AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

JOINT ESTIMATION OF RISK PREFERENCES AND TECHNOLOGY:
FLEXIBLE UTILITY OR FUTILITY?

Sergio H. Lence
Iowa State University

Contact information: Sergio H. Lence
368E Heady Hall
Iowa State University
Ames, IA 50011-1070
Phone: 515-294-6644
Fax: 515-294-4545
Email: shlence@iastate.edu

May 25, 2007

*Selected Paper prepared for presentation at the American Agricultural Economics Association
Annual Meeting, Portland, OR, July 29-August 1, 2007*

*Copyright 2007 by Sergio H. Lence. All rights reserved. Readers may make verbatim copies of
this document for non-commercial purposes by any means, provided that this copyright notice
appears on all such copies.*

JOINT ESTIMATION OF RISK PREFERENCES AND TECHNOLOGY:
FLEXIBLE UTILITY OR FUTILITY?

Abstract

The present study sets up a thought experiment calibrated to represent risks of a high-risk production activity (farming), and investigating whether the structure of risk aversion (i.e., the changes in absolute or relative risk aversion associated with changes in wealth) can be estimated with reasonable precision. Findings strongly suggest that typical production data are unlikely to allow identification of the structure of risk aversion. A flexible utility parameterization is found to worsen technology parameter estimates. Findings also indicate that even under a restricted utility specification, the quality of utility parameters estimated from small samples is very poor.

Keywords: expected utility, risk preferences, production analysis, risk attitudes.

JEL Codes: C13, D24, D81, Q12.

JOINT ESTIMATION OF RISK PREFERENCES AND TECHNOLOGY: FLEXIBLE UTILITY OR FUTILITY?

Risk and uncertainty are prominent features of agricultural production and marketing. Not surprisingly, numerous studies in agricultural economics have focused on producer behavior under uncertainty (Just and Pope 2002). One of the most popular topics of study in this field has been the estimation of decisions makers' risk preferences, both by direct elicitation from experimental data or responses to hypothetical questions (e.g., Binswanger, Robison), or by analyzing observed production and/or investment choices (e.g., Brink and McCarl; Antle 1987; Antle 1989; Love and Buccola; Saha, Shumway, and Talpaz; Chavas and Holt; Kumbhakar 2002a).

The seminal studies estimating risk preferences from actual production and/or investment decisions focused on the level of risk aversion, by estimating risk preferences separately from technology (e.g., Simmons and Pomareda, Brink and McCarl) and assuming restrictive utility functions (e.g., mean variance analysis). Such studies have been superseded by work where risk preferences are estimated simultaneously with technology, as doing so can increase estimation efficiency and may avoid inconsistency problems (e.g., Love and Buccola; Saha, Shumway, and Talpaz; Saha; Coyle; Kumbhakar 2002a; Kumbhakar 2002b; Chavas and Holt), even though Antle (1989) argued that there are some advantages in separating the estimation of technology and risk preferences. In addition, starting with Saha, Shumway, and Talpaz, the literature has emphasized the estimation of decision makers' "structure" of risk aversion (i.e., the changes in absolute or relative risk aversion associated with changes in wealth) by allowing for more flexible utility functions.

Knowledge about the structure of risk aversion is of interest because it determines, among other things, decision makers' responses to background risk, whether risky assets can be considered normal goods, and whether agents save for precautionary purposes (Gollier). Importantly, however, Kallberg and Ziemba (p. 1257) concluded that "... utility functions having

different functional forms and parameter values but ‘similar’ absolute risk aversion indices have ‘similar’ optimal portfolios.” More recently, Černý argued that, except for investments involving very large and skewed risks, decisions by agents with similar levels of “local” risk aversion are almost identical, regardless of their levels of relative risk aversion.

Estimation of the structure of risk aversion in production models is based on the premise that, in the presence of uncertainty, optimal input choices vary according to the decision maker’s structure of risk aversion. However, the studies by Kallberg and Ziemba, and Černý suggest that, given the same level of risk aversion (as measured by either the absolute risk aversion index or the level of local relative risk aversion), differences in optimal input decisions induced by different structures of risk aversion are negligible except for very large and skewed risks. This implies that, unless risks are very large and skewed, identification of the structure of risk aversion in production models may rely on too weak a source of information to allow the kind of econometric estimation that the literature has been pursuing.

The purpose of the present study is to investigate whether it is indeed feasible to estimate the structure of risk aversion given the risks underlying the data usually employed by researchers in empirical production analysis. To this end, a thought experiment is performed with risks calibrated using historical farm data. Farming is chosen because it involves substantial risks, thereby increasing the chances for the structure of risk aversion to exert a noticeable impact on production decisions. In addition, many of the empirical studies concerned have relied on data from production agriculture. Importantly, the experiment is designed to favor the likelihood of obtaining good estimates of the risk aversion structure, so that failure to get reasonable estimates provides strong evidence against the hypothesis that the structure of risk aversion can be recovered from production data. Succinctly, the experiment involves the generation of simulated data sets from simple and widely-used utility and production functions, and subsequently employing such data to simultaneously estimate utility and technology parameters (under the assumption that the econometrician knows the true functional forms of utility and production).

The study contributes to the literature by providing evidence against the hypothesis that typical production data contain enough information to allow identification of the structure of risk aversion. Further, a flexible utility parameterization seems to worsen the estimates of technology parameters. Overall, our findings greatly undermine the case for estimating the structure of risk aversion in studies of production. More generally, the method employed here may be useful in other situations where identification of the parameters and/or models of interest is suspected to be too weak to be useful, by allowing researchers to discard doomed-to-fail estimation projects as soon as possible.

I. Model

We adopt standard assumptions in the production literature by postulating that at decision time $t = 0$, a competitive producer chooses the amounts of inputs ($\mathbf{x}$) that maximize the expected utility of end-of-period random wealth:

$$(1.1) \quad \mathbf{x}^* = \operatorname{argmax}_{\mathbf{x}} E_W \{ U[\tilde{W}(\mathbf{x})] \},$$

where $\mathbf{x}^*$ denotes the vector of optimal input amounts, $E_W(\cdot)$ is the expectation operator with respect to random variable $\tilde{W}$, $U(\cdot)$ is the producer's utility function, and $\tilde{W}(\mathbf{x})$ is his end-of-period random wealth. The latter is defined to be the agent's initial wealth (W_0) plus random profits from production:

$$(1.2) \quad \tilde{W}(\mathbf{x}) \equiv \tilde{p} \tilde{y}(\mathbf{x}) - \mathbf{r} \mathbf{x} + W_0,$$

where $\tilde{p}$ is random end-of-period output price, $\tilde{y}(\mathbf{x})$ is random output, and $\mathbf{r}$ is the vector of variable input prices.

For the present purposes, model (1.1)-(1.2) is too general to be operational. To be able to make headway from an empirical standpoint, it is necessary to be more specific about the

utility function $U(\cdot)$, the technology $\tilde{y}(\mathbf{x})$, and the nature of randomness in price and output. Such issues are addressed in the following subsections.

I.1. The Decision Maker's Utility Function

The producer is assumed to be characterized by the hyperbolic absolute risk aversion (HARA) utility function (1.3):

$$(1.3) \quad U(W) = (1 - \gamma_1)^{-1} (\gamma_0 + W)^{1-\gamma_1},$$

which is defined on the domain of W satisfying $(\gamma_0 + W) > 0$. The negative of parameter γ_0 represents the agent's lowest admissible wealth. Parameter γ_1 is the agent's "baseline" risk aversion (Černý), and must be strictly positive if (1.3) is to represent risk-averse preferences. HARA utility is adopted here because it comprises the most popular functional forms used in expected utility analysis (i.e., the exponential, quadratic, and power utilities) (Gollier).

Importantly, quite different structures of risk aversion can be obtained under appropriate parameterizations of (1.3). To see this, note that the HARA coefficient of relative risk aversion is given by (1.4):

$$(1.4) \quad R(W) = \gamma_1 W (\gamma_0 + W)^{-1},$$

so that $\partial R(W)/\partial W = \gamma_1 \gamma_0 (\gamma_0 + W)^{-2}$. Since the sign of $\partial R(W)/\partial W$ is equal to the sign of parameter γ_0 , it follows that the HARA agent is characterized by decreasing, constant, or increasing relative risk aversion (DRRA, CRRA, and IRRA) if and only if parameter γ_0 is negative, zero, or positive, respectively.¹ Furthermore, as shown later, it is straightforward to

¹The coefficient of absolute risk aversion for HARA utility (1.3) is $A(W) = \gamma_1 (\gamma_0 + W)^{-1}$. Hence, such utility is characterized by decreasing absolute risk aversion, as $\partial A(W)/\partial W = -\gamma_1 (\gamma_0 + W)^{-2} < 0$. A three-parameter HARA utility would allow for preferences depicting increasing absolute risk aversion (e.g., Gollier, p.26). However, such generalization is not pursued here because the implied preferences are widely regarded as inconsistent with real-world attitudes toward risk (Gollier, p. 238).

parameterize (1.3) so as to test Černý's claim that optimal decisions are essentially the same regardless of whether the agent's utility is characterized by DRRA, CRRA, or IRRRA, except for decisions involving very large and skewed risks. More specifically, Černý labels $R(W)$ as the agent's "local" relative risk aversion, and argues that the key determinant for the optimal decision is $R(W_0)$ (i.e., the local risk aversion evaluated at the "safe" wealth level W_0).

I.2. Production Technology

The production technology $\tilde{y}(\mathbf{x})$ assumed for the analysis is of a Cobb-Douglas form:²

$$(1.5) \quad y(x_A, x_B; \tilde{e}_y) = \alpha_0 x_A^{\alpha_A} x_B^{\alpha_B} \tilde{e}_y,$$

where α_0 , α_A , and α_B are technology parameters, and $\tilde{e}_y$ is a random variable whose distribution is discussed in the next section. The main reason for adopting technology (1.5) is its simplicity. It seems highly unlikely for utility parameters γ_0 and γ_1 to be less accurately estimated under (1.5) than under more complex technologies. In addition, the Cobb-Douglas technology is arguably the most widely used production function in economic analysis.

II. Simulation Design, Calibration, and Simulated Data Generation

Given the model set up in the previous section, the present study's objective is to investigate whether it is possible to estimate the vector of utility parameters $\boldsymbol{\gamma} \equiv [\gamma_0, \gamma_1]$ simultaneously with the vector of technology parameters $\boldsymbol{\alpha} \equiv [\alpha_0, \alpha_A, \alpha_B]$ for levels of uncertainty --as reflected by the probability distributions of $\tilde{e}_y$ and $\tilde{p}$ -- usually found in production agriculture. The null hypothesis is that the typical data sets used in the empirical production literature do allow

²The present analysis is also performed using the Just-Pope production function $y(x_A, x_B; \tilde{e}_y) = \alpha_0 x_A^{\alpha_A} x_B^{\alpha_B} + \exp(\beta_A x_A + \beta_B x_B) \tilde{e}_y$ (Just and Pope 1978) instead of (1.5), parameterized with $\alpha_0 = 3$, $\alpha_A = 0.2$, $\alpha_B = 0.6$, $\beta_A = 0.06$, $\beta_B = -0.03$, and $\tilde{e}_y$ distributed as described in Appendix B. (Note that $\beta_B < 0$ means that input B is risk-reducing).

Results for the alternative production specification are omitted to save space, as the conclusions are essentially the same as for the simpler production function (1.5).

estimation of the structure of risk aversion (i.e., the simultaneous identification of γ_0 and γ_1). The validity of such hypothesis is evaluated by means of a thought experiment. The experiment consists of generating data corresponding to the postulated decision-making model by means of simulations, and then employing such data to estimate the underlying utility and technology parameters.

It is important to emphasize that the thought experiment is designed to obtain simulated data so as to favor the odds of being able to accurately estimate γ . That is, the data are deliberately constructed so as to exhibit nicer properties than actual field data. For example, the true behavioral model underlying the simulated data is very simple (i.e., (1.1) through (1.5)), all observations are identically and independently distributed (so that there is no time or cross-section correlation reducing the informational content of the data), there are neither optimization nor data-recording mistakes, inputs are not subject to physical constraints like field size or integer quantities, etc. Employing ideally well-behaved data is critical for our purposes, because it favors estimation of the structure of risk aversion. That is, the data are generated so as to protect the null hypothesis of identification to the extent possible. Rejecting the null hypothesis under such conditions should provide much stronger support for the claim that the structure of risk aversion cannot be recovered from field data.

II.1. Simulation Design

Simulated data are generated for nine basic scenarios, involving the combination of three structures of risk aversion (DRRA, CRRA, and IRRA) with three levels of uncertainty regarding random variables $\tilde{e}_y$ and $\tilde{p}$ (low, medium, and high variance). For each basic scenario, the simulated data consist of one million vectors, where the n th vector contains observations corresponding to the n th production decision $\mathbf{v}_n \equiv [W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}, x_{A,n}^*, x_{B,n}^*, p_n, y_n]$. Vector $\mathbf{v}_n$ comprises the relevant “exogenous” variables known at the time of making the n th decision (i.e., initial wealth $W_{0,n}$, output price at decision-making time $p_{0,n}$, and input prices $r_{A,n}$ and $r_{B,n}$), the corresponding optimal input amounts ($x_{A,n}^*$ and $x_{B,n}^*$), and the associated end-of-

period realizations of the variables that were random when the n th decision was made (i.e., output price p_n and output y_n).

II.2. Calibration

The DRRA, CRRA, and IRRA scenarios are obtained by parameterizing the HARA utility (1.4) with $[\gamma_0, \gamma_1]$ equal to $[-18.4, 1]$, $[0, 3]$, and $[43, 6]$, respectively. Parameters γ_0 and γ_1 are purposely set at substantially different levels across scenarios to facilitate obtaining different values at the estimation stage. The CRRA scenario with $[\gamma_0, \gamma_1] = [0, 3]$ implies a constant coefficient of relative risk aversion $R(W) = 3$, which reflects a reasonable level of relative risk aversion (Kocherlakota, Gollier, pp. 31 and 289). The advocated DRRA γ_1 parameterization constitutes a benchmark, as the coefficient of absolute prudence divided by the coefficient of absolute risk aversion equals exactly two when $\gamma_1 = 1$ in (1.3). Gollier (pp. 147, 287, 288, and 387) derives a series of risk preference properties that hinge upon whether the coefficient of absolute prudence divided by the coefficient of absolute risk aversion is smaller or greater than two.

Importantly, the aforementioned DRRA, CRRA, and IRRA parameterizations are chosen so that they all yield an average value of $R(W_0)$ equal to 3 under the assumed probability density function (pdf) of initial wealth (see (2.1) below). The latter feature allows us to investigate whether $R(W_0)$ does indeed reflect all that matters regarding the impact of risk preferences on optimal decisions, as argued by Černý. If he is correct, the DRRA and IRRA scenarios simulated here should yield optimal input amounts almost identical to the optimal input levels corresponding to the CRRA decision maker with coefficient of relative risk aversion $R(W) = 3$.

Since monetary units can be arbitrarily chosen, we standardize all prices by setting their unconditional means equal to one. The prices known to the agent when making the n th decision (i.e., $p_{0,n}$, $r_{A,n}$, and $r_{B,n}$) are obtained by assuming that they are identically and independently log-normally distributed with mean -0.03125 and variance 0.0625 , which implies a mean equal to one and a coefficient of variation of 25.4%. The twin assumptions of

independence and relatively large variability of decision-time prices are adopted to facilitate the estimation of utility parameters, as real-world data typically exhibit dependence and less variability than is being postulated here.

Production function (1.5) is parameterized with $\alpha_A = 0.2$ and $\alpha_B = 0.6$, which implies decreasing returns to scale. Having clearly different magnitudes for α_A and α_B may provide information about the impact of the parameter size on the precision of its estimate. To yield reasonable values of rates of return on variable inputs, scale parameter α_0 is fixed at a value of 3. If there were no uncertainty regarding $\tilde{e}_y$ and $\tilde{p}$, and they were fixed at their mean values of one, optimal inputs corresponding to the mean decision-time prices (i.e., $\bar{p}_0 = \bar{r}_A = \bar{r}_B = 1$) would equal $x_A^* = 2.0995$ units and $x_B^* = 6.2986$ units. The latter figures translate into costs of 8.3981 and gross revenues of 10.4976, for a net rate of return of 25% ($= 10.4976/8.3981 - 1$) on variable inputs.

Simulated observations on initial wealth are generated from (2.1):

$$(2.1) \quad W_0 = 18.9 + 69.2 z,$$

where z is random variable distributed according to the standard beta pdf $Beta(0.87, 1.27)$.

Hence, initial wealth has lower bound $\underline{W}_0 = 18.9$, upper bound $\overline{W}_0 = 88.1 (= 18.9 + 69.2)$, mean 47.03, and standard deviation 19.18. The rationale for using (2.1) is twofold. First, (2.1) provides a distribution of initial wealth consistent with real-world data under the advocated standardization of prices and quantities (see Appendix A for details). Second, as pointed out earlier, the mean value of $R(W_0)$ corresponding to (2.1) equals 3 under the adopted DRRA, CRRA, and IRRA parameterizations.

Output shocks $\tilde{e}_y$ are assumed to have a probability distribution whose shape mimics the empirical distribution of farm-level corn yields in Iowa. Estimation of the probability distribution of $\tilde{e}_y$ is explained in Appendix B. Output shocks have a mean of one; in the medium-variance scenario their standard deviation is 0.207, whereas in the low- and high-

variance scenarios their standard deviations are 0.104 and 0.310, respectively.³

Finally, the postulated random generating process for end-of-period crop prices is (2.2):

$$(2.2) \quad \ln(\tilde{p}) = \mu_p + 0.5 \ln(p_0) - 0.3 \ln(\tilde{e}_y) + \tilde{e}_p,$$

where p_0 is the (known) price at the time of decision making, and $\tilde{e}_p$ is a zero-mean normally distributed random variable. The term involving $\ln(\tilde{e}_y)$ accounts for the stylized fact that output shocks tend to have a negative impact on output prices. The standard deviations of $\tilde{e}_p$ for the low-, medium-, and high-variance scenarios are 0.1, 0.2, and 0.3, respectively.⁴ To obtain an unconditional mean of $\tilde{p}$ equal to 1, the intercept term μ_p is fixed at -0.0232 in the medium-variance scenario, and 0.0014 and -0.0659 in the low- and high-variance scenarios, respectively. Expression (2.2) is based on historical price and yield-shock data, and its estimation is described in Appendix C.

II.3. Generation of Simulated Data for the n th Production Decision

Conceptually, vector $\mathbf{v}_n \equiv [W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}, x_{A,n}^*, x_{B,n}^*, p_n, y_n]$ comprising simulated data for the n th production decision is calculated in three steps. In the first step, a random draw from (2.1) is used to compute initial wealth ($W_{0,n}$), and random draws from the respective log-normal distributions are used to obtain decision-time output price ($p_{0,n}$) and input prices ($r_{A,n}$ and $r_{B,n}$). In the second step, numerical methods are employed to solve for $x_{A,n}^*$ and $x_{B,n}^*$, i.e., the input amounts that maximize expected utility given the information available at decision time. The latter comprises the vector $[W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}]$, the production technology, and the probability distributions of $\tilde{e}_y$ and $\tilde{p}$. In the third step, actual realizations of output (y_n) and output price (p_n) are obtained. Realized output is calculated by drawing an output shock ($e_{y,n}$) from the

³That is, the standard deviation in the low-variance (high-variance) scenario is 50% smaller (larger) than the standard deviation in the medium-variance scenario.

⁴Again, standard deviations in the low- and high-variance scenarios are respectively 50% smaller and 50% larger than the standard deviation in the medium-variance scenario.

probability distribution of $\tilde{e}_y$, and plugging it along with optimal inputs ($x_{A,n}^*$ and $x_{B,n}^*$) into production function (1.5). As per realized output price, it is computed by drawing a price shock ($e_{p,n}$) from the pdf of $\tilde{e}_p$, and substituting it together with the decision-time price ($p_{0,n}$) and the output shock ($e_{y,n}$) into price equation (2.2).

To make results as comparable as possible for the alternative (DRRA, CRRA, IRRA) $\times$ (low-, medium-, high-variance) scenarios, the same vector of “exogenous” decision-time variables $[W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}]$ is used across all nine basic scenarios for the n th production decision. A different procedure is required for end-of-period shocks $e_{y,n}$ and $e_{p,n}$, however, because having low-, medium-, and high-variance cases prevents us from using the same $[e_{y,n}, e_{p,n}]$ values across all scenarios. Hence, to compute end-of-period shocks so as to enhance comparability across settings, we first draw from appropriate standard pdfs (uniform for $\tilde{e}_y$ and normal for $\tilde{e}_p$), and then convert such draws into the shocks corresponding to the alternative scenarios by performing suitable transformations.⁵

Calculation of optimal inputs is the step that requires the greatest computational effort. To describe how the numerical optimization is performed, let function $u(x_{A,n}, x_{B,n}, \tilde{e}_y, \tilde{e}_p; W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n})$ represent the utility of random end-of-period wealth corresponding to the n th production decision. With this notation, the optimization problem consists of maximizing (2.3) with respect to $x_{A,n}$ and $x_{B,n}$:

$$(2.3) \quad E_{e_y, e_p} [u(x_{A,n}, x_{B,n}, \tilde{e}_y, \tilde{e}_p; W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n})] \equiv \iint u(x_{A,n}, x_{B,n}, e_y, e_p; \cdot) f_{y,p}(e_y, e_p) de_y de_p,$$

where $f_{y,p}(\cdot)$ denotes the joint pdf of $\tilde{e}_y$ and $\tilde{e}_p$. However, output shocks ($\tilde{e}_y$) are assumed to be independent of price shocks ($\tilde{e}_p$), so that $f_{y,p}(e_y, e_p) = f_y(e_y) f_p(e_p)$, where $f_y(\cdot)$ and $f_p(\cdot)$ are the marginal pdfs of output and price shocks, respectively. Further, since the expectation in (2.3) has no analytical solution, its computation requires numerical quadrature methods (Miranda and

⁵To illustrate this point, consider the generation of the n th end-of-period price shocks. We first draw a realization $z_{p,n}$ from the standard normal pdf, and then calculate $e_{p,n}$ for the low-, medium-, and high-variance scenarios as $0.1 z_{p,n}$, $0.2 z_{p,n}$, and $0.3 z_{p,n}$, respectively (see explanation of (2.2) above).

Fackler, ch. 4). Hence, the objective function employed for numerical optimization is the one shown on the right-hand-side of (2.4):

$$(2.4) \quad E_{e_y, e_p} [u(x_{A,n}, x_{B,n}, \tilde{e}_y, \tilde{e}_p; W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n})] \equiv \sum_s \sum_q u(x_{A,n}, x_{B,n}, e_{y,q}, e_{p,s}; \cdot) \pi_{y,q} \pi_{p,s},$$

where $e_{y,q}$ and $e_{p,s}$ are quadrature nodes, and $\pi_{y,q}$ and $\pi_{p,s}$ are the respective quadrature weights. For output shocks, 100 nodes (and weights) are used, as described in Appendix B. Nodes and weights for price shocks are determined by a 10-point Gaussian quadrature.⁶

Numerical optimization of the objective function on the right-hand side of (2.4) is performed by means of the “fmincon” function in MATLAB version 7.0.4.365. Function “fmincon” is set up to find a minimum of a constrained nonlinear multivariable objective function. In the present setup, the only constraint imposed on the optimization is that costs $(r_{A,n} x_{A,n}^* + r_{B,n} x_{B,n}^*)$ do not exceed $\gamma_0 + W_{0,n}$; otherwise, the HARA condition $\gamma_0 + W_n > 0$ would be violated for sufficiently small levels of revenue $(p_n y_n)$.

The accuracy of the “fmincon” optimal inputs is verified by comparing them with the input amounts obtained by performing a grid-search optimization for a test set, consisting of the 625 ($= 5^4$) combinations of initial wealth and decision-time prices resulting from the Cartesian product of the 5-point Gaussian quadrature nodes for W_0 , p_0 , r_A , and r_B . Test set results are also used to construct initial values (through regressions of the corresponding x_A^* and x_B^* on W_0 , p_0 , r_A , and r_B) for the optimization concerning $x_{A,n}$ and $x_{B,n}$.

II.4. Generation of Simulated Samples

The procedure described in the preceding subsection is used to generate 0.5 million $\mathbf{v}_n$ vectors (i.e., data for 0.5 million production decisions) for each of the nine basic scenarios analyzed. Such vectors are then grouped into 5,000 (1,000; 500) samples of 100 (500; 1,000) vectors each.

⁶By construction, an S -point Gaussian quadrature computes the first $2S - 1$ moments of $\tilde{e}_p$ without error.

In turn, each of the simulated samples is used to perform an econometric estimation of the utility and technology parameters, resulting in 5,000 (1,000; 500) estimates for each parameter from the samples with 100 (500; 1,000) observations, respectively.

Since the objects of interest are the distributions of the parameter estimates, to enhance the estimation of such distributions the data set is augmented by means of antithetic replications (Geweke). More specifically, an additional set of 0.5 million antithetic decision vectors $\mathbf{v}_{n+} \equiv [W_{0,n+}, p_{0,n+}, r_{A,n+}, r_{B,n+}, x_{A,n+}^*, x_{B,n+}^*, p_{n+}, y_{n+}]$ are constructed and used for estimation. Variables $W_{0,n+}$, $p_{0,n+}$, $r_{A,n+}$, and $r_{B,n+}$ are the antithetic replications of $W_{0,n}$, $p_{0,n}$, $r_{A,n}$, and $r_{B,n}$, respectively. Variables $x_{A,n+}^*$ and $x_{B,n+}^*$ are the optimal inputs corresponding to $W_{0,n+}$, $p_{0,n+}$, $r_{A,n+}$, and $r_{B,n+}$. Finally, p_{n+} is calculated from (2.2) using $p_{0,n+}$ and the antithetic replications of $e_{y,n}$ and $e_{p,n}$, whereas y_{n+} is obtained by plugging $x_{A,n+}^*$, $x_{B,n+}^*$, and the antithetic replication of $e_{y,n}$ into (1.5). Antithetic replications allow us to obtain a total of 10,000 (2,000; 1,000) estimates for each parameter from the samples with 100 (500; 1,000) observations.⁷

II.5. “Mixed-Variance” Scenarios

In addition to the nine aforementioned basic scenarios, a set of three “mixed-variance” scenarios are generated for DRRA, CRRA, and IRRA preferences. In the mixed-variance settings, the n th production decision vector is augmented by incorporating variable σ_n (i.e., $\mathbf{v}_n \equiv [W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}, \sigma_n, x_{A,n}^*, x_{B,n}^*, p_n, y_n]$), where σ_n equals 0.5, 1, or 1.5 depending on whether the distribution of price and output shocks for the n th decision have small-, medium-, or large-variance, respectively. The mixed-variance scenarios are aimed at incorporating heteroskedasticity in the output and price shocks. The variance of shocks may change from period to period, but is assumed to be known at the time of decision making. Hence, the n th period’s optimal choice incorporates the level of uncertainty corresponding to the n th period’s output and price shocks. Data for the n th period mixed-variance scenario are straightforward to generate, by randomly

⁷Data corresponding to one million production decisions can obviously be generated without resorting to antithetic replications. However, they allow us to obtain more accurate estimates of the distributions of interest.

selecting $\mathbf{v}_n$ from the small-, medium-, and large-variance scenarios with probabilities 1/3 each, and augmenting such vector by the associated σ_n .

III. Estimation

To favor the null hypothesis that utility parameters can be recovered from the production decision data, we assume that the econometrician knows the specific form of the utility and production functions, and is only interested in estimating their corresponding parameters. Therefore, the present estimation is not affected by issues pertaining to functional form approximations.

Clearly, shocks enter the decision maker's objective function in a highly nonlinear fashion. Thus, even if shocks followed a standard (e.g., normal or log-normal distribution), maximum likelihood estimation would require a linear approximation to render the problem tractable (Jagannathan, Skoulakis, and Wang). Further, the postulated distribution of output shocks is not standard and it would be highly unrealistic to assume that it is known to the econometrician. This hampers the use of maximum likelihood (ML), as ML estimates may not be consistent when the distribution of the observable variables is misspecified (Hansen and Singleton). For such reasons, estimation is performed by means of the generalized method of moments (GMM).⁸

Useful references for the theory underlying GMM and its numerous applications include Hansen, Davidson and Mackinnon (ch. 17), Ogaki, and Cliff. The optimal decision making framework lends itself nicely to apply the GMM method. Succinctly, estimation is based on a system of three regression equations corresponding to the production function (i.e., (3.1)) and the first-order conditions (FOCs) for optimization of (2.3) (i.e., (3.2)):

$$(3.1) \quad \varepsilon_{y,n}(\boldsymbol{\alpha}) \equiv y_n - \alpha_0 - \alpha_A x_{A,n}^{\alpha_A} - \alpha_B x_{B,n}^{\alpha_B},$$

⁸We are not aware of a method other than GMM to obtain consistent parameter estimates in the present setup.

$$(3.2) \quad \varepsilon_{j,n}(\alpha_j, \gamma) \equiv (\gamma_0 + p_n y_n - r_{A,n} x_{A,n}^* - r_{B,n} x_{B,n}^* - W_{0,n})^{-\gamma_1} (p_n \alpha_j x_{j,n}^{*-1} y_n - r_{j,n}),$$

for $j = A$ and B , and $n = 1, \dots, N$, where N is the sample size. Note that $\varepsilon_{y,n}$ provides no information to identify utility parameters, and that $\varepsilon_{A,n}$ ($\varepsilon_{B,n}$) contains no information to identify production parameters α_0 and α_B (α_A). For the small-, medium-, and large-variance scenarios, the set of instruments used for each of the equations above consists of vector $\mathbf{t}_n \equiv [1, W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}, x_{A,n}^*, x_{B,n}^*]'$, comprising a constant (standardized to unity) and the variables known at the time of decision making. For the mixed-variance scenario, vector $\mathbf{t}_n$ is augmented by the variance variable σ_n .⁹

The rationale for using GMM in the present study is the same as for the typical use of GMM to test asset pricing relationships via Euler equations (e.g., Altuğ and Labadie, ch. 3). That is, for any variable z known at the time of decision making, it must be the case that $E_{e_y, e_p}(\tilde{\varepsilon}_j z) = E_{e_y, e_p}(\tilde{\varepsilon}_j) z = 0$ for $j = y, A, B$. This is true because z is non-random from the perspective of the decision time, $E_{e_y, e_p}(\tilde{\varepsilon}_y) = 0$ from production function (1.5), and $E_{e_y, e_p}(\tilde{\varepsilon}_A) = E_{e_y, e_p}(\tilde{\varepsilon}_B) = 0$ from FOCs. By the law of iterated expectations, it follows that the unconditional expectations are also zero, $E(\tilde{\varepsilon}_j z) = E[E_{e_y, e_p}(\tilde{\varepsilon}_j z)] = 0$ for $j = y, A, B$. Given a sample of size N , the set of sample counterparts of $E(\tilde{\varepsilon}_j z)$ is the vector $\mathbf{g}_N(\boldsymbol{\alpha}, \boldsymbol{\gamma}) \equiv 1/N \sum_n [\varepsilon_{y,n}(\boldsymbol{\alpha}) \varepsilon_{A,n}(\alpha_B, \boldsymbol{\gamma}) \varepsilon_{B,n}(\alpha_A, \boldsymbol{\gamma})]' \otimes \mathbf{t}_n$, where $\otimes$ denotes the Kronecker product. Since parameters $[\boldsymbol{\alpha}, \boldsymbol{\gamma}]$ are the only elements of $\mathbf{g}_N(\cdot)$ unknown to the econometrician, the GMM estimates $[\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\gamma}}]$ are chosen so as to render $\mathbf{g}_N(\cdot)$ as close to zero as possible by minimizing a quadratic form in $\mathbf{g}_N(\cdot)$ with respect to the unknown parameters:

$$(3.3) \quad [\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\gamma}}] = \arg\max_{[\boldsymbol{\alpha}, \boldsymbol{\gamma}]} [\mathbf{g}_N(\boldsymbol{\alpha}, \boldsymbol{\gamma}) \mathbf{g}_N(\boldsymbol{\alpha}, \boldsymbol{\gamma})' \mathbf{V}_N \mathbf{g}_N(\boldsymbol{\alpha}, \boldsymbol{\gamma})].$$

In (3.3), $\mathbf{V}_N$ is a positive definite weighting matrix which converges in probability to a positive

⁹Including σ_n in the set of instruments implies that the econometrician knows with certainty whether the variance associated with the n th period's shocks is small, medium, or large. This is obviously unrealistic, but is assumed here to favor the null hypothesis of identification.

definite matrix V_0 . Loosely speaking, the GMM estimates yield the sample counterparts of orthogonality conditions $E(\tilde{\varepsilon}_j z)$ as close to zero as possible.

Estimation is performed using the GMM and MINZ software libraries in MATLAB version 7.0.4.365 (Cliff), which can be downloaded from <http://mcliff.cob.vt.edu/progs.html>. Results reported in the next section are obtained from two-step GMM estimation. The true parameter values are used to initialize the numerical GMM optimization, so as to facilitate convergence. Also to enhance convergence, the weighting matrix for the first GMM step is set equal to $V_N = (\mathbf{1}_k \otimes \mathbf{t}' \mathbf{t})^{-1}$, where $\mathbf{1}_k$ is the $(k \times k)$ identity matrix, k is the number of regression equations (three in the present application), and $\mathbf{t} = [\mathbf{t}_1, \dots, \mathbf{t}_N]'$ (Cliff). In addition, to ensure that $\hat{\gamma}_0$ does not violate the constraint $(\hat{\gamma}_0 + W) > 0$ associated with HARA utility (1.3), for estimation purposes we impose the restriction that $\hat{\gamma}_0 > 0.001 - \min(W_{1,n}, n = 1, \dots, N)$, where $W_{1,n} \equiv p_n y_n - r_{A,n} x_{A,n}^* - r_{B,n} x_{B,n}^* - W_{0,n}$.

IV. Results and Discussion

Estimation results for the flexible utility specification are summarized in Tables 1 through 3. The tables contain the mean, the median (within parentheses), and the 2.5% and 97.5% quantiles (within brackets) for each of the utility and technology parameter estimates. In the following discussion, the 2.5%-97.5% quantile intervals are referred to as the 95% confidence intervals (CIs). Tables 1, 2, and 3 report results for the small-, large-, and mixed-variance scenarios, respectively. Outcomes for the medium-variance scenario are not reported to save space, as they typically lie between the small- and large-variance results.

The estimates for technology parameters exhibit patterns to be expected. The precision of $\hat{\alpha}_0$, $\hat{\alpha}_A$, and $\hat{\alpha}_B$ --as measured by the width of the 95% CIs-- increases with the sample size and decreases with the variance of the output and price shocks. Parameter α_B (whose true value is 0.6) tends to be estimated with slightly greater precision than α_A (which has a true value of 0.2). Technology estimates for the samples with 100 observations and large variance shocks are imprecise. This is true because the 95% CIs for $\hat{\alpha}_A$ overlap with the respective 95% CIs

Table 1. Parameter estimates for flexible utility estimation specification, small variance scenario.

Risk Structure	Sample Size	Parameter Estimates ^a				
		Utility		Technology		
		$\hat{\gamma}_0$	$\hat{\gamma}_1$	$\hat{\alpha}_0$	$\hat{\alpha}_A$	$\hat{\alpha}_B$
DRRA ^b	100	−19.8 ^c (−19.8) [−21.2, −15.3]	−63.8 ^c (−21.5) [−805.9, −4.8]	3.03 ^c (3.03) [2.03, 3.47]	0.20 ^c (0.20) [0.05, 0.33]	0.60 ^c (0.60) [0.46, 0.72]
DRRA ^b	500	−19.4 (−19.4) [−19.8, −19.1]	−61.4 (−13.4) [−44.5, −3.3]	3.06 (3.06) [2.90, 3.22]	0.20 (0.20) [0.15, 0.25]	0.59 (0.59) [0.54, 0.64]
DRRA ^b	1,000	−19.3 (−19.3) [−19.6, −19.0]	−11.8 (−11.8) [−20.6, −3.7]	3.06 (3.06) [2.96, 3.16]	0.20 (0.20) [0.16, 0.23]	0.59 (0.59) [0.56, 0.62]
CRRA ^b	100	1155.5 (32.5) [2.3, 299.2]	71.6 (32.6) [5.5, 279.7]	3.01 (3.01) [2.74, 3.29]	0.20 (0.20) [0.11, 0.29]	0.60 (0.60) [0.52, 0.68]
CRRA ^b	500	124.9 (31.1) [6.4, 210.7]	146.9 (31.1) [8.2, 218.8]	3.01 (3.01) [2.90, 3.12]	0.20 (0.20) [0.16, 0.24]	0.60 (0.60) [0.57, 0.63]
CRRA ^b	1,000	65.4 (27.4) [6.8, 205.5]	56.5 (27.3) [8.5, 230.3]	3.01 (3.01) [2.94, 3.11]	0.20 (0.20) [0.17, 0.24]	0.60 (0.60) [0.57, 0.62]
IRRA ^b	100	7732.4 (43.9) [7.2, 424.9]	35.8 (13.0) [2.3, 100.2]	3.01 (3.01) [2.74, 3.29]	0.20 (0.20) [0.11, 0.30]	0.60 (0.60) [0.52, 0.68]
IRRA ^b	500	61.4 (44.4) [43.1, 87.7]	153.3 (17.0) [7.2, 253.9]	3.01 (3.01) [2.90, 3.11]	0.20 (0.20) [0.16, 0.24]	0.60 (0.60) [0.57, 0.63]
IRRA ^b	1,000	55.4 (44.6) [43.1, 81.6]	96.0 (17.7) [7.1, 262.2]	3.01 (3.01) [2.94, 3.08]	0.20 (0.20) [0.18, 0.22]	0.60 (0.60) [0.58, 0.62]

^aFor each parameter, the table reports the average, the median (within parentheses), and the 2.5% and 97.5% quantiles (within brackets).

^bDRRA, CRRA, and IRRA risk structures correspond to $[\gamma_0, \gamma_1]$ equal to $[-18.4, 1]$, $[0, 3]$, and $[43, 6]$, respectively.

^cMean is calculated using only observations falling between the 2.5% and 97.5% quantiles, to avoid influence of extreme observations.

Table 2. Parameter estimates for flexible utility estimation specification, large variance scenario.

Risk Structure	Sample Size	Parameter Estimates ^a				
		Utility		Technology		
		$\hat{\gamma}_0$	$\hat{\gamma}_1$	$\hat{\alpha}_0$	$\hat{\alpha}_A$	$\hat{\alpha}_B$
DRRA ^b	100	−19.5 ^c (−19.5) [−20.9, −18.8]	−29.8 ^c (−10.5) [−428.8, 0.2]	3.14 ^c (3.12) [2.11, 4.30]	0.18 ^c (0.18) [−0.13, 0.52]	0.57 ^c (0.57) [0.24, 0.86]
DRRA ^b	500	−15.6 (−19.0) [−19.5, −18.7]	−6.3 (−6.6) [−14.8, −0.2]	3.24 (3.24) [2.82, 3.62]	0.18 (0.18) [0.07, 0.28]	0.55 (0.55) [0.45, 0.66]
DRRA ^b	1,000	−18.8 (−18.9) [−19.3, −18.7]	−4.5 (−5.4) [−10.2, −0.3]	3.25 (3.25) [2.99, 3.51]	0.18 (0.18) [0.12, 0.23]	0.54 (0.55) [0.48, 0.61]
CRRA ^b	100	1531.6 (29.0) [7.2, 258.4]	96.2 (43.7) [8.0, 335.7]	3.01 (3.01) [2.28, 3.85]	0.20 (0.21) [−0.05, 0.47]	0.58 (0.60) [0.37, 0.82]
CRRA ^b	500	50.9 (28.4) [10.2, 204.6]	53.9 (38.5) [9.4, 204.5]	3.01 (3.01) [2.72, 3.32]	0.21 (0.20) [0.11, 0.31]	0.59 (0.60) [0.51, 0.68]
CRRA ^b	1,000	46.3 (24.1) [11.2, 214.9]	50.3 (28.5) [9.9, 191.1]	3.01 (3.01) [2.82, 3.21]	0.20 (0.20) [0.14, 0.27]	0.59 (0.60) [0.54, 0.66]
IRRA ^b	100	7016.4 (43.6) [7.0, 168.7]	33.0 (14.7) [3.9, 148.0]	3.01 (3.00) [2.33, 3.75]	0.21 (0.20) [−0.03, 0.46]	0.61 (0.61) [0.39, 0.82]
IRRA ^b	500	46.7 (44.1) [23.3, 62.0]	31.5 (14.8) [7.0, 101.0]	3.00 (3.00) [2.71, 3.30]	0.20 (0.20) [0.10, 0.31]	0.60 (0.60) [0.51, 0.68]
IRRA ^b	1,000	48.8 (44.2) [43.1, 83.3]	39.6 (14.6) [7.3, 143.0]	3.00 (3.00) [2.81, 3.22]	0.20 (0.20) [0.13, 0.27]	0.60 (0.60) [0.54, 0.66]

^aFor each parameter, the table reports the average, the median (within parentheses), and the 2.5% and 97.5% quantiles (within brackets).

^bDRRA, CRRA, and IRRA risk structures correspond to $[\gamma_0, \gamma_1]$ equal to $[-18.4, 1]$, $[0, 3]$, and $[43, 6]$, respectively.

^cMean is calculated using only observations falling between the 2.5% and 97.5% quantiles, to avoid influence of extreme observations.

Table 3. Parameter estimates for flexible utility estimation specification, mixed variance scenario.

Risk Structure	Sample Size	Parameter Estimates ^a				
		Utility		Technology		
		$\hat{\gamma}_0$	$\hat{\gamma}_1$	$\hat{\alpha}_0$	$\hat{\alpha}_A$	$\hat{\alpha}_B$
DRRA ^b	100	−19.7 ^c (−19.6) [−21.0, −18.8]	−42.4 ^c (−14.3) [−578.1, −0.3]	3.09 ^c (3.07) [2.37, 3.90]	0.19 ^c (0.19) [−0.04, 0.43]	0.59 ^c (0.59) [0.35, 0.80]
DRRA ^b	500	−19.2 (−19.2) [−19.7, −18.8]	−8.2 (−8.0) [−19.6, −0.9]	3.14 (3.14) [2.85, 3.40]	0.19 (0.19) [0.11, 0.27]	0.58 (0.58) [0.51, 0.66]
DRRA ^b	1,000	−19.1 (−19.1) [−19.4, −18.7]	−5.7 (−5.3) [−13.1, −1.2]	3.15 (3.15) [2.97, 3.31]	0.19 (0.19) [0.14, 0.24]	0.58 (0.58) [0.53, 0.62]
CRRA ^b	100	3220.0 (32.1) [6.7, 228.9]	62.6 (38.0) [7.7, 247.8]	3.00 (3.00) [2.50, 3.55]	0.18 (0.20) [0.03, 0.37]	0.55 (0.60) [0.44, 0.75]
CRRA ^b	500	68.1 (37.6) [11.1, 204.9]	58.3 (37.2) [10.4, 149.3]	3.00 (3.00) [2.78, 3.24]	0.20 (0.20) [0.13, 0.28]	0.58 (0.60) [0.53, 0.66]
CRRA ^b	1,000	83.7 (34.7) [12.0, 177.9]	46.1 (33.3) [11.1, 135.2]	3.01 (3.00) [2.82, 3.20]	0.20 (0.20) [0.14, 0.27]	0.59 (0.60) [0.54, 0.65]
IRRA ^b	100	13505.0 (43.7) [7.1, 313.9]	34.1 (13.3) [2.6, 115.0]	3.01 (3.01) [2.54, 3.52]	0.20 (0.20) [0.04, 0.36]	0.60 (0.60) [0.46, 0.74]
IRRA ^b	500	50.5 (44.2) [15.8, 95.1]	52.6 (16.3) [6.8, 202.9]	3.00 (3.00) [2.80, 3.20]	0.20 (0.20) [0.13, 0.27]	0.60 (0.60) [0.54, 0.66]
IRRA ^b	1,000	171.1 (44.7) [20.0, 146.1]	404.6 (19.8) [7.0, 196.6]	3.00 (3.00) [2.87, 3.14]	0.20 (0.20) [0.16, 0.24]	0.60 (0.60) [0.56, 0.64]

^aFor each parameter, the table reports the average, the median (within parentheses), and the 2.5% and 97.5% quantiles (within brackets).

^bDRRA, CRRA, and IRRA risk structures correspond to $[\gamma_0, \gamma_1]$ equal to $[-18.4, 1]$, $[0, 3]$, and $[43, 6]$, respectively.

^cMean is calculated using only observations falling between the 2.5% and 97.5% quantiles, to avoid influence of extreme observations.

for $\hat{\alpha}_B$; further, the lower bounds for the former are negative. For CRRA and IRRRA preferences, median technology estimates are essentially the same as the corresponding true values. For DRRA preferences, however, the medians of $\hat{\alpha}_0$, $\hat{\alpha}_A$, and $\hat{\alpha}_B$ tend to be slightly different from the true values when the variance of shocks is large, and such differences increase with the sample size. In such instances, medians for $\hat{\alpha}_0$ tend to be larger than the true values, and the opposite is true for the $\hat{\alpha}_A$ and $\hat{\alpha}_B$ medians. The distributions of technology estimates for the mixed-variance scenarios fall between those for the small- and large-variance counterparts.

Unlike technology estimates, the estimates of the preference parameters bear little resemblance to the true values. Consider the CRRA scenario first. Both γ_0 and γ_1 are grossly overestimated, as the smallest $\hat{\gamma}_0$ and $\hat{\gamma}_1$ medians are respectively 24.1 and 27.3 (1000-observation samples for large- and small-variance scenarios, respectively), compared to true values of $\gamma_0 = 0$ and $\gamma_1 = 3$. Prominently, none of the 95% CIs for $\hat{\gamma}_0$ and $\hat{\gamma}_1$ include the true values. Further, as the sample size increases, the distance between the true values and the closest (i.e., lowest) bounds of the 95% CIs increases. For example, when the variance is large and samples have 100, 500, and 1000 observations, the 2.5% quantiles for $\hat{\gamma}_0$ ($\hat{\gamma}_1$) are 7.2, 10.2 and 11.2 (8.0, 9.4, and 9.9), respectively. It is also worth noting that the distributions of $\hat{\gamma}_0$ and $\hat{\gamma}_1$ are substantially skewed to the right.

IRRA estimates $\hat{\gamma}_0$ and $\hat{\gamma}_1$ tend to be slightly better behaved than their CRRA counterparts. This is true because the 95% CIs sometimes include the true values, and the medians for $\hat{\gamma}_0$ are relatively close to the true value $\gamma_0 = 43$. Overall, however, $\hat{\gamma}_0$ and $\hat{\gamma}_1$ display very poor behavior. For example, the 95% CIs are extremely wide, and the medians are further away from the true values as the sample sizes increase. As in the CRRA case, the distributions of $\hat{\gamma}_0$ and $\hat{\gamma}_1$ are noticeably right-skewed.

The worst-behaved preference estimates are the $\hat{\gamma}_1$ s under DRRA. In this instance, not only the 95% CIs are quite wide and lie strictly below the true value $\gamma_1 = 1$, but also they fall entirely on the negative orthant (except for a minuscule segment of the 95% CI for 100-

observation samples in the large-variance setting). The latter means that the $\hat{\gamma}_1$ s would mislead the econometrician into concluding that the decision maker exhibits risk-loving preferences. The 95% CIs for γ_0 under DRRA are much narrower than under CRRA or IRRA. However, the former almost always lie entirely below the true value of $\gamma_1 = -18.4$. The reason for the relatively narrow 95% CIs for γ_0 under DRRA is that the estimation constraint $\hat{\gamma}_0 > 0.001 - \min(W_{1n}, n = 1, \dots, N)$ is almost always binding.

The mixed-variance scenario reported in Table 3 has the potential to provide additional information to aid in the identification of the structure of risk aversion. This is true because decision makers differing in their risk preferences will generally change their input choices in different ways in response to changes in risks. However, comparison of the distributions of $\hat{\gamma}_0$ and $\hat{\gamma}_1$ in Table 3 with the ones in Tables 1 and 2 shows that heteroskedasticity is of little help in the estimation of risk preferences. Additional support for this conclusion is furnished by the fact that the estimates in Table 3 assume that the econometrician knows exactly the level of uncertainty embedded in the decision maker's optimal choices, which is highly unrealistic. Therefore, it seems safe to conclude that in real-world applications, the structure of risk aversion is very unlikely to be rendered identifiable by the presence of heteroskedasticity.

IV.1. Estimation under Restricted Utility Specifications

As noted earlier, Černý argues that decisions corresponding to utilities with the same value of $R(W_0)$ are very similar, unless random shocks are very skewed and have substantial variances. If such argument applies to agricultural production under uncertainty, observed input choices would be consistent with an infinite number of γ parameterizations yielding the same $R(W_0)$. Since the econometrician must infer γ_0 and γ_1 from the observed input choices, this means that attempting to simultaneously estimate both γ_0 and γ_1 is likely to encounter difficulties due to lack of identification.

A simple solution to the identification problem is to perform estimation under a restricted utility specification, by fixing either γ_0 or γ_1 at a reasonable value, and estimate the remaining

utility parameter. An appealing specification consists of setting $\gamma_0 = 0$ and estimating only γ_1 , in fact approximating the true DRRA and IRRA utilities with CRRA preferences. The nice feature of such specification is that the resulting $\hat{\gamma}_1$ provides an approximate estimate of both the coefficients of “baseline” and “local” relative risk aversion, where the latter is the key determinant of decisions under uncertainty. Given the calibration used here, restricted estimation assuming CRRA preferences should yield $\hat{\gamma}_1 \cong 3$.

Results from the restricted utility specification are provided in Tables 4, 5, and 6, which correspond respectively to Tables 1, 2, and 3 under unrestricted estimation. Compared to the latter, the former exhibit substantial differences regarding the estimates of both technology and utility parameters. Regarding α_0 , α_A , and α_B , the most noticeably impact of the restricted estimation is the substantially narrower 95% CIs. In addition, in the instances where the medians tend to be most different from the true values (e.g., for α_0), such differences are typically smaller for the restricted estimates than for the unrestricted ones.

Inspection of the estimates of the coefficient of relative risk aversion shown in Tables 4 through 6 uncovers a number of interesting outcomes. First, estimates of γ_1 obtained using samples with 100 observations appear to be very poor, as the associated 95% CIs are very wide. The 2.5% quantiles under large variance are negative (i.e., indicating risk affinity rather than risk aversion), and the 97.5% quantiles under small variance are unreasonably large. These findings cast doubt on the ability of being able to accurately estimate risk preferences (even if restricted) from small samples.

A second interesting feature of the $\hat{\gamma}_1$ distributions is that the means and medians seem to systematically overestimate the true value of the coefficient of local relative risk aversion (i.e., 3). The magnitude of the positive bias in $\hat{\gamma}_1$ declines with the number of observations in the sample, but for DRRA preferences it is still substantial even when samples have as many as 1000 observations.

Third, the quality of the γ_1 estimates is noticeably worse in the small variance scenario than in the large variance scenario. Greater variance increases the overall noise in the

Table 4. Parameter estimates for restricted utility estimation specification ($\gamma_0 = 0$), small variance scenario.

Risk Structure	Sample Size	Parameter Estimates ^a			
		Utility	Technology		
		$\hat{\gamma}_1$	$\hat{\alpha}_0$	$\hat{\alpha}_A$	$\hat{\alpha}_B$
DRRA ^b	100	115.3 ^c (65.1) [12.6, 1560.0]	3.01 ^c (3.01) [2.70, 3.35]	0.20 ^c (0.20) [0.11, 0.30]	0.60 ^c (0.60) [0.51, 0.69]
DRRA ^b	500	34.7 (32.8) [15.4, 72.7]	2.97 (2.97) [2.86, 3.06]	0.20 (0.20) [0.20, 0.21]	0.61 (0.61) [0.60, 0.62]
DRRA ^b	1,000	31.9 (31.1) [17.0, 56.2]	2.96 (2.96) [2.88, 3.03]	0.20 (0.20) [0.20, 0.21]	0.61 (0.61) [0.60, 0.62]
CRRA ^b	100	25.7 ^c (10.2) [1.2, 313.2]	3.01 ^c (3.01) [2.79, 3.26]	0.20 ^c (0.20) [0.13, 0.29]	0.60 ^c (0.60) [0.53, 0.67]
CRRA ^b	500	4.9 (4.6) [2.2, 8.8]	3.00 (3.00) [2.94, 3.07]	0.20 (0.20) [0.20, 0.20]	0.60 (0.60) [0.59, 0.61]
CRRA ^b	1,000	4.3 (4.2) [2.5, 6.3]	3.00 (3.00) [2.96, 3.04]	0.20 (0.20) [0.20, 0.20]	0.60 (0.60) [0.60, 0.61]
IRRA ^b	100	22.2 ^c (9.6) [0.6, 277.8]	3.00 ^c (3.00) [2.78, 3.24]	0.20 ^c (0.20) [0.12, 0.28]	0.60 ^c (0.60) [0.54, 0.67]
IRRA ^b	500	4.5 (4.2) [1.5, 8.4]	3.00 (3.00) [2.95, 3.06]	0.20 (0.20) [0.20, 0.20]	0.60 (0.60) [0.59, 0.61]
IRRA ^b	1,000	3.7 (3.7) [1.7, 6.0]	3.01 (3.01) [2.96, 3.05]	0.20 (0.20) [0.20, 0.20]	0.60 (0.60) [0.60, 0.61]

^aFor each parameter, the table reports the average, the median (within parentheses), and the 2.5% and 97.5% quantiles (within brackets).

^bDRRA, CRRA, and IRRA risk structures correspond to $[\gamma_0, \gamma_1]$ equal to $[-18.4, 1]$, $[0, 3]$, and $[43, 6]$, respectively.

^cMean is calculated using only observations falling between the 2.5% and 97.5% quantiles, to avoid influence of extreme observations.

Table 5. Parameter estimates for restricted utility estimation specification ($\gamma_0 = 0$), large variance scenario.

Risk Structure	Sample Size	Parameter Estimates ^a			
		Utility	Technology		
		$\hat{\gamma}_1$	$\hat{\alpha}_0$	$\hat{\alpha}_A$	$\hat{\alpha}_B$
DRRA ^b	100	5.8 (5.8) [−0.6, 12.1]	2.95 (2.93) [2.34, 3.64]	0.21 (0.21) [0.17, 0.24]	0.63 (0.63) [0.50, 0.72]
DRRA ^b	500	4.7 (4.8) [1.8, 7.2]	2.90 (2.88) [2.67, 3.16]	0.21 (0.21) [0.19, 0.22]	0.62 (0.63) [0.57, 0.66]
DRRA ^b	1,000	4.6 (4.7) [2.5, 6.7]	2.90 (2.89) [2.72, 3.09]	0.21 (0.21) [0.20, 0.22]	0.62 (0.62) [0.59, 0.65]
CRRA ^b	100	5.4 (5.1) [−0.6, 12.8]	3.05 (3.05) [2.54, 3.63]	0.20 (0.20) [0.17, 0.23]	0.61 (0.61) [0.50, 0.69]
CRRA ^b	500	3.4 (3.4) [1.3, 5.9]	3.02 (3.01) [2.84, 3.19]	0.20 (0.20) [0.19, 0.21]	0.60 (0.60) [0.57, 0.63]
CRRA ^b	1,000	3.2 (3.2) [1.8, 4.6]	3.01 (3.01) [2.89, 3.13]	0.20 (0.20) [0.19, 0.21]	0.60 (0.60) [0.58, 0.62]
IRRA ^b	100	5.2 (5.0) [−0.6, 12.2]	3.06 (3.06) [2.43, 3.69]	0.20 (0.20) [0.17, 0.23]	0.61 (0.61) [0.50, 0.70]
IRRA ^b	500	3.4 (3.3) [1.0, 5.7]	3.03 (3.03) [2.86, 3.21]	0.20 (0.20) [0.19, 0.21]	0.60 (0.60) [0.57, 0.63]
IRRA ^b	1,000	3.0 (3.0) [1.6, 4.4]	3.03 (3.03) [2.92, 3.14]	0.20 (0.20) [0.19, 0.20]	0.60 (0.60) [0.58, 0.62]

^aFor each parameter, the table reports the average, the median (within parentheses), and the 2.5% and 97.5% quantiles (within brackets).

^bDRRA, CRRA, and IRRA risk structures correspond to $[\gamma_0, \gamma_1]$ equal to $[-18.4, 1]$, $[0, 3]$, and $[43, 6]$, respectively.

Table 6. Parameter estimates for restricted utility estimation specification ($\gamma_0 = 0$), mixed variance scenario.

Risk Structure	Sample Size	Parameter Estimates ^a			
		Utility	Technology		
		$\hat{\gamma}_1$	$\hat{\alpha}_0$	$\hat{\alpha}_A$	$\hat{\alpha}_B$
DRRA ^b	100	7.8 (7.0) [0.8, 18.5]	3.01 (3.00) [2.65, 3.34]	0.20 (0.20) [0.19, 0.22]	0.61 (0.61) [0.56, 0.65]
DRRA ^b	500	5.4 (5.4) [2.7, 8.7]	2.93 (2.94) [2.72, 3.09]	0.20 (0.20) [0.20, 0.21]	0.61 (0.61) [0.59, 0.63]
DRRA ^b	1,000	5.1 (5.0) [3.1, 7.5]	2.90 (2.90) [2.75, 3.04]	0.20 (0.20) [0.20, 0.21]	0.61 (0.61) [0.59, 0.63]
CRRA ^b	100	5.5 (4.8) [0.7, 13.9]	3.04 (3.05) [2.78, 3.31]	0.20 (0.20) [0.19, 0.22]	0.60 (0.60) [0.57, 0.64]
CRRA ^b	500	3.5 (3.4) [2.0, 5.2]	3.02 (3.01) [2.94, 3.08]	0.20 (0.20) [0.20, 0.20]	0.60 (0.60) [0.59, 0.61]
CRRA ^b	1,000	3.3 (3.2) [2.3, 4.4]	3.01 (3.01) [2.94, 3.08]	0.20 (0.20) [0.20, 0.20]	0.60 (0.60) [0.59, 0.61]
IRRA ^b	100	5.6 (5.2) [0.4, 13.0]	3.05 (3.04) [2.77, 3.33]	0.20 (0.20) [0.19, 0.22]	0.60 (0.60) [0.56, 0.64]
IRRA ^b	500	3.6 (3.6) [2.0, 5.5]	3.02 (3.02) [2.92, 3.11]	0.20 (0.20) [0.20, 0.20]	0.60 (0.60) [0.59, 0.61]
IRRA ^b	1,000	3.3 (3.3) [2.2, 4.7]	3.01 (3.01) [2.95, 3.08]	0.20 (0.20) [0.20, 0.20]	0.60 (0.60) [0.59, 0.61]

^aFor each parameter, the table reports the average, the median (within parentheses), and the 2.5% and 97.5% quantiles (within brackets).

^bDRRA, CRRA, and IRRA risk structures correspond to $[\gamma_0, \gamma_1]$ equal to $[-18.4, 1]$, $[0, 3]$, and $[43, 6]$, respectively.

estimation system. However, at reasonable levels it enhances the ability to recover the coefficient of risk aversion because the larger the latter, the greater the impact of the variance on the optimal decisions. (Note that in the limit when there is no uncertainty, optimal decisions for risk-averse and risk-neutral agents are the same.)

Finally, the distributions for $\hat{\gamma}_1$ s under CRRA are much more similar to the distributions under IRRA than under DRRA. Further, the CRRA and IRRA estimates are more consistent with the true coefficient of local risk aversion than the DRRA estimates. The main explanation for this result is the constraint that costs $(r_{A,n} x_{A,n}^* + r_{B,n} x_{B,n}^*)$ do not exceed $\gamma_0 + W_{0,n}$, imposed at the optimization stage to calculate optimal input choices. Such constraint is binding in approximately 20%, 5%, and 0% of the observations under DRRA, CRRA, and IRRA, respectively. The effect of the constraint is to yield input choices smaller than they would be otherwise. Since the constrained optimization is not modeled for estimation purposes, the smaller input levels are in effect translated into a greater coefficient of relative risk aversion at the estimation stage.

As in the unrestricted utility estimation case, restricted utility estimates do not seem to be improved by heteroskedasticity (compare $\hat{\gamma}_1$ s in Table 6 with $\hat{\gamma}_1$ s in Tables 4 and 5). This finding lends additional support for the conclusion that heteroskedasticity is highly unlikely to be provide the information needed to identify the structure of risk aversion.

V. Concluding Remarks

Numerous studies have focused on the simultaneous estimation of technology and risk preferences from actual production data. Importantly, many of them have specifically aimed at uncovering the structure of risk aversion (i.e., the changes in absolute or relative risk aversion associated with changes in wealth) by estimating flexible utility functions. However, some works in finance argue that the structure of risk aversion significantly affects choices under uncertainty only when risks are very large and skewed. This means that, unless production risks are very large and skewed, recovering the structure of risk aversion from production data

should be difficult.

The present study explores the apparent disconnect between the production and finance literatures, by setting up a thought experiment calibrated to match the characteristics of risks faced by decision makers in a high-risk production activity (farming), and investigating whether the structure of risk aversion can be estimated with reasonable precision. Farming data are used for calibration not only because most of the studies concerned have employed data pertaining to production agriculture, but also because the high and skewed risks involved provide the most potential for the structure of risk aversion to considerably affect optimal decisions. The thought experiment is designed to facilitate estimation of the structure of risk aversion. Failure to reasonably estimate the structure of risk aversion under such “ideal” conditions can then be construed as strong evidence against the hypothesis that the risk aversion structure can be recovered from actual production data.

The study demonstrates that the simultaneous estimation of the two parameters of a standard hyperbolic absolute risk aversion (HARA) utility function and a three-parameter Cobb-Douglas production function yields extremely poor estimates of the utility parameters, even when samples comprise as many as 1000 observations. The 95% confidence intervals (CIs) for the utility parameters are very wide and usually fail to include the true values. Further, under certain parameterizations the 95% CIs lie entirely in the range corresponding to risk-loving (as opposed to risk-averse) attitudes.

When estimation is performed under a restricted utility specification, consisting of a one-parameter constant-relative risk aversion (CRRA) utility function, both utility and technology parameter estimates exhibit substantial improvements. Importantly, this is true even for the scenarios where CRRA utility only provides an approximation of the true two-parameter HARA utility originating the data. It is worth pointing out, however, that when samples have 100 observations, estimates of the single CRRA parameter are very poor (even if the true utility is CRRA). This finding is relevant for the literature, as studies that have simultaneously estimated technology and the structure of risk aversion often relied on real-world samples comprising

fewer than 100 observations.

In summary, the findings of the present study call into question the wisdom of attempting to estimate the structure of risk aversion simultaneously with technology using production data. In the purposely simple set up postulated here, allowing for a flexible utility specification yields utility estimates that bear no resemblance to the true parameters. Further, the resulting technology estimates are worse than those obtained under a restricted utility specification (even if the restricted utility is only an approximation of the actual utility generating the data). The findings also suggest that even in the restricted utility specification case, the quality of the utility parameters estimated from small samples (a common practice in econometric studies of production under risk) is very poor.

Overall, the results suggest that the emphasis on the estimation of flexible risk preferences in production studies has been misplaced, and that future efforts are likely to be more fruitfully employed elsewhere. Of course, the results from the present thought experiment need not apply to some real-world scenarios. Indeed, it may be the case that in specific situations there are gains to be made by estimating the structure of risk aversion simultaneously with technology. However, our findings strongly suggest that such gains are likely to be unwarranted in most instances, and that the burden of proof should fall on those claiming the contrary.

More generally, the method employed here can be applied in other circumstances where there is the suspicion that the data may provide too little information to successfully identify parameters and/or models of interest. This is important because assessing *ex ante* whether a particular estimation project is worth pursuing may prevent the waste of scarce resources gathering and analyzing data that are highly unlikely to yield the information researchers are looking for.

Appendix A: Estimation of the Probability Density Function of Initial Wealth

The initial wealth pdf (2.1) is estimated using the balanced panel employed by Hart and Lence, which contain annual initial wealth observations for 317 Iowa farms over the period 1991-1998. Since monetary data in the simulation model are standardized by setting unconditional mean prices equal to unity, the 2350 strictly positive initial wealth observations in the panel are multiplied by the ratio $7/73788$ to obtain standardized initial wealth values. In the standardizing ratio, the numerator 7 equals the approximate median costs (i.e., $\text{median}(\mathbf{r} \mathbf{x}^*)$) for the CRRA simulations (which do not depend on initial wealth data), whereas the denominator 73788 is the median operating expense (i.e., the real-world data analog of $\text{median}(\mathbf{r} \mathbf{x}^*)$) for the strictly positive initial wealth observations.

A strictly positive lower bound $\underline{W}_0$ is necessary to conduct DRRA simulations, as HARA utility (1.3) requires $\gamma_0 + W_0 > 0$, and DRRA entails $\gamma_0 < 0$. For the DRRA scenario, we fix $\gamma_1 = 1$ to clearly differentiate it from its CRRA counterpart and to provide a useful benchmark (see “Calibration” subsection), and set γ_0 at the value that yields a mean value of $R(W_0)$ equal to 3 for the estimated $Beta(\cdot)$ pdf for W_0 . Since such γ_0 value must satisfy the restriction $\underline{W}_0 > -\gamma_0$, we simultaneously calculate $\underline{W}_0$ and γ_0 by means of the following iterative procedure:

Step 1. Set iteration counter at $j = 1$.

Step 2. Obtain a sample of $2350 - 2j$ observations with lower bound $\underline{W}_0^{(j)}$ and upper bound $\overline{W}_0^{(j)}$, by discarding the smallest j and the largest j observations on standardized initial wealth (so that the sample median stays constant).

Step 3. Use the sample from Step 2 to estimate the standard beta pdf $Beta(\cdot)^{(j)}$ via maximum likelihood, by means of the “betafit” function in MATLAB version 7.0.4.365.

Step 4. Given $Beta(\cdot)^{(j)}$, $\underline{W}_0^{(j)}$, and $\overline{W}_0^{(j)}$, calculate $R^{(j)} = \text{mean}[R(W_0)|\gamma_0^{(j)} = 0.5 - \underline{W}_0^{(j)}, \gamma_1 = 1]$.

Step 5. If $R^{(j)} > 3$ (note that $R^{(j)} < R^{(j-1)} \forall j$), stop and fix $\underline{W}_0 = \underline{W}_0^{(j)}$, $\overline{W}_0 = \overline{W}_0^{(j)}$, $Beta(\cdot) = Beta(\cdot)^{(j)}$, and $\gamma_0 = \gamma_0^{(j)}$. Otherwise, set $j = j + 1$ and go back to Step 2.

In the present sample, iterations stop at $j = 720$. It should be clear that the only way to stop at a smaller j while having a mean value of $R(W_0) = 3$ is by adopting a DRRA parameterization more

similar to the CRRA scenario (i.e., by setting the DRRA $[\gamma_0, \gamma_1]$ closer to $[0, 3]$).

Appendix B: Estimation of the Probability Density Function of Output Shocks

The probability distribution of $\tilde{e}_y$ is derived from Hart and Lence's balanced panel, which has annual corn yields for 407 Iowa farms from 1991 through 1998. For each farm, standardized yields are calculated by dividing actual yields by the farm's average yield. Standardized yields are then pooled across all farms to obtain a sample of 3256 observations used to obtain the vector $[e_{y,0.5}, e_{y,1.5}, \dots, e_{y,98.5}, e_{y,99.5}]$, where $e_{y,q}$ is the q th quantile of standardized yields. The probability distribution of $\tilde{e}_y$ for the medium-variance scenario consists of $[e_{y,0.5}, e_{y,1.5}, \dots, e_{y,98.5}, e_{y,99.5}]$ with probabilities $[\pi_{y,0.5}^M, \pi_{y,1.5}^M, \dots, \pi_{y,98.5}^M, \pi_{y,99.5}^M] = [0.01, 0.01, \dots, 0.01, 0.01]$.

For the low- and high-variance scenarios, Prelec's probability weighting function (Prelec) is used to assign the probabilities $\pi_{y,q}^L$ and $\pi_{y,q}^H$ corresponding to $e_{y,q}$. For the low-variance case, the distribution of $\tilde{e}_y$ is given by $e_{y,q}$ with probability $\pi_{y,q}^L \equiv \pi(q + 0.5; \phi_1^L, \phi_2^L) - \pi(q - 0.5; \phi_1^L, \phi_2^L)$ for $\pi(q; \phi_1, \phi_2) \equiv \exp\{-[-\ln(q/100)/\phi_1]^{1/\phi_2}\}$, $[\phi_1^L, \phi_2^L] = [1.02, 0.51]$, and $q = 0.5, 1.5, \dots, 98.5, 99.5$. Values for ϕ_1^L and ϕ_2^L are derived by trial-and-error so as to yield the same mean but a standard deviation 50% smaller than the standard deviation under the medium-variance scenario. Analogously, the distribution of $\tilde{e}_y$ under the high-variance scenario is $e_{y,q}$ with probabilities $\pi_{y,q}^H \equiv \pi(q + 0.5; \phi_1^H, \phi_2^H) - \pi(q - 0.5; \phi_1^H, \phi_2^H)$ for $[\phi_1^H, \phi_2^H] = [0.93, 1.73]$. Compared to the medium-variance probabilities $\pi_{y,q}^M$, the low-variance (high-variance) probabilities shift weight from the extremes (middle) of vector $[e_{y,0.5}, e_{y,1.5}, \dots, e_{y,98.5}, e_{y,99.5}]$ to its middle (extremes), so as to reduce (increase) the standard deviation by 50% while maintaining the mean unchanged at 1.¹⁰

Appendix C: Estimation of the Probability Density Function of Output Prices

Expression (2.2) is an approximation based on the regression estimates reported in tables A.1 and

¹⁰Note that $\phi_{y,q}^M = \pi(q + 0.5; \phi_1^M, \phi_2^M) - \pi(q - 0.5; \phi_1^M, \phi_2^M)$ for $[\phi_1^M, \phi_2^M] = [1, 1]$.

A.2. The former table shows results for regressions employing U.S. aggregate yield shocks, whereas the latter displays analogous results but for yield shocks corresponding to the Hart and Lence farm-level data.

Table A1. Price regressions for U.S. aggregate data, 1970-2005.

Crop	Regression	Regression	R^2
		Std. Error	
Corn	$\ln(p_t) = 47 - 0.0248 t + 0.48 \ln(p_{t-1}) - 0.85 \ln e_{y,t} + e_{p,t}$ (11) (0.0061) (0.13) (0.26)	0.156	0.914
Wheat	$\ln(p_t) = 42 - 0.0219 t + 0.47 \ln(p_{t-1}) - 1.09 \ln e_{y,t} + e_{p,t}$ (11) (0.0058) (0.13) (0.46)	0.178	0.870
Soybeans	$\ln(p_t) = 48 - 0.0248 t + 0.45 \ln(p_{t-1}) - 1.06 \ln e_{y,t} + e_{p,t}$ (11) (0.0058) (0.12) (0.31)	0.146	0.919

Note: Variable p_t is the U.S. season-average price received by farmers in marketing year t based on monthly prices weighted by monthly marketings (source: U.S. Department of Agriculture), deflated by the corresponding U.S. Consumer Price Index, all items, U.S. city average, not seasonally adjusted (1982-84 = 100) (source: Bureau of Labor Statistics). Variable $\ln e_{y,t}$ is the ordinary least-squares residual of the regression of $\ln(\text{yield}_t)$ on a constant and t , where yield_t is the U.S. yield per acre in marketing year t (source: U.S. Department of Agriculture). Each regression comprises 35 observations. Standard errors are shown within parentheses below the respective coefficients.

Table A2. Price regressions for Iowa farm-level yield data, 1991-1998.

Crop	Regression	Regression	R^2
		Std. Error	
Corn	$\ln(p_t) = 36.4 - 0.0197 t + 0.303 \ln(p_{t-1}) - 0.248 \ln(e_{y,f,t}) + e_{p,f,t}$ (2.1) (0.0011) (0.024) (0.011)	0.135	0.266
Soybeans	$\ln(p_t) = 30.3 - 0.01657 t + 0.145 \ln(p_{t-1}) - 0.1532 \ln(e_{y,f,t}) + e_{p,f,t}$ (1.8) (0.00091) (0.029) (0.0093)	0.114	0.199

Note: Variable p_t is the U.S. season-average price received by farmers in marketing year t based on monthly prices weighted by monthly marketings (source: U.S. Department of Agriculture), deflated by the corresponding U.S. Consumer Price Index, all items, U.S. city average, not seasonally adjusted (1982-84 = 100) (source: Bureau of Labor Statistics). Variable $e_{y,f,t}$ is farm f 's crop yield in marketing year t divided by farm f 's average crop yield over 1991-1998 (source: Hart and Lence dataset). Each regression comprises 3,200 observations. Standard errors are shown within parentheses below the respective coefficients.

References

- Altuğ, S., and P. Labadie. *Dynamic Choice and Asset Markets*. San Diego, CA: Academic Press, 1994.
- Antle, J. M. "Econometric Estimation of Producers' Risk Attitudes." *American Journal of Agricultural Economics* 69(1987):509-522.
- Antle, J. M. "Nonstructural Risk Attitude Estimation." *American Journal of Agricultural Economics* 71(1989):774-784.
- Binswanger, H. P. "Attitudes toward Risk: Experimental Measurement In Rural India." *American Journal of Agricultural Economics* 62(1980):395-407.
- Brink, L., and B. McCarl. "The Tradeoff Between Expected Return and Risk Among Cornbelt Farmers." *American Journal of Agricultural Economics* 60(1978):258-263.
- Černý, A. *Mathematical Techniques in Finance: Tools for Incomplete Markets*. Princeton, NJ: Princeton University Press, 2004.
- Chavas, J. P., and M. T. Holt. "Economic Behavior Under Uncertainty: A Joint Analysis of Risk Preferences and Technology." *Review of Economics and Statistics* 78(1996):329-335.
- Cliff, M. T. "GMM and MINZ Program Libraries for MATLAB." Unpublished manuscript, Krannert Graduate School of Management, Purdue University, Indiana, 2003. Available at <http://mcliff.cob.vt.edu/progs/gmmdoc.pdf> (accessed May 20, 2007).
- Coyle, B. T. "Risk Aversion and Yield Uncertainty in Duality Models of Production: A Mean-Variance Approach." *American Journal of Agricultural Economics* 81(1999):553-567.
- Davidson, R., and J. G. Mackinnon. *Estimation and Inference in Econometrics*. New York: Oxford University Press, 1993.
- Geweke, J. "Antithetic Acceleration of Monte Carlo Integration in Bayesian Inference." *Journal of Econometrics* 38(1988): 73-90.
- Gollier, C. *The Economics of Risk and Time*. Cambridge, MA: The MIT Press, 2001.
- Hansen, L. P. "Large Sample Properties of Generalized Method of Moment Estimators." *Econometrica* 50(1982):1029-1054.

- Hansen, L. P., and K. J. Singleton. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica* 50(1982):1269-1286.
- Hart, C. E., and S. H. Lence. "Financial Constraints and Farm Investment: A Bayesian Examination." *Journal of Business and Economic Statistics* 22(2004):51-63.
- Jagannathan, R., G. Skoulakis, and Z. Wang. "Generalized Method of Moments: Applications in Finance." *Journal of Business and Economic Statistics* 20(2002):470-481.
- Just, R. E., and R. D. Pope, Editors. *A Comprehensive Assessment of the Role of Risk in U.S. Agriculture*. Boston, MA: Kluwer Academic Publishers, 2002.
- Just, R. E., and R. D. Pope. "Stochastic Specification of Production Functions and Economic Implications." *Journal of Econometrics* 7(1978):67-89.
- Kallberg, J. G., and W. T. Ziemba. "Comparison of Alternative Utility Functions in Portfolio Selection Problems." *Management Science* 29(1983):1257-1276.
- Kocherlakota, N. "The Equity Premium: It's Still A Puzzle." *Journal of Economic Literature* 34(1996):42-71.
- Kumbhakar, S. C. "Risk Preference and Productivity Measurement under Output Price Uncertainty." *Empirical Economics* 27(2002a):461-472.
- Kumbhakar, S. C. "Specification and Estimation of Production Risk, Risk Preferences and Technical Efficiency." *American Journal of Agricultural Economics* 84(2002b):386-390.
- Love, H. A., and S. T. Buccola. "Joint Risk Preference-Technology Estimation with a Primal System." *American Journal of Agricultural Economics* 73(1991):765-774.
- Miranda, M., and P. Fackler. *Applied Computational Economics and Finance*. Cambridge, MA: MIT Press, 2002.
- Ogaki, M. "Generalized Method of Moments: Econometric Applications." *Handbook of Statistics*, Volume 11. S. Maddala, C. R. Rao, and H. D. Vinod, eds., pp. 455-488. New York: Elsevier Science Publishers, 1993.
- Prelec, D. "The Probability Weighting Function." *Econometrica* 66(1998):497-528.
- Robison, L. J. "An Appraisal of Expected Utility Hypothesis Tests Constructed from

Responses to Hypothetical Questions and Experimental Choices.” *American Journal of Agricultural Economics* 64(1982):367-375.

Saha, A. “Risk Preference Estimation in the Nonlinear Mean Standard Deviation Approach.” *Economic Inquiry* 35(1997):770-782.

Saha, A., C. R. Shumway, and H. Talpaz. “Joint Estimation of Risk Preference Structure and Technology Using Expo-Power Utility.” *American Journal of Agricultural Economics* 76(1994):173-184.

Simmons, R. L., and C. Pomareda. “Equilibrium Quantity and Timing of Mexican Vegetable Exports.” *American Journal of Agricultural Economics* 57(1975):472-479.