exercise is to find unobservables as if consumers base their purchase decisions on noisy signals of price. For this exercise, I use the imperfect perception framework of Caplin and Martin [2014a], which has three unobservables: a perception function π, a choice function C, and a utility function U.

The perception function summarizes the consumer's private information about price. It is a function,

$$\pi : P \to \Delta\left(\Gamma\right)$$

that maps prices into probability distributions with finite support over posterior beliefs $\gamma \in \Gamma$.[5] Let $\Gamma\left(\mu\right)$ be the set of possible posterior beliefs for a given set of prior beliefs μ, and γ_p be the posterior belief of price p. The perception function is assumed to satisfy Bayes' Rule, so that,

$$\gamma_p = \frac{\mu_p \pi_p\left(\gamma\right)}{\sum_{r\in P} \mu_r \pi_r\left(\gamma\right)} \text{ for all } \gamma \in \Gamma\left(\mu\right) \text{ and } p \in P.$$

Next, the choice function $C : \Gamma \to \Delta\left(X\right)$ maps posterior beliefs into bundle choices and allows for mixed strategies. Finally, the utility function $U : X \to \mathbb{R}_+$ over bundles is assumed to be nonsatiated, continuous, concave, and monotonic.

[4] Because prices are almost always treated as known to decision makers, expected budget constraints are not standard in the literature. One exception is Feenstra [1986], who considers a precautionary model with an expected budget constraint because individual must make several decisions before price uncertainty is resolved. Another is Varian [1988], who considers choices from Arrow-Debreu contingent commodities that have consumption and prices which vary by state.

[5] Kamenica and Gentzkow [2011] establish that it is possible to work directly with posterior beliefs, so it is not necessary to specify the exact signal process.

For choices from budget sets, a Bayesian expected utility maximizing (BEU) representation is defined by the following three conditions.

Definition 1 *(π, C, U) is a* ***BEU representation*** *of (X, P, μ, q) if it satisfies:*

1. *Data Matching: For all $p \in P$ and $x \in X$,*

$$q_p^x = \sum_{\gamma \in \Gamma(\pi)} \pi_p(\gamma) \, C^x(\gamma).$$

2. *Bayesian Updating: For all $p \in P$ and $\gamma \in \Gamma(\pi)$,*

$$\gamma_p = \frac{\mu_p \pi_p(\gamma)}{\sum_{r \in P} \mu_r \pi_r(\gamma)}.$$

3. *Maximization: For all $\gamma \in \Gamma(\pi)$ and $x \in X$ such that $C^x(\gamma) > 0$,*

$$U(x) \geq U(y)$$

for all $y \in X$ such that,

$$\sum_{p \in P} \gamma_p p x \geq \sum_{p \in P} \gamma_p p y.$$

There are two extreme cases nested inside of this Bayesian approach. The first is full information, in which prices are known with certainty. In this case, consumers can still mix between indifferent bundles, so the representation does not collapse to the standard deterministic one. The second is no information (beyond the prior), which is often called the "null" information structure. In this case, the perception function does not vary with price, so by Data Matching, the data should also not vary with price.

This Bayesian approach also nests a type of "salience" in which consumers only notice extreme prices (relative to the prior) as in Bordalo, Gennaioli, and Shleifer (2013). With such an information structure, the perception function for an extreme price (one that is far from the expected price) would put a large weight on posteriors that give a high probability to that extreme price. This would satisfy Bayes' rule if the overall probability of that posterior reflected the probability of encountering the extreme price. In the languages of signals, this would be as if the consumer gets a very informative signal only in the case of extreme prices.

However, some behavioral biases related to signals are not included in the Bayesian approach. For example, an optimistic consumer who overweighs positive signals about price (signals of low price) or an overconfident consumer who overweighs the precision of their signals about price may not satisfy Bayes' rule.

To characterize the BEU representation, I first define the direct revealed preference relation BR^0 as xBR^0y if,

$$\sum_{p\in P} \mu_p q_p^x px \geq \sum_{p\in P} \mu_p q_p^x py.$$

This relation says that x is revealed preferred to y if the expenditure of bundle x, *given the expected price whenever x is chosen*, is at least as high as the expenditure of bundle y, *given the expected price whenever x is chosen*. The central feature of this relation is looking across all decision problems where the bundle x is chosen.

Let BR be the transitive closure of BR^0. Further, the strict relation BP is defined as $xBPy$ if,

$$\sum_{p\in P} \mu_p q_p^x px > \sum_{p\in P} \mu_p q_p^x py.$$

These two relations can be used to produce the following revealed preference condition.

Definition 2 (X, P, μ, q) *satisfies* ***Bayesian GARP (BGARP)*** *if whenever $xBRy$, then not $yBPx$.*

The following theorem establishes that this condition is both necessary and sufficient for the decision maker to be modeled as basing their purchase decisions on noisy signals of price. The necessity of BGARP is straightforward, but the sufficiency is more surprising. If BGARP is satisfied, there always exists a signal process that can rationalize the observed data.

Theorem 1 (X, P, μ, q) *has a BEU representation if and only if it satisfies* ***Bayesian GARP (BGARP)****.*

Proof. Necessity: Suppose that (π, C, U) define a BEU of (X, P, μ, q). First, assume $xBRy$, so that,

$$\sum_{p\in P} \mu_p q_p^x px \geq ... \geq \sum_{p\in P} \mu_p q_p^x py.$$

By Data Matching, division by $\sum_{r\in P}\mu_r\pi_r(\gamma)$, and Bayesian Updating,

$$\sum_{p\in P}\sum_{\gamma\in\Gamma(\pi)}\mu_p\pi_p(\gamma)\,C^x(\gamma)px \geq \sum_{p\in P}\sum_{\gamma\in\Gamma(\pi)}\mu_p\pi_p(\gamma)\,C^x(\gamma)py.$$

For this to hold, it must be that for some γ where $C^x(\gamma)>0$,

$$\sum_{p\in P}\gamma_p px \geq \sum_{p\in P}\gamma_p py.$$

Thus, by Maximization $U(x)\geq U(y)$. Next, assume that BGARP does not hold, so that $yBPx$ or,

$$\sum_{p\in P}\mu_p q_p^y py > \sum_{p\in P}\mu_p q_p^y px.$$

Following the steps above, this implies that $U(y)>U(x)$, a contradiction.

Sufficiency: Assume that (X,P,μ,q) satisfies BGARP. The next steps will identity $\bar{\pi}$, $\bar{C}$, and $\bar{U}$ such that $(\bar{\pi},\bar{C},\bar{U})$ provides a BEU representation of (X,P,μ,q). Given $x\in X$ and $p\in P$, define the posterior $\bar{\gamma}_p^x$ by,

$$\bar{\gamma}_p^x \equiv \frac{\mu_p q_p^x}{\sum_{r\in P}\mu_r q_r^x}.$$

Because actions can have identical posteriors, it is necessary to partition the set of chosen bundles into $H\leq|X|$ sets, where all bundles in a partition have identical posteriors. Let $\bar{X}(h)$ be the h-th partition of bundles with identical posteriors, and $\bar{\gamma}(h)$ be the distinct posterior for that partition. Technically, $x,y\in\bar{X}(h)$ if and only if $\bar{\gamma}^x=\bar{\gamma}^y=\bar{\gamma}(h)$. Hence for $x,y\in\bar{X}(h)$ and $p\in P$,

$$\sum_{r\in P}\mu_r q_r^y = \mu_p q_p^y\left[\frac{\sum_{r\in P}\mu_r q_r^x}{\mu_p q_p^x}\right]. \tag{1}$$

Define the domain of the perception function as $\Gamma(\bar{\pi})=\cup_{h=1}^H\bar{\gamma}(h)$. Define the perception function itself as,

$$\bar{\pi}_p(\bar{\gamma}(h)) = \sum_{y\in\bar{X}(h)} q_p^y. \tag{2}$$

Define the choice function as,

$$\bar{C}^x(\bar{\gamma}(h)) = \begin{cases} \frac{\sum_{r\in P}\mu_r q_r^x}{\sum_{y\in\bar{X}(h)}\sum_{r\in P}\mu_r q_r^y}\in(0,1] \text{ if } x\in\bar{X}(h); \\ 0 \text{ if } x\notin\bar{X}(h). \end{cases}, \tag{3}$$

Finally, to define the utility function $\bar{U}$, first define the expected price $\overline{p}^x$ as,

$$\overline{p}^x = \sum_{p\in P}\mu_p q_p^x p.$$

Thus, xBR^0y if $\overline{p}^x x \geq \overline{p}^x y$ and $xBPy$ if $\overline{p}^x x > \overline{p}^x y$. From BGARP, we know that GARP is satisfied for the observations $\{(\overline{p}^x, x)\}_{x\in X}$. Hence, using the main results from Afriat (1967) and Varian (1985), it is possible to construct a nonsatiated, continuous, concave, and monotonic $\bar{U}$ with $\bar{U}(x) \geq \bar{U}(y)$ for all $y \in X$ such that $\overline{p}^x x \geq \overline{p}^x y$.

To verify $(\bar{\pi}, \bar{C}, \bar{U})$ provides a BEU representation of (X, P, μ, q), it is necessary to show that it satisfies Data Matching, Bayes' Rule, and Maximization. To show that Data Matching is satisfied, first note that by construction, given $p \in P$ and $x \in \bar{X}(h)$, the choice function sets $\bar{C}^x(\bar{\gamma}(h)) = 0$ unless $x \in \bar{X}(h)$. Also, from substitution of equations (2) and (3), for each $p \in P$ and $x \in \bar{X}(h)$,

$$\sum_{\gamma\in\Gamma(\bar{\pi})} \bar{\pi}_p(\gamma)\bar{C}^x(\gamma) = \sum_{h=1}^{H} \bar{\pi}_p(\bar{\gamma}(h))\bar{C}^x(\bar{\gamma}(h)) = \sum_{h=1}^{H}\left(\sum_{y\in\bar{X}(h)} q_p^y\right)\left[\frac{\sum_{r\in P}\mu_r q_r^x}{\sum_{y\in\bar{X}(h)}\sum_{r\in P}\mu_r q_r^y}\right]. \tag{4}$$

Further, from equation (1),

$$\sum_{y\in\bar{X}(h)}\sum_{r\in P}\mu_r q_r^y = \sum_{y\in\bar{X}(h)}\mu_p q_p^y\left[\frac{\sum_{r\in P}\mu_r q_r^x}{\mu_p q_p^x}\right]. \tag{5}$$

By substituting equation (5) into the denominator of equation (4),

$$\begin{aligned}\sum_{h=1}^{H}\bar{\pi}_p(\bar{\gamma}(h))\bar{C}^x(\bar{\gamma}(h)) &= \sum_{h=1}^{H}\left(\sum_{y\in\bar{X}(h)} q_p^y\right)\left[\frac{\mu_p q_p^x}{\sum_{y\in\bar{X}(h)}\mu_p q_p^y}\right] = \sum_{h=1}^{H}\frac{\sum_{y\in\bar{X}(h)}\mu_p q_p^y q_p^x}{\sum_{y\in\bar{X}(h)}\mu_p q_p^y}\\ &= q_p^x\sum_{h=1}^{H}\left[\frac{\sum_{y\in\bar{X}(h)}\mu_p q_p^y}{\sum_{y\in\bar{X}(h)}\mu_p q_p^y}\right] = q_p^x,\end{aligned}$$

which confirms Data Matching.

To show that Bayesian Updating is satisfied, note from Data Matching that for all $p \in P$, $1 \leq h \leq H$, $\bar{\gamma}(h) \in \Gamma(\pi)$, and $x \in \bar{X}(h)$,

$$\bar{\gamma}_p(h) = \frac{\mu_p q_p^x}{\sum_{r\in P}\mu_r q_r^x} = \frac{\mu_p\bar{\pi}_p(\bar{\gamma}(h))\bar{C}^x(\bar{\gamma}(h))}{\sum_{r\in P}\mu_r\bar{\pi}_r(\bar{\gamma}(h))\bar{C}^x(\bar{\gamma}(h))} = \frac{\mu_p\bar{\pi}_p(\bar{\gamma}(h))}{\sum_{r\in P}\mu_r\bar{\pi}_r(\bar{\gamma}(h))}.$$

Finally, to show that Maximization is satisfied, first note that by construction for all $x, y \in X$ and $\bar{\gamma}(h) \in \Gamma(\bar{\pi})$, if $\bar{C}^x(\bar{\gamma}(h)) > 0$ and $\bar{U}(x) \geq \bar{U}(y)$ then $\sum_{p\in P}\mu_p q_p^x px \geq \sum_{p\in P}\mu_p q_p^x py$. Then note that for each $p \in P$, $1 \leq h \leq H$ and $x \in \bar{X}(h)$,

$$\mu_p q_p^x = \bar{\gamma}_p(h)\sum_{r\in P}\mu_r q_r^x,$$

so,

$$\sum_{p\in P}\bar{\gamma}_p(h)\sum_{r\in P}\mu_r q_r^x px \geq \sum_{p\in P}\bar{\gamma}_p(h)\sum_{r\in P}\mu_r q_r^x py.$$

Division by the constant $\sum_{r \in P} \mu_r q_r^x > 0$ then yields,

$$\sum_{p \in P} \bar{\gamma}_p(h) px \geq \sum_{p \in P} \bar{\gamma}_p(h) py.$$

■

As mentioned previously, BGARP can be applied even when μ is not known or decision makers are assumed to hold subjective priors. This is accomplished by modifying the theorem to require that there exists a prior such that BGARP is satisfied. Varian [1983] suggests a similar approach in his tests of expected utility theory. Clearly, this modification makes BGARP less restrictive, but it does not lose its restrictiveness entirely.

2.1 Fixed Budget Constraint

A behaviorally equivalent representation to the BEU representation is one in which consumers (1) have a fixed budget, (2) purchase an unobservable residual good of a fixed price $q > 0$ using the remainder of their unspent budget set, and (3) get utility $\tilde{u}$, which is strictly increasing over their utility from the observable bundle x and their expectation of the residual good. Let $y \in \mathbb{R}_+$ represent their expectation of the residual good. For this alternative representation, Maximization is replaced with: there exists $y \geq 0$ and $q > 0$ such that for all $\gamma \in \Gamma(\pi)$ and $x \in X$ such that $C^x(\gamma) > 0$,

$$\max_{x \in X} \tilde{u}\left(U\left(x\right), y\right)$$

subject to,

$$\sum_{p \in P} \gamma_p px + qy \leq m.$$

The results in Polisson and Quah [2013] can be applied in the proof above to show that this representation holds if and only if BGARP is satisfied.

3 Application

I test GARP and BGARP on a balanced panel of consumption choices for 977 households in the Denver metropolitan area. This data set, which is derived from a data set used by Aguiar and Hurst [2007], contains all packaged grocery purchases over a two year period (February 1993 to

February 1995).[6] Products are aggregated into three categories (beverages, meals, and snacks), and to reduce storage concerns, consumption is aggregated at the monthly level, producing 24 observations for each household. A market level price index for each month is then constructed using observed purchases. Since the exact same bundle is very rarely purchased twice, the set of bundles X is constructed by rounding the quantity consumed to the nearest quartile. This produces 64 representative bundles. As a robustness check, I show how the results change if quantities are rounded to the nearest octile, for which there are 512 representative bundles.

Of the 977 households, only 57.5% pass GARP in the baseline case, and only 36.1% in the robustness check. One possibility is that these violations are due to uncertainty about prices. In fact, Chetty, Looney, and Kroft [2009] conduct a survey for similar goods in similar settings and find uncertainty about prices.[7] It is possible to account for this possibility by using the BGARP condition.

However, before testing BGARP, it is first necessary to estimate the stochastic choice function q and probability distribution of prices μ. The data set q is estimated by determining the distribution of bundle choices across households at each vector of prices (the price in each period). In order to pool choices across households, I make a number of additional assumptions that are frequently used in the applied literature, but are admittedly unrealistic. For instance, I assume homogeneity (of perception and preferences) and stationarity (of preferences and the price distribution).[8] Homogeneity is needed to create a distribution of choices for each price line. Stationarity is needed to pool household choices across time.[9] Finally, to produce μ, I assume that all observed prices have equal likelihood.[10]

Given the estimated q and μ, it is possible to test BGARP.[11] The results are presented in table 1. In the baseline case, BGARP is satisfied, even though there 2,068 pairs in the relation BR. On

[6] For a detailed description of the data set used in this analysis, see Dean and Martin [2014].

[7] This is especially true when tax is not included in the price, as in this area.

[8] For other data sets, some of these assumptions may not be needed. For example, in the field experiments of Choi, Kariv, Müller, and Silverman [2012] choices are taken from many budget sets in a very short amount of time and the distribution of prices is set exogenously.

[9] In many settings, there is seasonal variation in the distribution of priors, especially in grocery purchases. To make the stationarity assumption more palatable, one could test choices in each season separately.

[10] In the analysis that follows, I demonstrate the restrictiveness of BGARP first for the equiprobable prior and then for the case of subjective priors.

[11] The estimation of q and μ might introduce error into the BGARP test. See Caplin and Dean [2014] for a statistical approach to testing stochastic choice axioms when there are just two actions.

	Baseline	Robustness
Number of households	977	977
Pass GARP	57.5%	36.1%
Pass BGARP	Yes	No
Number of bundles	64	512
Number of relations	2,068	131,736
Number of violations	0	649

Table 1: Summary statistics for GARP and BGARP tests

	Type 1	Type 2	Type 3
Number of simulations	1,000	1,000	1,000
Pass BGARP	8.3%	2.7%	1.9%
Average relations	2,081	2,082	2,081
Average violations	5.2	8.5	8.0
Pass BGARP (some prior)	93.3%	57.1%	59.3%

Table 2: Restrictiveness of GARP and BGARP tests

the other hand, BGARP is not satisfied in the robustness check. However, less than 0.5% of pairs in the relation BR generate a violation. Given the robustness of these results, it is fair to ask about the restrictiveness of the condition.

To show the restrictiveness of BGARP for this data set, simulated choices can be run through the same test. In the baseline case, I find that BGARP is almost always violated for three types of simulated choices. For the first type of simulation, the stochastic choice function q_p^x is determined by taking for each price p the distribution of bundle choices for some other p (chosen with equal probability). For the second, the stochastic choice function q_p^x is determined by taking for each bundle x the joint distribution for some other x (chosen with equal probability). For the last, the stochastic choice function q_p^x is determined by taking for each price p a random distribution of bundle choices (chosen with uniform probability from all possible distributions).

The results of these tests are presented in table 2. Although the average number of violations is not large, almost all of the simulations fail BGARP.

If decision makers are allowed to hold subjective priors, the restrictiveness of BGARP can be determined by looking to see whether there exists a prior that would rationalize the simulated data. To do this, I generate 100 uniform random priors for each simulated data set and test BGARP for those priors. The results are presented in the bottom row of table 2. Except for the first type of randomization, BGARP continues to be restrictive, in that over 40% of simulations fail BGARP for all 100 of the random priors.

4 Concluding Remarks

In the language of Varian [1982], this paper develops a test of "consistency" with a model in which consumers base their purchase decisions on noisy signals of price. It is also possible to use the results of this paper to pursue "recoverability", both in terms of the utility function and the perception function. The utility function can be bounded by combining the expected prices that are "revealed" by choice with the techniques demonstrated by Varian [1982]. The perception function can be recovered using the "revealed" posteriors, as demonstrated in Caplin and Dean [2014].

Also, as mentioned in the introduction, the existing literature includes tests that assume a specific form for error in the measurement of prices. In this paper, I do not assume there is measurement error in prices, which is realistic for the detailed scanner data used in the application. Instead, I assume that the decision maker is uncertain about prices, and I place no functional form assumptions on the signals about price they receive.

References

Afriat, S. (1967), "The Construction of a Utility Function from Demand Data", *International Economic Review*, 8, 67-77.

Aguiar, M. and E. Hurst (2007), "Life-Cycle Prices and Production", *American Economic Review,* 97(5), 1533-1559.

Bordalo, P., N. Gennaioli, and A. Shleifer (2013), "Salience and Consumer Choice", *Journal of Political Economy*, 121, 803-843.

Caplin, A., and M. Dean (2014), "Revealed Preference, Rational Inattention, and Costly Information Acquisition", Unpublished.

Caplin, A., and D. Martin (2014a), "A Testable Theory of Imperfect Perception", *Economic Journal*, forthcoming.

Caplin, A., and D. Martin (2014b), "Framing Effects and Attention in Complex Choices", Unpublished.

Chetty, R., A. Looney, and K. Kroft (2009), "Salience and Taxation: Theory and Evidence", *American Economic Review*, 99(4),1145-1177.

Choi, S., S. Kariv, W. Müller, and D. Silverman (April 2011), "Who Is (More) Rational?", Unpublished.

Cherchye, L., I. Crawford, B. De Rock, and F. Vermeulen (December 2008), "The Revealed Preference Approach to Demand", Unpublished.

Dean, M., and D. Martin (2014), "Measuring Rationality with the Minimum Cost of Revealed Preference Violations", Unpublished.

Feenstra, R. (1986), "Functional equivalence between liquidity costs and the utility of money", *Journal of Monetary Economics*, 17(2), 271-291.

Kamenica, E., and M. Gentzkow. (2011), "Bayesian Persuasion", *American Economic Review*, 101(6), 2590-2615.

Hoderlein, S., and J. Stoye (2013), "Testing Stochastic Rationality and Predicting Stochastic Demand: The Case of Two Goods", Unpublished.

Manzini, P., and M. Mariotti (2014), "Stochastic choice and consideration sets", *Econometrica*, forthcoming.

Polisson, M., and J. Quah (2013), "Revealed Preference in a Discrete Consumption Space", *American Economic Journal: Microeconomics*, 5(1), 28-34.

Varian, H. (1982), "The Nonparametric Approach to Demand Analysis", *Econometrica*, 5(4).

Varian, H. (1983), "Nonparametric Tests of Models of Investor Behavior", *Journal of Financial and Quantitative Analysis*, 18(3), 269-278

Varian, H. (1985), “Nonparametric analysis of optimizing behavior with measurement error”, *Journal of Econometrics*, 30, 445–458.

Varian, H. (1988), “Estimating risk aversion from Arrow-Debreu portfolio choice”, *Econometrica*, 973-979.