Models of Revealed Preference

ABI ADAMS and IAN CRAWFORD

Abstract

Revealed preference theory is concerned with what we can learn about the process by which economic agents make decisions using simply the features of the world that we observe: choices. Different economic models place different restrictions on these choices. The revealed preference literature derives these restrictions and then puts them to use. Research effort has recently been extended far beyond the axiomatic characterization of neoclassical models of choice to consider data-consistency and preference-recoverability for a wide class of models. This essay places these recent developments in context, giving a brief introduction to the revealed preference approach before elaborating on recent research that has dramatically extended the domain and ambition of the discipline. We address the widening of the scope of revealed preference theory to new classes of models, which have introduced novel techniques, and challenges, to the discipline, before charting emerging trends within the areas of identification and power, which have arisen as revealed preference has emerged as an empirical method.

MODELS OF REVEALED PREFERENCE

One could object to the presence of an article entitled "Models of Revealed Preference" in a project entitled "Emerging Trends." The approach has a long and distinguished history, dating back to Samuelson and an initial citation date of 1938. Varian describes it as "one of the most influential ideas in economics" (2006, p. 99). The essential idea is very simple. Economic models of individual choice are based on the premise that people are endowed with coherent preferences. Yet, preferences are not directly observable; all we can observe are the choices people make. Revealed preference theory is concerned with what we can learn about the process by which economic agents make decisions using the features of the world that we observe, namely, choices. Different economic models place different restrictions on how individuals' make their choices. The revealed preference literature derives these restrictions and puts them to use.

The constructive application of revealed preference techniques is a novel development within the literature and has prompted the emergence of the

Emerging Trends in the Social and Behavioral Sciences. Edited by Robert Scott and Stephen Kosslyn.
 ISBN 978-1-118-90077-2.

discipline as an empirical method. Research effort has now extended far beyond the axiomatic characterization of theoretical models of choice to address questions of data-consistency and preference-recoverability for a wide class of models. A revealed preference approach offers the prospect of eliminating the gap between the economic theorist and the applied economist who deals with data, providing a purer way to assess the empirical performance of a theory and complementing the growing sentiment that the identification of a model's parameters is not an "all-or-nothing" affair.

This essay places these developments in context, giving a brief introduction to the revealed preference approach before elaborating on recent research that has dramatically extended the domain and ambition of the discipline. We first address the widening of the scope of revealed preference theory. The approach has now been extended to new classes of models, which has introduced novel techniques, and challenges, to the discipline. We then consider developments and trends within the areas of identification and power that have arisen as revealed preference has emerged as an empirical method.

THE ESSENTIAL IDEA

General models, which offer a simplified version of an interesting aspect of reality, have a central place in the social sciences. We live in a complicated world. Thus, there can be great power in the elimination of irrelevant detail. An appeal to Occam's razor is generally apt in our setting: all things equal, a simpler explanation should be preferred to a more complex one.

However, there is a trade-off between accuracy and parsimony. We wish to adopt the simplest model that remains capable of illuminating the situation under examination. This desire leads the social scientist to ask: Can a model accurately account for the aspect of reality under consideration? As Sen remarks: "The primary concern here is not with the relation of postulated models to the real economic world but with the accuracy of answers to well-defined questions posed with preselected assumptions which severely constrain the nature of the models" (1977, p. 322). Addressing this concern requires an empirical test of the theory: is the model consistent, or not, with data we have on the situation observed?

Yet, many models deal with motivations and mechanisms that we cannot acquire data on. This is certainly true of rational consumer choice theory, the birthplace of the revealed preference approach. The utility maximization model, at the heart of this theory, simply states that individuals choose the course of action that makes them "best off." The notion of "best off" is formalized via the specification of a "utility function," a numerical representation of an individual's preferences over different goods and actions. Formally, the

utility maximization model posits that individuals decide how much of each good to buy in order to maximise their utility function.

To fix ideas, imagine a world in which there are only two goods: whisky and bread. An individual's utility function assigns a number to any bundle of whisky and bread that we can imagine, allowing us to rank these bundles according to an individual's relative preference. For example, let the amount of whisky that an individual consumes be given by W, and the amount of bread that an individual consumes be given by B. The utility function, which describes their total amount of happiness, takes the form:

$$\text{Utility}(W, B) = f(W, B)$$

For example,

$$\text{Utility}(W, B) = W + B$$

In this example, an individual is indifferent over the split of a bundle between whisky and widgets; she just wants more of both goods.

Although the utility maximization model is a simple way to model choice behavior, it is hard to test. We do not observe an individual's utility function and, therefore, the choice that actually makes them best off. Given this, how are we to determine if rational choice theory can account for the behavior we observe in the marketplace?

Standard approaches start by placing structure on the aspects of an economic model that go unspecified by theory and are unobserved in reality, i.e. the precise structure of the utility function. Therefore, following this approach, we would assume that preferences can be described by some specific utility function (as in the example above) and then determine if the choices we observe in reality best advance these preferences.

The problem with this approach is that it is not purely a test of economic theory. It is a test of the theory plus a test of the hypothesis that the structure of the utility function takes the assumed form. To highlight the problem with this approach, imagine that we found that observed choices differed from those dictated by the hypothesized utility function. We would not be justified in discarding rational choice theory on this basis. The observed inconsistency could arise from the fact that preferences in reality do not coincide with the utility function we have assumed, rather than arising from any irrationality. Continuing to try different specifications for the utility function is not guaranteed to help either. There are an infinite number of ways to formulate preferences and thus, following this approach, we are unable to reject the model in a finite number of steps (see Popper, 1959 and Miller, 1996 for a deeper discussion).

Revealed preference theory offers an alternative method to assess how well a model performs when confronted with reality. The approach derives the

restrictions on the observable features of the world that are logically implied by the model alone. No appeal to any specific utility function is necessarily involved. If reality respects these conditions, one can conclude that it is within a model's ability to account for what we see. However, if the revealed preference restrictions are violated by observations on the world, one is justified in concluding that the model does not provide an adequate description of reality. The approach thus provides us with "effective" tests of theory.

FOUNDATIONS OF REVEALED PREFERENCE

Samuelson (1938, 1948) was the first to consider the restrictions that rationality alone imposes on observed choice behavior, concluding that: "if an individual selects batch one over batch two, he does not at the same time select two over one" (1938, p. 7). If at some point in time, an individual selects batch one over two, then they reveal their preference for one over two. Therefore, to be consistent with pure rational choice theory, that same individual cannot reveal a preference for two over one at a later date. Therefore, regardless of what form an individual's utility function does take, the model imposes a requirement of consistency on choice behavior and can be rejected if this requirement is not upheld in reality.

The foundational revealed preference literature largely concerned itself with the utility maximization model and the refinement of Samuelson's conditions. Samuelson's Weak Axiom of Revealed Preference, introduced intuitively above, was extended by Houthakker (1950) to impose consistency on choices that reveal bundles to be *indirectly revealed preferred* to one another. For example, if a rational individual chooses batch X over Y, but Y over Z, then they must choose X over Z when presented with this choice. However, few attempts were made to implement the derived tests despite their being motivated by a quest for empirical falsifiability. In the words of Mas-Colell, revealed preference represented as "foundational and purely theoretical a subject as one can find" (1982, p. 72).

This changed with the translation (Diewert, 1973) of a ground breaking, but "virtually impenetrable" (Pollak, 1990, p. 142) paper by Afriat (1967), which opened the door to greater emphasis on implementation and identification within the literature. Although the question of consistency between a model and reality is an important one, there is often a desire to go beyond this. Typically, we seek knowledge of the structure of a model's unobserved components and want to predict behavior in new situations. For example, imagine we find that an individual's choices satisfy the revealed preference conditions implied by rational choice theory. It is then interesting to ask what *kind* of preferences this individual holds. Afriat (1967) developed a method to reconstruct preferences rationalizing observed choices given finite data. He

showed how simple linear programming techniques could be used to both implement revealed preference tests of the theory and then recover individual preferences.

Once the value of Afriat's approach was recognized and made intelligible by Diewert (1973),[1] and further extended and refined by Varian (1982, 1983) the stage was set for the literature to move beyond revealed preference to "preferences revealed" and to the prediction of behavior.

Varian's (1982) contribution was central to this movement, providing a clear exposition of how a revealed preference methodology offers a viable alternative to statistical methods for applied demand analysis. However, recoverability and prediction using a revealed preference methodology rarely leads to a single answer. Unlike methods that place prior structure on the unobserved components of a model, model features are "set identified." For example, given a change in prices, revealed preference methods will not yield a unique prediction for the new choice, but will identify ranges within which the new choices must lie to remain consistent with rationality. As we will explore in further in this review, achieving tighter identification while maintaining a revealed preference approach is a key theme in modern work.

In summary, revealed preference theory has the potential to develop into a powerful tool and offers a way of reducing the distance between econometrician and theorist. The approach allows us to assess whether the simplifications a model makes are good ones by addressing the consistency of theory and reality. However, it can facilitate much more than this, providing a methodology to recover unobservable features of models and predict action in new environments.

MODERN THEMES AND EMERGING TRENDS

This section considers modern research and where we believe the future of the discipline lies. We structure our discussion thematically as the literature has expanded along a number of dimensions. We begin by considering the extension of the methodology to more complex models before addressing developments that have occurred as revealed preference has emerged as an empirical method.

Scope

The scope of the literature has extended far beyond the canonical utility maximization model. Chambers, Echenique, and Shmaya (2012) prove that there

1. This essay was born out of a referee report that he wrote concerning a paper that Afriat sent to the *Journal of Political Economy*. Diewert recognized the immense value of the paper but recommended that it be revised to make it more understandable. Afriat did not resubmit the paper to that journal.

exist revealed preference necessary and sufficient conditions, defined only in terms of observables, for any model that can be expressed as a series of universal statements. Thus, in principle, there awaits a huge class of models to be given a revealed preference treatment. However, deriving the conditions for many models requires the use of mathematical techniques that are not part of the mainstream economist's toolkit, potentially slowing the expansion of the methodology.

A large portion of early research energy was devoted to characterizing the restrictions implied by different preference structures, reflected in the general form of the utility function. The structure of utility functions is enormously important for both theoretical and applied attempts to model consumer behavior, with implications for topics as apparently diverse as choice under uncertainty and intertemporal choice[2] and the measurement of inflation.[3] Varian (1983) introduced revealed preference characterizations for some key preference structures, which made it possible for researchers to widen the span of revealed preference methods to a number of important areas. For example, building on Varian's work, Browning (1989) provided a revealed preference analysis of an intertemporal model of consumer choice, whereas Green and Srivastava (1986) considered models of choice under uncertainty.

The derivation of the testable restrictions implied by different hypotheses regarding the "separability" of goods in the utility function has led to a number of interesting developments. Separability assumptions partition the choice space into subsets that the individual evaluates independently. Separability is a matter of degree. At one end of the scale is the characteristics model of Gorman (1956) and Lancaster (1966), under which preferences are defined with respect to characteristics rather than goods, i.e. we have a desire for energy (a characteristic of a good) and so buy different food items to satisfy this. Blow, Browning, and Crawford (2008) were the first to give the model a revealed preference characterization and, in so doing, extended the scope and ambition of the discipline. The notion of latent separability was introduced in Gorman (1968) and thoroughly developed in Blundell and Robin (2000). It too is an important functional structure as it underlies Samuelson's (1956) account of the representative consumer and was characterized by Cherchye, Crawford, De Rock, and Vermeulen (2012) in their revealed preference exploration of aggregation.

Recent years have seen an extension of the revealed preference methodology to the analysis of group behavior. The key issue here is that we typically only observe the aggregate outcome of a group's decision-making process.

2. Expected utility theory and the discounted utility model are both based on additive separability.

3. The existence and uniqueness of subcost of living indices depends on weak separability and homotheticity.

Given that the aggregate behavior of a group of rational consumers is not necessarily rational (see Kirman, 1992 for example), the idea that it might be possible to test collective models using revealed preference techniques seems incredible. However, Cherchye, De Rock, and Vermeulen (2007) have shown that it is possible and have derived a number of subtle results in this area. The extension of the revealed preference methodology to collective models has opened the door to an extensive research programme that aims at opening the "black box" of the family. A wide variety of models, which make varying assumptions about the degree of cooperation and commitment between family members, have now been given the revealed preference treatment (Cherchye *et al*., 2007, 2012; Cherchye, De Rock, & Vermeulen, 2009; Adams, Cherchye, De Rock, & Verriest, 2014).

Testing the conditions associated with these models has introduced new techniques to the discipline. Typically, the revealed preference restrictions for this class of models are nonlinear and, therefore, difficult to test. A key contribution of Cherchye *et al*. (2007) was to show how to reformulate the revealed preference restrictions as a mixed integer linear programme (MILP) that can be solved using standard, robust techniques. Given a MILP formulation of the problem, one searches directly for each family member's revealed preference relation. This is a powerful recharacterization of the problem and the introduction of these new programming techniques will make it easier to operationalize tests for other complex models in the future.

A further research agenda, which has significantly widened the scope of revealed preference theory, is that considering the stochastic demand function. All the work cited above assumes that each consumer has a single set of preferences. However, simple introspection suggests that this is typically not the case. Individual preferences are often thought to embody a random component. Alternatively, an absence of information on the context of choice can cause observed choices to appear somewhat random. Bandyopadhyay, Dasgupta, and Pattanaik (1999) and McFadden (2005) derive revealed preference characterizations of models of stochastic demand. These results are important and their value has not yet been fully utilised by the empirically focused strain to the literature. Given the concern with distributions of preferences, this work provides a coherent way to extend revealed preference results to a population setting and facilitate greater use of cross-section data sets. Many of these extended revealed stochastic preference problems have not yet been studied but this area is likely to prove something of a treasure trove to economic researchers.

The gradual widening of the scope of revealed preference analysis has not yet taken in behavioral economics to any great extent. This promises to be an important area of research in the future. Behavioral economics has recently become a prominent subfield and moves somewhat away from the

neoclassical tradition of treating people as always-rational decision makers. The approach focuses on models that have more plausible sociological and psychological foundations to behavior, allowing for people to be influenced by others, make mistakes, and regret their choices. Behavioral economics promises much in terms of its potential to help us understand behavior that would otherwise prove resistant to straightforward explanation by standard rational choice models.

However, whether the important classes of models in behavioral economics will succumb to a revealed preference analysis is currently an open question. In many respects, the models do not seem likely to provide firm enough predictions, nor the necessary stability in behavior, that revealed preference methods typically make use of. If it can be shown that these models are not falsifiable using revealed preference arguments then it may undermine them—it being all very well to claim that such models can fit the data but quite another thing if it turns out that it is impossible for them to do otherwise. Alternatively, if they do turn out to be rejectable using revealed preference methods, it is likely that this will speed their acceptance amongst traditionalist nonconverts.

Identification

As should now be clear, a wide variety of economic models have been given a revealed preference treatment, and their consistency with real-life decision-making behavior ascertained. Yet, revealed preference techniques do not just provide tools to address whether a model can rationalize choices. They also provide a methodology to recover the aspects of a model that go unobserved in reality and allow us to predict behavior in new circumstances.

Recoverability has been addressed in a variety of contexts. Varian (1982) yet again provides the key foundational reference for the utility maximization model, providing a detailed exposition of how to recover a consumer's utility function and predict their theory-consistent demands at new budget regimes. Revealed preference techniques have also been used to throw light on the power relations presiding within the household and the differences in preferences that exist within the family unit. Cherchye *et al*. (2008) identify bounds on the "sharing rule," a parameter that summarizes the split of resources within household, whereas Adams *et al*. (2014) identify minimum differences in the levels of patience of different family members.

However, the methodology typically identifies *ranges* within which an answer must lie if behavior is to be characterized by a particular model. For example, a revealed preference analysis is only able to predict that, in reaction to an increase in the price of widgets by 10%, an individual will alter their behavior to demand somewhere in the range of, say, 5 and 10

widgets. The potential response is not uniquely identified; we are not able to give a unique prediction of, for example, 8 widgets.

A key problem for researchers is that survey data typically imply uninformative bounds on predicted responses and other model features. The informational content of consumption data is removed by income variation as this negates any notion of "selecting over" from the problem at hand. Recent research has explored ways of integrating more standard estimation techniques to the revealed preference methodology, in an attempt to get the best of both worlds: the functional form agnosticism of revealed preference and the tighter identification of traditional, statistical techniques.

Blundell, Browning, and Crawford (2003, 2008) exploit information on how choices vary with income to tighten the revealed preference bounds on predicted demand responses. Using their method, one first uses more traditional methods to estimate Engel curves, functions describing how choices change with income while keeping everything else constant, to, in effect, remove the impact of income variation from choice behavior. One then performs the usual revealed preference analysis on these "corrected" demands. Their method is further extended by Blundell, Browning, Cherchye, Crawford, De Rock, and Vermeulen (2012) to exploit the information that can be gleaned from the transitivity of preferences. Cherchye, De Rock, Lewbel, and Vermeulen (2012) also combine revealed preference and statistical methods in order to identify household sharing rules from cross-section data. The results presented in these works signal the enormous benefits that can flow from blending statistical econometric methods and revealed preference techniques. The bounds on demand responses calculated by Blundell *et al*. (2003, 2008, 2012) are significantly tighter, and therefore, much more useful, than those generated by a pure revealed preference approach and, without prior estimation of the household demand function, Cherchye *et al*. (2012) would not even have been able to proceed with a revealed preference analysis of the family choice problem. We consider the emergence of the trend toward greater integration of these separate branches of econometrics an important one and expect this to be an active research topic in the coming years.

We also expect a greater engagement with literature from information theory, statistical physics, and engineering. A range of techniques exist in these fields in order to select a unique answer from the feasible solution set. Infusing revealed preference with these techniques could help further narrow the distance between it and mainstream econometrics by facilitating the point identification of model features. Exactly what form the knowledge crossovers between these disciplines will take is unclear but further research in this area offers substantial promise.

Power

The emergence of revealed preference as an empirical method has prompted discussion over how one should interpret the results gained from the application of the methodology to real-world data. Revealed preference conditions provide very "crisp" tests of optimizing models: if the data satisfy the revealed preference restrictions then the data are consistent with the model and when they do not satisfy the restrictions they are not. However, it can be a little difficult to know what to make of the results of these sorts of tests. One particularly important question is whether the optimizing model in question is not, in fact, the "true" behavioral model: can we be sure that revealed preference tests are sensitive enough to detect this? In statistical hypothesis testing, this question concerns the "power" of the test: the probability of rejecting the null hypothesis when the null is not true.

Revealed preference tests are typically nonstochastic and the statistical notion of power is not applicable. Yet, there remains a need to consider the sensitivity of the testing procedure to nonmodel-consistent data-generating processes. Just as is the case with statistical power calculations, the results will depend on the alternative hypothesis considered. The difficulty is that there are many alternatives to rational choice models and no obvious benchmark. The question of the right alternative model was considered by Becker (1962), and in particular, he focused on the alternative hypothesis that consumer make random choices subject to their budget constraints. This idea was applied to revealed preference tests by Bronars (1987) who asked how likely it would be for a revealed preference test to detect a negative result when the true data-generating hypothesis was random behavior with a uniform probability distribution. More recent contributions (notably Andreoni, Gillen, & Harbaugh, 2012) consider more data-driven alternatives to random choice with uniform probability—they use the empirical distribution of choices to allow for a more realistic alternative hypothesis. Work in this area is ongoing, but the leading approach remains Bronars' (1987) method. It seems that the field has yet to settle on an agreed procedure for evaluating the sensitivity of revealed preference methods in this manner due mainly to the difficulty of finding a compelling, nonrational benchmark for the alternative data-generating process.

Given this, researchers have developed different approaches that do not require an alternative model of behavior. Andreoni *et al*. (2012), for example, also do this. One of the most interesting and intuitively appealing among these is what they term the Afriat Power Index. Given a dataset in which no revealed preference violations are detected, this measures how much the consumer's budget would have to be adjusted in order to induce a violation.

If the required adjustment is small, then the test is considered to be sensitive, if it is high then it is not.

A third type of approach, recently applied by Beatty and Crawford (2011), derives from experimental game theory and is due to Selten (1991). Under a revealed preference characterization, economic models typically delineate well-defined sets of choices that are consistent with the model of interest. It is then useful to consider the size of the theory-consistent set of behaviors relative to the size of the set of all possible behaviors. The essential idea is that if the set of observations explainable by the model is large relative to the set of possible behaviors, then simply noting that many of the observed choices are theory-consistent is not a very demanding requirement; they could hardly have done otherwise. Thus, "fit," the proportion of the sample which passes the relevant test, is not a sufficient basis for ranking the empirical performance of alternative theories. If this were the case, then a meaningless theory that ruled nothing out could not be outperformed. A better approach is to consider the trade-off between the pass rate and a measure of how sensitive the test is. We can measure the size of the theory-consistent set relative to the possibility set for the model of interest. The relative area of the empty set is zero and the relative area of all outcomes is one. We can also measure the proportion of the data that satisfies the restrictions of the model of interest (the "hit rate"). Selten (1991) argues that a good measure of the predictive success of a theory trades off these elements: the size, as it were, of the target against the number of times the data manage to hit it. In fact, he provides an axiomatic argument that the trade-off between these two elements should be the simple difference measure[4]: hit rate minus the size of the target.

Axiomatisations other than Selten's would produce different forms for the measure of the predictive success of a theory and are yet to be explored. However, the basic idea that the measure should combine both the pass rate and some measure of sensitivity remains and appears to be an important and promising area for further work.

CONCLUSION

Revealed preference has undergone something of a resurgence in recent years and is now established as an empirical discipline. This essay has charted the extension of the discipline into new classes of models and has described recent developments in the areas of identification, power, and inference, which have occurred as revealed preference has emerged as an empirical method. Among those discussed here, we consider the integration of revealed preference and traditional statistical techniques to

4. See also Beatty and Crawford (2011) for an application of this to revealed preference test

be the key emerging trend to be on the look out for. Research in this area has the potential to open up a variety of new avenues and prompt exciting developments in the areas of identification and inference.

REFERENCES

Adams, A., Cherchye, L., De Rock, B., & Verriest, E. (2014). Consume now or later? Time inconsistency, collective choice and revealed preference. *American Economic Review*, *104*(12), 4147–4183. doi:10.1257/aer.104.12.4147

Afriat, S. (1967). The construction of utility functions from expenditure data. *International Economic Review*, *8*(1), 67–77. doi:10.2307/2525382

Andreoni, J., Gillen, B., & Harbaugh, W. (2012). *The power of revealed preference tests: Ex-post evaluation of experimental design*. Pasadena, CA: California Institute for Technology.

Bandyopadhyay, T., Dasgupta, I., & Pattanaik, P. (1999). Stochastic revealed preference and the theory of demand. *Journal of Economic Theory*, *84*(1), 95–110. doi:10.1006/jeth.1998.2499

Beatty, T., & Crawford, I. (2011). How demanding is the revealed preference approach to demand? *American Economic Review*, *101*, 2782–2795. doi:10.1257/aer.101.6.2782

Becker, G. (1962). Irrational behavior and economic theory. *Journal of Political Economy*, *70*, 1–13. doi:10.1086/258584

Blow, L., Browning, M., & Crawford, I. (2008). Revealed preference analysis of characteristics models. *Review of Economic Studies*, *75*(2), 371–389. doi:10.1111/j.1467-937X.2008.00485.x

Blundell, R., Browning, M., Cherchye, L., Crawford, I., De Rock, B., & Vermeulen, F. (2012). *Sharp for SARP: Nonparametric bounds on the behavioural and welfare effects of price changes*. London, UK: Institute for Fiscal Studies.

Blundell, R., Browning, M., & Crawford, I. (2003). Nonparametric Engel curves and revealed preference. *Econometrica*, *71*(1), 205–240. doi:10.1111/1468-0262.00394

Blundell, R., Browning, M., & Crawford, I. (2008). Best nonparametric bounds on demand responses. *Econometrica*, *76*(6), 1227–1262.

Blundell, R., & Robin, J. (2000). Latent separability: Grouping goods without weak separability. *Econometrica*, *68*(1), 53–84. doi:10.1111/1468-0262.00093

Bronars, S. (1987). The power of nonparametric tests of preference maximization. *Econometrica*, *55*(3), 693–698. doi:10.2307/1913608

Browning, M. (1989). A nonparametric test of the life-cycle rational expectations hypothesis. *International Economic Review*, *30*(4), 979–992. doi:10.2307/2526762

Chambers, C., Echenique, F., & Shmaya, E. (2012). *General revealed preference theory*. Pasadena, CA: California Institute for Technology.

Cherchye, L., Crawford, I. D., Rock, B., & Vermeulen, F. (2012). *Aggregation without the aggravation? A nonparametric analysis of representative consumers*. Leuven, Belgium: Katholieke Universiteit Leuven.

Cherchye, L., De Rock, B., Lewbel, A., & Vermeulen, F. (2012). *Sharing rule identification for general collective consumption models*. Oxford, UK: Department of Economics.

Cherchye, L., De Rock, B., & Vermeulen, F. (2007). The collective model of household consumption: A nonparametric characterization. *Econometrica*, *75*(2), 553–574. doi:10.1111/j.1468-0262.2006.00757.x

Cherchye, L., De Rock, B., Sabbe, J., & Vermeulen, F. (2008). Nonparametric tests of collectively rational consumption behavior: an integer programming procedure. *Journal of Econometrics*, *147*, 258–265.

Cherchye, L., De Rock, B., & Vermeulen, F. (2009). Opening the black box of intrahousehold decision making. *Journal of Political Economy*, *117*(6), 1074–1104. doi:10.1086/649563

Diewert, E. W. (1973). Afriat and revealed preference theory. *The Review of Economic Studies*, *40*(3), 419–425. doi:10.2307/2296461

Gorman, W. (1968). The structure of utility functions. *Review of Economic Studies*, *35*(4), 367–380. doi:10.2307/2296766

Gorman, W. M. (1956). A possible procedure for analysing quality differentials in the eggs market. *Review of Economic Studies*, *47*(5), 843–856.

Green, R., & Srivastava, S. (1986). Expected utility maximisation and demand behavior. *Journal of Economic Theory*, *38*(2), 313–323. doi:10.1016/0022-0531(86)90121-3

Houthakker, H. (1950). Revealed preferences and the utility function. *Economica*, *17*(66), 159–174. doi:10.2307/2549382

Kirman, A. (1992). Whom or what does the representative individual represent? *Journal of Economic Perspectives*, *6*(2), 117–136. doi:10.1257/jep.6.2.117

Lancaster, K. (1966). A new approach to consumer theory. *Journal of Political Economy*, *74*(2), 132–157. doi:10.1086/259131

Mas-Colell, A. (1982). Revealed preference after Samuelson. In G. Feiwei (Ed.), *Samuelson and neoclassical economics* (pp. 72–82). Amsterdam, The Netherlands: Kluwer.

McFadden, D. (2005). Revealed stochastic preference: A synthesis. *Economic Theory*, *26*(2), 245–264. doi:10.1007/s00199-004-0495-3

Miller, D. (1996). What use is empirical confirmation? *Economics and Philosophy*, *12*(2), 197–206. doi:10.1017/S0266267100004168

Pollak, R. (1990). Distinguished fellow: Houthakker's contribution to economics. *Journal of Economic Perspectives*, *4*(2), 141–156. doi:10.1257/jep.4.2.141

Popper, K. (1959). *The logic of scientific discovery*. London, England: Hutchinson and Company.

Samuelson, P. (1938). A note on the pure theory of consumer behavior. *Economica*, *5*(17), 61–71. doi:10.2307/2548836

Samuelson, P. (1948). Consumption theory in terms of revealed preference. *Economica*, *15*(60), 243–253. doi:10.2307/2549561

Samuelson, P. (1956). Social indifference curves. *Quarterly Journal of Economics*, *70*(1), 1–22. doi:10.2307/1884510

Selten, R. (1991). Properties of a measure of predictive success. *Mathematical Social Sciences*, *21*(2), 153–167. doi:10.1016/0165-4896(91)90076-4

Sen, A. (1977). Rational fools: A critique of the behavioral foundations of economic theory. *Philosophy and Public affairs*, *6*(4), 317–344.
Varian, H. (1982). The nonparametric approach to demand analysis. *Econometrica*, *50*(4), 945–973. doi:10.2307/1912771
Varian, H. (1983). Non-parametric tests of models of consumer behaviour. *Review of Economic Studies*, *50*(1), 99–110. doi:10.2307/2296957
Varian, H. (2006). Revealed preference. In M. Szenberg & A. Gottesman (Eds.), *Samuelsonian economics and the twenty-first century* (pp. 99–115). Oxford, England: Oxford University Press.

ABI ADAMS SHORT BIOGRAPHY

Abi Adams is a Junior Research Fellow at Merton College, University of Oxford and a Research Associate of the Institute for Fiscal Studies. She works on revealed preference analysis of individual and household behaviour.

IAN CRAWFORD SHORT BIOGRAPHY

Ian Crawford is a Professor of Economics at the University of Oxford and a Research Fellow of the Institute of Fiscal Studies, where he was previously Deputy Director. He works on nonparametric/revealed preference analysis of consumer behavior.

RELATED ESSAYS

Coevolution of Decision-Making and Social Environments *(Sociology)*, Elizabeth Bruch *et al*.
Mental Models *(Psychology)*, Ruth M. J. Byrne
Choice Architecture *(Psychology)*, Adrian R. Camilleri and Rick P. Larrick
Emerging Trends: Asset Pricing *(Economics)*, John Y. Campbell
Culture and Cognition *(Sociology)*, Karen A. Cerulo
The Inherence Heuristic: Generating Everyday Explanations *(Psychology)*, Andrei Cimpian
Language and Thought *(Psychology)*, Susan Goldin-Meadow
Behavioral Economics *(Sociology)*, Guy Hochman and Dan Ariely
Emotion and Decision Making *(Psychology)*, Jeff R. Huntsinger and Cara Ray
Genetic Foundations of Attitude Formation *(Political Science)*, Christian Kandler *et al*.
Cultural Neuroscience: Connecting Culture, Brain, and Genes *(Psychology)*, Shinobu Kitayama and Sarah Huff
From Individual Rationality to Socially Embedded Self-Regulation *(Sociology)*, Siegwart Lindenberg

Against Game Theory *(Political Science)*, Gale M. Lucas *et al*.
Implicit Attitude Measures *(Psychology)*, Gregory Mitchell and Philip E. Tetlock
Heuristics: Tools for an Uncertain World *(Psychology)*, Hansjörg Neth and Gerd Gigerenzer
Embodied Knowledge *(Psychology)*, Diane Pecher and René Zeelenberg
Economics and Culture *(Economics)*, Gérard Roland
Event Processing as an Executive Enterprise *(Psychology)*, Robbie A. Ross and Dare A. Baldwin