American Economic Review: Papers & Proceedings 2009, 99:2, 356–362
http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.2.356

NEW EMPIRICAL APPROACHES TO DECISION MAKING UNDER UNCERTAINTY[†]

Identifying Preferences under Risk from Discrete Choices

By Pierre-André Chiappori, Amit Gandhi, Bernard Salanié, and Francois Salanié*

When studying consumption choices, economists have often relied on the abstraction of a representative agent. Such an agent can indeed be shown to exist and to replicate the aggregate consumers' demand, but only under very strong (and actually quite unrealistic) assumptions (Alan P. Kirman 1992). There was also a justifiable reluctance to introduce heterogeneous preferences, as such a step might seem ad hoc when trying to explain different consumption behaviors. The rise of empirical studies based on microdata has opened new perspectives. The microeconomic importance of uninsurable risks is now recognized, and threatens the foundations of the representative agent hypothesis often used in macroeconomics. The continuing controversies surrounding the question of individual attitudes toward risk have motivated many empirical studies and observations; most of them find a bewildering diversity of individual preferences (Robert B. Barsky et al. 1997; Alma Cohen and Liran Einav 2002; Luigi Guiso and Monica Paiella 2008; Syngjoo Choi et al. 2007; Chiappori and Paiella 2007). Clearly, the identifiability of the heterogeneous distribution of preferences becomes a crucial issue in this perspective.

This paper proposes conditions under which heterogeneous individual attitudes toward risk can be nonparametrically identified from individual- or market-level data on the choices made by agents over risk prospects. Our main result establishes that given data that is usually available (essentially market shares of the different risky prospects present within a market, plus the realizations of the final outcomes of agents), the analyst can recover the whole distribution of individual preferences so long as preferences can be indexed by a one-dimensional parameter that satisfies a fairly weak single-crossing condition. We then discuss several applications of our general methodology.

I. Theory

We consider an economic situation in which a population of privately informed agents faces a finite menu of risky prospects. In this section, we assume that from the perspective of the econometrician, this population appears homogeneous. Equivalently, the econometrician has data on a bigger population, including some variables X describing the population under study, and the subgroup we consider here appears homogeneous because we implicitly control for the values of X.

A. *The Economic Model*

Agents are heterogeneous; each agent is characterized by some parameter θ, his type, which is his private information. Within a market, all agents are proposed the same finite menu $\mathcal{C} = \{c_1, \ldots, c_n\}$. Each prospect c_i in this set can be thought of as a lottery (or a "bet"), which in

[†]*Discussants:* Shachar Kariv, University of California-Berkeley; Matt Shum, Johns Hopkins University; Victor Matheson, College of the Holy Cross.

*Chiappori: Economics Department, Columbia University, 1009A International Affairs Building, MC 3308, 420 West 118th Street, New York, NY 10027 (e-mail: pc2167@columbia.edu); Gandhi: Department of Economics, University of Wisconsin-Madison, 7442 Social Science Building, 1180 Observatory Dr., Madison, WI 53706 (e-mail: agandhi@ssc.wisc.edu); B. Salanié: Economics Department, Columbia University, International Affairs Building, MC 3308, 420 West 118th Street, New York, NY 10027 (e-mail: bs2237@columbia.edu); F. Salanié: Toulouse School of Economics and (LERNA/INRA), 21 Allée de Brienne, 31000 Toulouse, France (e-mail: salanie@toulouse.inra.fr). F. Salanié gratefully acknowledges financial support from Fondation du risque/SCOR, through its "Marchés des risques et création de valeur" grant.

practice may be an insurance contract, a job proposal, or any choice involving risky outcomes. There is a set of consequences $\mathcal{M} \subset \mathbb{R}^k$, a typical element of which is a real vector $\mathbf{m}$; lotteries on this set are characterized by their c.d.f. F. In the most typical case, the consequence $\mathbf{m}$ is simply a scalar $m \in \mathbb{R}$ measured in monetary units. By assumption, agents care only about consequences, so that we endow each agent with preferences represented by the utility functional $W(F, \theta)$. This functional may be expected utility:

$$W(F, \theta) = \int u(m, \theta) \; dF(m),$$

but we do not require this assumption; in fact, we will put much stress on identifying possible departures from expected utility in the population of agents.

When faced with the menu $\mathcal{C}$ of prospects, agent θ associates a distribution $F(\cdot \,|\, c_i, \theta)$ to each feasible choice c_i, $i = 1, \ldots, n$; note, in particular, that the (perceived) probability distribution induced by a given action may depend on the agent's type, as would be the case, for instance, when riskiness is agent-specific. Given the agent's preferences W, she will choose the option that gives the highest value to $V(c_i, \theta) = W(F(\cdot \,|\, c_i, \theta), \theta)$, which yields a choice function $C(\theta, \mathcal{C})$.

This model of choice is very general, indeed so much so that we cannot make much progress without specializing it. In this paper we discuss the one-dimensional case with single-crossing. So we take θ to be a scalar, whose distribution we normalize to be uniform over $[0, 1]$[1]; and we assume the existence of a complete ordering $\prec$ of prospects such that:

Single Crossing Assumption (SC): For any $\theta < \theta'$ and any pair of prospects $c \prec c'$,

$$V(c, \theta) \leq V(c', \theta)$$

implies

$$V(c, \theta') < V(c', \theta').$$

The SC Assumption is standard. It states, in our context, that one can rank the prospects of the menu $\mathcal{C}$ in such a way that agents with a higher type prefer prospects with a higher rank. One may, for instance, identify θ to a risk preference parameter, and rank prospects by increasing risk. Then our assumption requires, first, that the diversity of preferences be accounted for by a single parameter, and, second, that they can be ranked globally by decreasing dislike for risk. But our setting can accommodate other interpretations. For instance, in a nonexpected utility framework, θ could index some deformation of probability; alternatively, we may think of a model entailing robust control à la Lars Hansen and Thomas Sargent (2007), and θ could then indicate the cost the agent put on probability distortions.[2]

Without loss of generally, we may assume that the ranking characterized in the SC Assumption is the natural one—i.e., $c_i \prec c_j$ if and only if $i < j$. Note also that the single-crossing condition bears on a complex object, as the type θ might enter both preferences and distributions of consequences. That is, θ enters both $W(F, \theta)$ and $F(\cdot \,|\, c_i, \theta)$. We will discuss subcases below.

B. *The Data*

Ideally, we would like to recover the distribution of preferences $W(F, \theta)$ over the set of types. Obviously this requires us to assume at a minimum that we have stable preferences, i.e., we observe the same population distribution making choices over different menus of contracts. We can think of each such menu $\mathcal{C}$ as defining a market, and the basic idea we pursue below is how cross market variations in the choice set facing agents identifies $W(F, \theta)$ nonparametrically. For each market $\mathcal{C}$ in the support of the data generating process, the econometrician observes:

- The market share $s_i(\mathcal{C})$ of each prospect, which is given by

$$s_i(\mathcal{C}) = \Pr\,[C(\theta, \mathcal{C}) = i];$$

- The empirical distribution of m, for each prospect:

$$G_i(m \,|\, \mathcal{C}) = E[F(m \,|\, c_i, \theta) \,|\, i = C(\theta, \mathcal{C})].$$

[1] If θ is distributed according to some c.d.f. Φ, then $\overline{\theta} = \Phi(\theta)$ is uniformly distributed on $[0, 1]$, and we can define $\overline{V}(c, \overline{\theta}) = V(c, \overline{\Phi}^{-1}(\overline{\theta}))$.

[2] We could also forgo preference heterogeneity and focus on heterogeneity of risk, for instance in a model of insurance under adverse selection.

The assumption on the observability of market shares does not call for much discussion. On the other hand, a couple of remarks are in order here. First, the same population must be observed in different choice situations; this clearly requires that there exist an unmodelled but exogenous shock on the supply side that varies the menu of contracts across markets. Think, for example, of the population of bettors in a horse race on Sundays, or of young male drivers in rural regions shopping for car insurance. In both cases it may be reasonable to assume that the distribution of preferences remains the same on each Sunday, or across regions. In the first case menus differ because different horses in different races have different winning probabilities. In the second case menus may vary from one region to another because loading factors differ, reflecting, for instance, various levels of competition across insurers by regions. Note that the econometrician does not need to model the precise reason these menus differ—as long as she can make the case that the underlying process is independent of preferences of the riskiness of agents at the same time.

Second, the equation on empirical distributions above requires that agents have correct beliefs when making their choices. For instance, if we analyze insurance contracts, we are requiring that an agent make decisions based on the true distribution of claims conditional on (c_i, θ). Such an assumption cannot be avoided: barring data on beliefs, most datasets could be rationalized by assigning strange beliefs to agents.

C. *Identification*

The beauty of single crossing, as shown in hundreds of papers since James A. Mirrlees (1971), is that it yields market segmentation: given a menu $\mathcal{C}$, agents with a higher θ will choose prospects with a higher rank. Then we can define a sequence of numbers $\theta_i(\mathcal{C})$ for $i = 1,\ldots,n-1$ such that:

- $\theta_0(\mathcal{C}) = 0$;
- $\theta_n(\mathcal{C}) = 1$;
- Prospect $i \in \{1,\ldots,n\}$ is chosen by the set of types

$$\theta \in [\theta_{i-1}(\mathcal{C}), \theta_i(\mathcal{C})].$$

In particular, the type $\theta_i(\mathcal{C})$ is defined by the property that such agents are just indifferent between choices c_i and c_{i+1}. By single crossing, all agents with a type above $\theta_i(\mathcal{C})$ strictly prefer c_{i+1} over c_i, while types below $\theta_i(\mathcal{C})$ have the opposite preference.

Since we normalized the distribution of types to be uniform in $[0, 1]$, the market share of prospect i among agents is

$$s_i(\mathcal{C}) = \theta_i(\mathcal{C}) - \theta_{i-1}(\mathcal{C}),$$

which can be rewritten as

$$\theta_i(\mathcal{C}) = \sum_{j \le i} s_i(\mathcal{C}).$$

We may thus assume that for each menu $\mathcal{C}$, the econometrician knows the indices of types indifferent between two consecutive prospects.

Similarly, we can rewrite the estimated distribution of consequences as

$$G_i(m \mid \mathcal{C}) = \frac{1}{s_i(\mathcal{C})} \int_{\theta_{i-1}(\mathcal{C})}^{\theta_i(\mathcal{C})} F(m \mid c_i, \theta)\, d\theta.$$

We may now provide a first intuition of the identification strategy. Consider the subset of markets whose menu includes the prospect c_i. For each such market, we know the values of the bounds in the integral and the function to be integrated. If there is "enough variation" in these bounds, in the sense that the support of the distribution of $(\theta_{i-1}(\mathcal{C}), \theta_i(\mathcal{C}))$ is "sufficiently large" (for example, $[0, 1] \times [0, 1]$ would be sufficient but certainly not necessary), we can recover from the data the value of the function $F(m \mid c_i, \theta)$, for any m and any θ such that there exists a market in which θ indeed chooses c_i.

Finally, we know that for any menu $\mathcal{C}$ and for any i we have

$$W[F(\cdot \mid c_i, \theta_i(\mathcal{C})), \theta_i(\mathcal{C})] = W[F(\cdot \mid c_{i+1}, \theta_i(\mathcal{C})), \theta_i(\mathcal{C})].$$

Now fix a value for both the agent's type and the prospect—say, $\overline{\theta}$ and $\overline{c}$. Consider the subset of markets for which $\overline{\theta}$ is indifferent between $\overline{c}$ and an adjacent prospect—technically, the set of menus $\mathcal{C}$ such that $c_i = \overline{c}$ and either $\theta_i(\mathcal{C})$ or $\theta_{i-1}(\mathcal{C})$ equals $\overline{\theta}$. Intuitively, each such market provides a point on the indifference curve of θ that goes through $F(\cdot \mid c, \theta)$. Provided such observations display enough variation, one can identify the indifference curves of agent θ, and thus her preferences up to an increasing transformation.

Our present goal is to introduce a general methodology. Consequently, we shall not attempt a general statement of the necessary conditions for identification; it is much easier to illustrate them on specific examples. Suffices to say that, in many cases, the corresponding assumptions are fairly simple. Moreover, the Single-Crossing Assumption, as well as other features of our model, can be tested.

II. Applications

We now turn to a few selected examples.

A. *Parimutuel Horse Betting*

Let us start with parimutuel betting on horse races, which we study in much more detail in a companion paper (Chiappori et al. 2008). With betting, we typically do not have bettor-level data; on the other hand, we can observe odds of each horse in a very large number of races, along with the results of the race. The former gives us information on the preferences of bettors, as odds result from a simple market mechanism; the latter informs us on the distribution of consequences of bets. To make this clear, consider the market for "win bets" in a race with n horses. Each bettor chooses on which horse he wants to bet a fixed amount. The odds R_i is the net return from a winning bet: if a bettor bets \$1 on horse i, his final wealth is increased by R_i if i wins, and reduced by \$1 if i loses.

Bets on races may be organized by bookmakers, as in Bruno Jullien and Salanié (2000); but parimutuel data are more useful for our purposes, as odds then mechanically reflect the market shares of horses. In parimutuel races, the money bet is placed in a pool; the organizers of the race (and/or the taxman) receive a share t known as the "take," and the lucky bettors share what is left in the pool. Hence, market shares (s_i) and odds (R_i) are related by the simple equality

$$R_i + 1 = \frac{1-t}{s_i},$$

and the value of t can be deduced from the fact that $\sum_i s_i = 1$.

Observe also that if we properly control for observables, the winning probability of horse i in a race with odds $(R_1, \ldots, R_n)$ can be a function only of these odds, since bettors care only about consequences[3]. Hence, this probability $p_i(R_1, \ldots, R_n)$ can be computed as the empirical frequency of the event "horse i wins," on the subset of races with such odds (assuming a large enough dataset).

We can now apply our theoretical work to horse races. The first step is to partition the set of races into different subsets; within each subset (e.g., races held on weekdays on a particular track), we assume that the population of bettors is the same in each race.[4]

We then have to specify what information agents have regarding each race; this is not an easy task, as information on horses is typically dispersed across betters, and odds are not known until the beginning of the race since they result from the amounts bet. We rely on a result by Amit K. Gandhi (2008), which shows that if the pooled information of the agents pins down the actual probabilities, then there is a unique and fully revealing rational expectations equilibrium of the underlying contingent claims market. Assuming that this equilibrium prevails, we obtain that beliefs will be correct in equilibrium, as if all betters knew the probabilities $p_i(R_1, \ldots, R_n)$ and the odds $(R_1, \ldots, R_n)$. Using the notation from the previous section, we see that the menu offered to bettors is simply $\mathcal{C} = \{(p_1, R_1), \ldots, (p_n, R_n)\}$ (each horse c_i being characterized by a probability of winning and a payment in case of success). The distribution of the final monetary outcome, $F(m \mid c, \theta)$, is simply a binomial distribution, in general independent of θ[5]: given i and $\mathcal{C}$, $m = R_i$ with probability p_i, and $m = -1$ with probability $(1 - p_i)$.

As agents care only for consequences, we can thus write

$$V(c_i, \theta) = V(R_i, p_i, \theta).$$

[3] This does not need to be true if odds are set by oligopolistic bookmakers, since they may set odds that depend on other considerations, such as the existence of other races. Under parimutuel betting, such a problem does not arise.

[4] We leave self-selection of bettors for further work.

[5] Note, however, that the initial wealth of an agent can be part of his type θ.

Chiappori et al. (2008) then use a simple condition on V to ensure that single crossing holds: agents with higher types θ must dislike risk less (in a precise sense that fits this general utility theoretic framework). As a consequence, one can identify the whole distribution of preferences.

This application has several features that greatly simplify matters. First, the particular rules of parimutuel betting limit the need for data; odds allow one to directly recover market shares and the take. Second, barring probability deformation, the distribution of consequences of a bet on a horse does not depend on the type θ of the bettor, but only on the odds of the horse R_i and on the probability p_i that the horse wins. As a consequence, it is very easy to identify the probability distribution on consequences from the data. Note, however, that more complex models can also be estimated in this framework; for instance, the decision criteria may entail nonexpected utility criteria that are not linear in probability, in which case the perceived distribution may depend on θ. Again, the reader is referred to Chiappori et al. (2008) for precise results and proofs.

B. *Compensating Wage Differences*

A very similar framework can be applied to the analysis of the statistical value of human life. Book I of the *Wealth of Nations* already emphasized that [...] the wages of labour vary with the ease or hardship, the cleanliness or dirtiness, the honourableness or dishonourableness, of the employment."

A fortiori, Sherwin Rosen (1988) emphasized that the wage differential needed to compensate a higher probability of accidental death allows one to estimate the value of life. It has been extensively used subsequently (see, for instance, Joseph E. Aldy and W. Kip Viscusi 2003). In all studies, however, homogeneity is a crucial prerequisite; i.e., identification relies on the rather doubtful assumption that all workers have the same monetary valuation of their life.

Our approach suggests that this assumption is simply not needed. Specifically, assume that agents are faced with a menu of jobs, each of which is defined by a wage and some parameters related to its riskiness; to keep things simple, let us assume that the risk is a fixed probability of a lethal accident, and that the latter can be estimated from data on the distribution of accidents in each job i. Also, let θ index the value that any given agent assigns to his life; that is, the value of agent θ's life is some $v(\theta)$ with $v(0)$ corresponding to a minimum level (say zero) and $v(1)$ to some upper bound. Here, a job $c_i = (w_i, p_i)$ is fully defined by a wage w_i and a death probability p_i. The resulting probability distribution $F(m \mid c, \theta)$ is again a binomial distribution. An agent choosing the job i in the available menu $\mathcal{C}$ receives $m = w_i$ with probability $1 - p_i$, and $m = -\theta$ with probability $(1 - p_i)$. The wage rate w_i of a job is typically observable, and the probability of death parameter p_i can readily be identified from available data on average casualty rate per job; also, the natural ranking is by increasing risk (and wage).

Considering now the estimation of preferences, we may start with the simplest case of risk-neutral agents (which most of the existing literature focuses on); then $\theta_i(\mathcal{C})$ is defined by

$$(1 - p_i)w_i + p_i v(\theta_i(\mathcal{C})) = (1 - p_{i+1})w_{i+1} + p_{i+1}\, v(\theta_i(\mathcal{C}));$$

therefore

$$v(\theta_i(\mathcal{C})) = \frac{(1 - p_i)w_i - (1 - p_{i+1})w_{i+1}}{p_{i+1} - p_i}$$

and the function v is exactly identified. Note that the single-crossing condition is obviously satisfied in that case.

We may, however, go one step further, by assuming first that individuals are characterized by some common utility u that is concave in monetary rewards. The equation becomes

$$(1 - p_i)u(w_i) + p_i v(\theta_i(\mathcal{C})) = (1 - p_{i+1})u(w_{i+1}) + p_{i+1}\, v(\theta_i(\mathcal{C})).$$

Now, assume we can observe different menus in which the same agent $\overline{\theta}$ is pivotal between the same job $(\overline{p}, \overline{w})$ and different alternatives (p, w). Then, it must be the case that for all these alternatives, the expected utility $(1 - p)u(w) + pv(\overline{\theta})$ is constant (this is the "indifference curve" mentioned above). In practice, if these various alternatives describe a wage/probability relationship of the form $p(w, \overline{\theta})$, then

$$v(\overline{\theta}) = \frac{(1 - \overline{p})u(\overline{w}) - (1 - p(w, \overline{\theta}))u(w)}{p(w, \overline{\theta}) - \overline{p}}$$

and if we normalize $u(\overline{w})$ to be zero, we have that

$$\frac{u(w)}{v(\overline{\theta})} = \frac{\overline{p} - p(w,\overline{\theta})}{1 - p(w,\overline{\theta})}$$

which identifies both u and v up to a common multiplicative constant and also generates strong overidentifying restrictions.

Finally, we may also introduce heterogeneous utilities by allowing u to depend on θ, although the single-crossing condition will then have to be tested ex post.

The intuition of this example is clear. While the existing literature tends to pick up arbitrarily one type of risky job and estimate the statistical value of life from this unique source under the assumption of identical consumers, we suggest that additional mileage could be obtained form the respective "market shares" of various dangerous professions.

C. *Insurance and Portfolio Choice*

As a third possible application, consider the choice of an insurance contract—a topic already studied by Chiappori et al. (2005) and Cohen and Einav (2002), among others. Here, a contract is, in the simplest case, defined by a pair consisting of a premium and a fixed deductible. The distribution induced by a particular choice may be more complex than before, since it need not be binomial: losses incurred by the agent can in principle take any value below the deductible–although several applied papers actually disregard this issue, certainly an acceptable simplification when the deductible is low. In any case, it can, as before, be identified from available data on accident realizations (and possibly conditional losses); this has actually been done in several papers.

The next step is to study individual demand for specific contracts. This task, again, has been performed in a few existing papers, although our approach introduces an important innovation. Specifically, existing work heavily relies on parametric assumptions on preferences (e.g., CARA expected utility). However, it follows from the previous discussions that the shape of individual preferences as a function of the heterogeneity parameter θ can be nonparametrically identified by a procedure close to (although more general than) the one just described. Again, the crucial requirement is that similar agents must be faced with different menus of contracts, which implies a source of contract variation that is orthogonal to individual preferences. Fortunately, the existence of such variations seems well established empirically.

In the literature, two types of heterogeneity have been considered so far. One is heterogeneity in risks. In such frameworks, agents have identical preferences but different accident probabilities, and the choice of an insurance contract partly reflects information about idiosyncratic risk. The vast empirical literature on asymmetric information in insurance, starting with Chiappori and Salanié (2000), belongs to this class. Unlike the previous examples, in these models the distribution of probability over outcomes, $F(\cdot \mid c_i, \theta)$, does depend on θ (which is correlated with risk) as well as c_i (say, because of ex ante or ex post moral hazard). Estimating this relationship is often the difficult part of the empirical analysis; clearly, the availability of different menus is crucial here. On the other hand, once the function F has been recovered, estimating the (common) utility from contract choices is usually simpler. Alternatively, one can consider a model in which agents have homogeneous risks (at least conditional on the information available to the insurer),[6] but differ in their preferences. Then F does not depend on θ (which is now interpreted, for instance, as a risk aversion parameter), and can readily be recovered from data on claims.[7] Heterogeneity is mainly reflected in contract choices; the single crossing condition typically expresses that more risk-averse agents always prefer more comprehensive coverage, a property most models will satisfy. Again, one can then try to estimate nonparametrically the (heterogeneous) shape of individual preferences from available data.

An obvious limitation of our approach, in the case of insurance contracts, is the one-dimensionality of the heterogeneity parameter θ. Strictly speaking, it requires either that individuals differ only in their risk or their risk aversion (but not both); or, if both dimensions

[6] Regarding automobile insurance, for instance, many insurers tend to believe that, if anything, they have a much better knowledge of a client's risk than the client himself.

[7] However, the decision to file a claim may be endogenous, and possibly correlated with risk aversion; see Chiappori et al. (2006) for a discussion.

are allowed to vary in the population, that there exist a perfect correlation between them, so that the resulting heterogeneity can be represented by one index only. While this setting is very restrictive, it is still possible to find interesting theoretical frameworks that are compatible with it. For instance, in Jullien, Salanié, and Salanié (2007), agents can exert an effort to reduce their accident probability, and heterogeneous attitudes to risk result in different effort levels, and therefore different realized risks. In that context, a (one-dimensional) distribution of risk aversion generates heterogeneity in both risk and contract choices conditional on risk.

III. Conclusion

In this paper, we propose a general methodology for estimating heterogeneous preferences in situations of discrete choice under risk. Our key requirements are (i) that the existing heterogeneity can be summarized in a one-dimensional parameter, and (ii) that there exists some exogenous source of variation in the menus faced by the agents. Neither of these assumptions is innocuous. It should, however, be stressed that they typically allow a nonparametric identification of the distribution of preferences (or risks); moreover, they are empirically testable, and actually generate strong overidentifying restrictions. Empirical work currently in progress should give us a better understanding of their relevance.

REFERENCES

Aldy, Joesph E., and W. Kip Viscusi. 2003. "The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World." *Journal of Risk and Uncertainty*, 27(1): 5–76.

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Mathew D. Shaprio. 1997. "Preference Parameters and Behavioral Heterogeneity: An Observational Approach in the Health and Retirement Study." *Quarterly Journal of Economics*, 112(2): 537–79.

Chiappori, Pierre André, Amit Gandhi, Bernard Salanié, and Francois Salanié. 2008. "You Are What You Bet: Eliciting Risk Attitudes from Horse Racing." Unpublished.

Chiappori, Pierre André, Bruno Jullien, Bernard Salanié, and Francois Salanié. 2006. "Asymmetric Information in Insurance: General Testable Implications." *The Rand Journal of Economics*, 37(4): 783–98.

Chiappori, Pierre André, and Monica Paiella. 2007. "Relative Risk-Aversion is Constant: Evidence from Panel Data." Unpublished.

Chiappori, Pierre André, and Bernard Salanié. 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, 108(1): 56–78.

Choi, Syngjoo, Raymond Fisman, Douglas Gale, and Sarchar Kariv. 2007. "Consistency and Heterogeneity of Individual Behavior under Uncertainty." *American Economic Review*, 97(5): 1921–38.

Cohen, Alma, and Liran Einav. 2007. "Estimating Risk Preferences from Deductible Choice." *American Economic Review*, 97(3): 745–88.

Guiso, Luigi and Monica Paiella. 2008. "Risk-aversion, Wealth and Background Risk." *Journal of the European Economic Association*, 6(6): 1109–50.

Gandhi, Amit K. 2008. "Rational Expectations at the Racetrack: Testing Expected Utility Using Prediction Market Prices." Unpublished.

Hansen, Lars, and Thomas Sargent. 2007. *Robustness.* Princeton, NJ: Princeton University Press.

Jullien, Bruno, and Bernard Salanié. 2000. "Estimating Preferences under Risk: The Case of Racetrack Bettors." *Journal of Political Economy*, 108(3): 503–30.

Jullien, Bruno, Bernard Salanié, and Francois Salanié. 2007. "Screening Risk-Averse Agents Under Moral Hazard: Single-crossing and the CARA case." *Economic Theory*, 30(1): 151–69.

Kirman, Alan P. 1992. "Whom or What Does the Representative Individual Represent?" *Journal of Economic Perspectives*, 6(2): 117–36.

Mirrlees, James A. 1971. "An Exploration in the Theory of Optimal Income Taxation." *Review of Economic Studies*, 38(114): 175–208.

Rosen, Sherwin. 1988. "The Value of Changes in Life Expectancy." *Journal of Risk and Uncertainty*, 1(3): 285–304.