HOW MUCH CAN WE LEARN ABOUT PRODUCERS' UTILITY FUNCTIONS FROM THEIR PRODUCTION DATA?

Sergio H. Lence*

Abstract

A thought experiment is designed to investigate 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 from agricultural production data. 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 slightly worsen technology parameter estimates. Results also indicate that even under a restricted utility specification, utility parameter estimates are biased. Further, their quality is much worse when shocks are not large or samples are small.

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

JEL Codes: C13, D24, D81, Q12.

Paper presented at Regional Research Committtee NC-1014 Meeting,
2008 Agricultural and Rural Finance Markets in Transition, September 25-26, 2008, Kansas City, Missouri.

September 1, 2008

Corresponding author: Sergio H. Lence
Department of Economics
368E Heady Hall
Iowa State University
Ames, IA 50011-1070
Phone: 515-294-6644
Email: shlence@iastate.edu

*Sergio Lence is Professor and Marlin Cole Chair of International Agricultural Economics, Department of Economics, Iowa State University, Ames.

HOW MUCH CAN WE LEARN ABOUT PRODUCERS' UTILITY FUNCTIONS FROM THEIR PRODUCTION DATA?

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 1980, Robison 1982), or by analyzing observed production and/or investment choices (e.g., Brink and McCarl 1978, Antle 1987, Antle 1989, Love and Buccola 1991, Saha, Shumway and Talpaz 1994, Chavas and Holt 1996, Kumbhakar 2002a).

The seminal studies estimating risk preferences from actual production and/or investment decisions have focused on the level of risk aversion, by estimating risk preferences separately from technology (e.g., Simmons and Pomareda 1975, Brink and McCarl 1978) 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 (e.g., Love and Buccola 1991, Coyle 1999), as doing so can increase estimation efficiency and may avoid inconsistency problems, 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 (1994), 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 (Chavas and Holt 1996, Saha 1997, Bar-Shira, Just and Zilberman 1997, Kumbhakar 2001, Kumbhakar 2002a, Kumbhakar 2002b, Kumbhakar and Tveterås 2003, Isik and Khanna 2003, Abdulkadri, Langemeier and Featherstone 2003).

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 2001). Importantly, however, Kallberg and Ziemba (1983, p. 1257) concluded that "… utility functions having different functional forms and parameter values but 'similar' absolute risk aversion indices have 'similar' optimal portfolios." They defined the index of absolute risk aversion of agent *i* under end-of-period wealth distribution *d* as the expected coefficient of absolute risk aversion corresponding to *i* under *d*. Kallberg and Ziemba's (1983) conclusions must be qualified by the fact that their study assumed multivariate normally distributed returns, a case relatively favorable to finding similar choices across alternative utility functions (Černý 2004). More recently, Černý (2004) argued that, except for investments involving very large and skewed risks, agents with similar values of relative risk aversion, evaluated at their initial wealth levels, make almost identical portfolio decisions, regardless of their risk aversion structures.

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 (1983) and Černý (2004) 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 sources of information too weak 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 a large number of simulated data sets from widely-used utility and production functions, and subsequently employs 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. If anything, a flexible utility parameterization slightly worsens 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 at an early stage.

I. Model

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

$$(1) \qquad \boldsymbol{x}^* = argmax_{\boldsymbol{x}} E_W\{U[\widetilde{W}(\boldsymbol{x})]\},$$

where $\boldsymbol{x}^*$ denotes the vector of optimal input amounts, $E_W(\cdot)$ is the expectation operator with respect to random variable $\widetilde{W}$, $U(\cdot)$ is the producer's utility function, and $\widetilde{W}(\boldsymbol{x})$ 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,

(2) $\widetilde{W}(\boldsymbol{x}) \equiv \tilde{p}\ \tilde{y}(\boldsymbol{x}) - \boldsymbol{r}\,\boldsymbol{x} + W_0,$

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

For present purposes, model (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 $\tilde{y}(\boldsymbol{x})$, 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,

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

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 2004), and must be strictly positive if (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 2001).

Importantly, quite different structures of risk aversion can be obtained under appropriate parameterizations of (3). To see this, note that the HARA coefficient of relative risk aversion is given by

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

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 parameterize (3) so as to test Černý's (2004) claim that optimal decisions are essentially the same, regardless of whether the agent's utility is characterized by DRRA, CRRA or IRRA, except for decisions involving very large and skewed risks. More specifically, Černý (2004) labels $R(W)$ as the agent's local relative risk aversion, and argues that the key determinant for optimal portfolio decisions is $R(W_0)$ (i.e., the local risk aversion evaluated at the "safe" wealth level W_0). Loosely speaking, this means that individuals with similar coefficients of relative risk aversion evaluated at initial wealth will behave similarly toward risk.

I.2. Production Technology

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

$$(5) \qquad y(x_A, x_B; \tilde{e}_y) = \alpha_0 \; x_A^{\alpha_A} \; x_B^{\alpha_B} \; \tilde{e}_y ,$$

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. Major reasons for adopting technology (5) are its simplicity and the fact that the Cobb-Douglas technology is arguably the most widely-used production function in economic analysis. Examples of studies employing two variable inputs are Saha, Shumway, and Talpaz (1994) and Saha (1997), who analyzed wheat farms in Kansas with capital and materials as inputs, and Kumbhakar and Tveterås (2003), who studied salmon farms in Norway using feed and labor as inputs.[2] More complex technologies, or a Cobb-Douglas production

[1]The coefficient of absolute risk aversion for HARA utility (3) is $A(W) = \gamma_1/(\gamma_0 + W)$. 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 2001, 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 2001, p. 238).

[2]Kumbhakar (2001, 2002a, 2002b) examined Norwegian salmon farms with a production function consisting of only one variable input (labor). Isik and Khanna (2003) studied the adoption of site-specific technologies by farmers in

function involving additional inputs, would require the estimation of additional technology parameters, thereby posing stronger challenges for the estimation of utility parameters γ_0 and γ_1.

The present analysis was 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 (5), parameterized with $\alpha_0 = 3$, $\alpha_A = 0.2$, $\alpha_B = 0.6$, $\beta_A = 0.06$, and $\beta_B = -0.03$ following Saha, Shumway, and Talpaz (1994), and with $log(\tilde{e}_y)$ distributed as described in Appendix B. (Note that $\beta_B < 0$ means that input *B* is risk-reducing). To save space, results for the Just-Pope specification are omitted, as they led to the same conclusions as the results for the simpler production function (5). Importantly, technology (5) also seems better suited than the Just-Pope setup for the present purposes. This is true because the latter is more complex and therefore more prone to the critique that it hinders the chances for risk preference identification.[3]

II. Simulation Design, Calibration, and Simulated Data Generation

Given the model set-up described 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 $\tilde{p}$ - often found in production agriculture. The null hypothesis is that the typical data sets used in the empirical production literature do allow estimation of the structure of risk aversion (i.e., the simultaneous identification of γ_0 and γ_1). The validity of such a 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

---

Illinois resorting to three variable inputs (nitrogen, phosphorous, and potassium). Abdulkadri, Langemeier, and Featherstone (2003) investigated wheat farms in Kansas with three variable inputs (materials, capital, and labor).

[3]Conditional on knowing the true values of β_A and β_B, the Just-Pope specification offers better opportunities for identifying the utility parameters. However, such advantage does not necessarily extend to the case where β_A and β_B are unknown, because in the latter instance, estimation is more involved than for (5). Just-Pope simulations also require more care, because each optimal input bundle must be checked to ensure that the additive shock does not lead to a negative amount of output with strictly positive probability. In our Just-Pope simulations, observations where input bundles violated the nonnegative output restriction were discarded.

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) through (5)), all observations are generated from the same utility and technology parameters,[4] 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. By construction, our simulated samples are free of selection bias (i.e., the bias due to the endogeneity of qualitative characteristics of inputs in crop choice). This is important, because inferences regarding risk behavior may be incorrect if selection bias is present but not accounted for (Koundouri and Nauges 2005). Further, all variables are unambiguously defined, so that no issues arise regarding the estimation and/or interpretation of the structure of risk aversion (see Meyer and Meyer 2005).

Consistent with the stated goal of stacking the odds in favor of identifying the underlying preferences, end-of-period wealth does not include income from other activities, or payments from government programs aimed at reducing yield, price, or revenue uncertainty. Including additional activities would reduce the chances for identification, as it would involve more unknown parameters (e.g., moments of the joint probability distributions of the associated shocks, and technological parameters if the additional activities entailed production).[5] Government payments would also hinder the identification of preferences, by making risk less important. In the limiting case of no yield uncertainty and government payments eliminating price uncertainty,

[4]Arguably, the issue of preference heterogeneity and reliable identification of individual preferences is just as important as reliable estimation of flexible functional forms. However, estimating 2 N individual utility parameters using a sample of size N will be at least as problematic as estimating only two utility parameters from the same sample, even if one includes individual characteristics as additional explanatory variables.

[5]Omission of non-agricultural production activities is a standard assumption in the empirical literature (e.g., Saha, Shumway, and Talpaz 1994; Saha 1997; Isik and Khanna 2003; Abdulkadri, Langemeier, and Featherstone 2003; Kumbhakar 2001, 2002a, and 2002b; and Kumbhakar and Tveterås 2003).

it would be impossible to identify the parameters of the utility function because farmer risk preferences would be irrelevant.

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 two million vectors, where the nth vector contains observations corresponding to the nth production decision $\boldsymbol{v}_n \equiv [W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}, x^*_{A,n}, x^*_{B,n}, p_n, y_n]$. Vector $\boldsymbol{v}_n$ comprises the relevant "exogenous" variables known at the time of making the nth 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 nth 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 (4) with $[\gamma_0, \gamma_1]$ equal to $[-18.4, 1]$, $[0, 3]$ and $[43, 6]$, respectively. Parameters γ_0 and γ_1 are purposefully 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 1996; Gollier 2001, pp. 31 and 289; Meyer and Meyer 2005, table 2 on p.

260). 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 (3). Gollier (2001, 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 three under the assumed probability density function (pdf) of initial wealth (see (6) 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 (2004). 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 a coefficient of relative risk aversion $R(W) = 3$.[6]

Since monetary units can be arbitrarily chosen, all prices are scaled by setting their unconditional means equal to one. The prices known to the agent when making the nth decision (i.e., $p_{0,n}$, $r_{A,n}$, are $r_{B,n}$) 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 (5) is parameterized with $\alpha_A = 0.2$ and $\alpha_B = 0.6$, so as to closely match the estimates obtained by Saha, Shumway, and Talpaz (1994) and Saha (1997) for capital

[6]Kallberg and Ziemba's (1983) study suggests that an alternative approach would be to use different utility function parameterizations yielding the same value for the index of absolute risk aversion. The main shortcoming of such an approach for present purposes, however, is that the index of absolute risk aversion is endogenous, in the sense that it cannot be defined without reference to the agent's choices. To see this, note that Kallberg and Ziemba's (1983) index requires the evaluation of the agent's coefficient of absolute risk aversion at his end-of-period wealth, but the latter is determined by the agent's decisions (which depend on his risk aversion). In contrast, computation of Černý's (2004) index of local risk aversion requires knowledge about the agent's beginning wealth (as opposed to end-of-period wealth), and this is predetermined.

and materials, respectively. The values chosen for α_A and α_B imply 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 three. 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.

For the baseline scenarios, simulated observations on initial wealth are generated from

$$W_0 = 18.9 + 69.2\, z, \tag{6}$$

where z is a 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 (6) is twofold. First, (6) provides a distribution of initial wealth consistent with real-world data under the advocated scaling of prices and quantities (see Appendix A for details). Second, as pointed out earlier, the mean value of $R(W_0)$ corresponding to (6) equals three under the adopted DRRA, CRRA, and IRRA parameterizations.[7] The corresponding average elasticities of $R(W_0)$ with respect to initial wealth are –2, 0, and 0.5, respectively. The absolute values of the DRRA and IRRA average elasticities are large (see Meyer and Meyer 2005, p. 260), so as to facilitate the estimation of the structure of risk aversion.

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

[7]The [5%, 50%, 95%] quantiles of $R(W_0)$ equal [9.26, 1.70, 1.30] for DRRA, [3, 3, 3] for CRRA, and [1.94, 3.06, 3.92] for IRRA.

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.[8] Having low- and high-variance scenarios allows us to analyze the sensitivity of the parameter estimates to the uncertainty in output shocks. This is important, among other things, to alleviate concerns that the medium-variance scenario may not adequately represent the true output uncertainty faced by farmers. A possible reason for such concern is that our estimation of the probability distribution of the output shocks assumes that it is exogenous, i.e., that the variability in farm-level corn yields in Iowa is not affected by the risk preferences of the surveyed farmers.

The levels of output uncertainty employed for the simulations are representative of farm-level production risks often found in the U.S. This is true because for Iowa corn, the case used to estimate the baseline scenario, the average rate for 65% actual production history crop insurance is 4.1 cents per dollar of liability (Risk Management Agency). By comparison, the analogous rates in cents per dollar of liability for other agricultural products in their respective main producing states are as follows: almonds 7.4, apples 6.0, avocados 5.8, cotton 5.5, dry beans 8.6, forage 6.1, grapefruit 4.3, grapes 7.6, lemons 3.0, macadamia nuts 1.4, peaches 4.8, peanuts 8.0, pears 5.6, potatoes 5.0, rice 5.6, soybeans 3.5, sugar beets 5.4, sugar cane 3.2, sweet oranges 6.7, tobacco 3.7, tomatoes 4.9, and walnuts 5.3 (Risk Management Agency).

Finally, the postulated random generating process for end-of-period crop prices is

$$(7) \qquad ln(\tilde{p}) = \mu_p + 0.5\, ln(p_0) - 0.3\, ln(\tilde{e}_y) + \tilde{e}_p,$$

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

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

low-, medium-, and high-variance scenarios are 0.1, 0.2, and 0.3, respectively.[9] 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 (7) is based on regressions employing historical price and yield-shock data, which are reported in Appendix C.

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

Conceptually, vector $\boldsymbol{v}_n \equiv [W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}, x^*_{A,n}, x^*_{B,n}, p_n, y_n]$, comprising simulated data for the nth production decision, is calculated in three steps. In the first step, a random draw from (6) 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 $[W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}]$, 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 probability distribution of $\tilde{e}_y$, and substituting it, along with optimal inputs ($x^*_{A,n}$ and $x^*_{B,n}$), into production function (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 (7).

To ensure that results are as comparable as possible for the alternative (DRRA, CRRA, IRRA) × (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 nth production decision. A different procedure is required for end-of-period shocks $e_{y,n}$ and $e_{p,n}$, however,

[9]Again, 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. The postulated levels of price uncertainty represent the price risks associated with many agricultural products in the U.S., including broilers, calves, corn, cotton, dry beans, eggs, grapefruit, hogs, lambs, lettuce, onions, oranges, pears, potatoes, rice, sorghum, soybeans, tomatoes, and wheat (Harwood et al. 1999, p. 11).

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.[10]

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 nth production decision. With this notation, the optimization problem consists of maximizing (8) with respect to $x_{A,n}$ and $x_{B,n}$:

$$(8) \qquad E_{e_y,e_p}[u(x_{A,n}, x_{B,n}, \tilde{e}_y, \tilde{e}_p; W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n})] \equiv \iint u(x_{A,n}, x_{B,n}, e_y, e_p; \cdot)\, f_{y,p}(e_y, e_p)\, de_y\, de_p,$$

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 (8) has no analytical solution, its computation requires numerical quadrature methods (Miranda and Fackler 2002, ch. 4). Hence, the objective function employed for numerical optimization is the one shown on the right-hand-side of (9):

$$(9) \qquad E_{e_y,e_p}[u(x_{A,n}, x_{B,n}, \tilde{e}_y, \tilde{e}_p; W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n})] \equiv \Sigma_s \Sigma_q\, u(x_{A,n}, x_{B,n}, e_{y,q}, e_{p,s}; \cdot)\, \pi_{y,q}\, \pi_{p,s},$$

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

[10]To illustrate this point, consider the generation of the nth 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_{pn}, 0.2 $z_{p,n}$, and 0.3 $z_{p,n}$, respectively (see explanation of (7) above).

weights for price shocks are determined by a 10-point Gaussian quadrature.[11]

Numerical optimization of the objective function on the right-hand side of (9) 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 one million $\boldsymbol{v}_n$ vectors (i.e., data for one million production decisions) for each of the nine basic scenarios analyzed. 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 1988). More specifically, an additional set of one million antithetic decision vectors $\boldsymbol{v}_{n+}$ $\equiv [W_{0,n+}, p_{0,n+}, r_{A,n+}, r_{B,n+}, x^*_{A,n+}, x^*_{B,n+}, p_{n+}, y_{n+}]$ 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 (7) using $p_{0,n+}$ and the antithetic replications of $e_{y,n}$ and

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

$e_{p,n}$, whereas y_{n+} is obtained by substituting $x^*_{A,n+}$, $x^*_{B,n+}$ and the antithetic replication of $e_{y,n}$ into (5).

The simulated data for the two million production decisions are then grouped into 20,000 (4,000; 2,000) samples of 100 (500; 1,000) production decisions each. In turn, each of the simulated samples is used to perform an econometric estimation of the utility and technology parameters, resulting in 20,000 (4,000; 2,000) estimates for each parameter from the samples with 100 (500; 1,000) observations, respectively.[12]

The sizes of the simulated samples were chosen to cover the sample magnitudes typically used in the literature. For the studies of the structure of risk aversion cited in the introduction, sample sizes ranged from a low of 32 observations in Chavas and Holt (1996), to a high of 1010 observations in Bar-Shira, Just and Zilberman (1997). Saha (1997), Saha, Shumway and Talpaz (1994), Isik and Khanna (2003), Kumbhakar and Tveterås (2003), Kumbhakar (2002a), Kumbhakar (2002b), and Kumbhakar (2001) used samples of 60, 60, 100, 224, 224, 226 and 300 observations, respectively. The estimation performed by Abdulkadri, Langemeier and Featherstone (2003) relied on three separate samples comprised of 125, 240 and 255 observations per sample.

II.5. "Mixed-Variance" Scenarios

In addition to the aforementioned nine (DRRA, CRRA, IRRA) × (low-, medium-, high-variance) scenarios, a set of three "mixed-variance" scenarios are generated for DRRA, CRRA, and IRRA preferences. In the mixed-variance settings, the nth production decision vector is augmented by incorporating variable σ_n (i.e., $\boldsymbol{v}_n \equiv [W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}, \sigma_n, x^*_{A,n}, x^*_{B,n}, p_n, y_n]$), where σ_n equals 0.5, 1, or 1.5 depending on whether the distribution of price and output shocks for the nth 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

[12]Data corresponding to two million production decisions can be generated without resorting to antithetic replications. However, the latter allow us to obtain more accurate estimates of the distributions of interest.

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

II.6. Scenarios with a Wider Distribution of Initial Wealth

The distribution of initial wealth is a major risk structure identifier. To analyze whether an initial wealth distribution substantially wider than (6) enhances the estimation of the utility parameters, simulations are also performed using observations on initial wealth generated from a uniform distribution with the same lower bound as (6) (i.e., 18.9), but a much larger upper bound of 500 (versus 88.1 for (6)). In this instance, initial wealth has a mean of 259.45 and a standard deviation of 138.88, and the upper bound divided by the lower bound yields a ratio of 26.6. To see how wide this uniform distribution is, note that for the distribution of family net worth for all U.S. farm households in 2004, the ratio of the 95% quantile (= $2.36 million) to the 10% quantile (= $150 thousand) is only 15.7 (Economic Research Service 2008).

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 (ML) estimation would require a linear approximation to render the

problem tractable (Jagannathan, Skoulakis and Wang 2002). 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 ML, as ML estimates may not be consistent when the distribution of the observable variables is misspecified (Hansen and Singleton 1982). For such reasons, the generalized method of moments (GMM) is adopted as the baseline estimation method.

Useful references for the theory underlying GMM and its numerous applications include Hansen (1982), Davidson and Mackinnon (1993, ch. 17), Ogaki (1993) and Cliff (2003). The present optimal decision making framework lends itself nicely to application of GMM. Succinctly, estimation is based on a system of three regression equations corresponding to the logarithmic transformation of the production function (5) (i.e., (10)) and the first-order conditions (FOCs) for optimization of (8) (i.e., (11)):[13]

$$(10) \qquad \varepsilon_{y,n}(\boldsymbol{\alpha}) \equiv log(y_n) - log(\alpha_0) - \alpha_A\, log(\, x^*_{A,n}\,) - \alpha_B\, log(\, x^*_{B,n}\,),$$

$$(11) \qquad \varepsilon_{j,n}(\alpha_j, \boldsymbol{\gamma}) \equiv \left(\frac{\gamma_0 + p_n\, y_n - r_{A,n}\, x^*_{A,n} - r_{B,n}\, x^*_{B,n} + W_{0,n}}{\gamma_0 + W_{0,n}} \right)^{-\gamma_1} \quad (\, p_n \alpha_j x^{*-1}_{j,n} y_n \quad - r_{j,n}),$$

for $j = A$ and B, and $n = 1, ..., N$, where N is the sample size. To compare the present model to the existing empirical work, note that (10) and (11) above are analogous, respectively, to estimation equations (16) and (14) in Saha, Shumway, and Talpaz (1994); (11) and (10) in Isik and Khanna

[13]Estimation was also performed employing the untransformed technology equation (i.e., $\varepsilon_{y,n}(\boldsymbol{\alpha}) \equiv y_n - \alpha_0 \quad x^{*\alpha_A}_{A,n}$ $x^{*\alpha_B}_{B,n}$ instead of (10)), but the quality of the resulting parameter estimates was slightly lower. The term $[(\gamma_0 + p_n\, y_n - r_{A,n} \quad x^*_{A,n} \quad - r_{B,n} \quad x^*_{B,n} \quad + W_0)/(\gamma_0 + W_0)]^{-\gamma_1}$ in FOC (11) represents the stochastic discount factor (Pennacchi 2007, Ch. 4). In the GMM literature, FOC equations are typically estimated by means of a stochastic factor specification (e.g., Altuğ and Labadie 1994, ch. 3; Hansen and Singleton 1982; Ogaki 1993). In the present simulations, using the stochastic discount factor instead of $(\gamma_0 + p_n\, y_n - r_{A,n} \quad x^*_{A,n} \quad - r_{B,n} \quad x^*_{B,n} \quad + W_0)^{-\gamma_1}$ in (11) yields substantial improvements in estimation.

(2003); (17) and (19) in Abdulkadri, Langemeier, and Featherstone (2003); (1) and (13) in Kumbhakar (2001); (6a) and (7a) in Kumbhakar (2002a); (8) and (12), and (18) and (21) in Kumbhakar (2002b); (2) and (3a)-(3b) in Kumbhakar and Tveterås (2003); and the equation alluded to on p. 774 and (8) in Saha (1997).

The term $\varepsilon_{y,n}$ provides no information to help identify utility parameters, and $\varepsilon_{A,n}$ ($\varepsilon_{B,n}$) contains no information for identifying 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 $\boldsymbol{\iota}_n \equiv [1, W_{0,n}, p_{0,n}, r_{A,n}, r_{B,n}, x^*_{A,n}, x^*_{B,n}]'$, comprising a constant (standardized to unity) and the variables known at the time of decision making. For the mixed-variance scenario, vector $\boldsymbol{\iota}_n$ is augmented by the variance variable σ_n.[14]

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 1994, 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 (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 $\boldsymbol{g}_N(\boldsymbol{\alpha}, \boldsymbol{\gamma}) \equiv 1/N \, \Sigma_n \, [\varepsilon_{y,n}(\boldsymbol{\alpha}) \; \varepsilon_{A,n}(\alpha_B, \boldsymbol{\gamma}) \; \varepsilon_{B,n}(\alpha_B, \boldsymbol{\gamma})]' \otimes \boldsymbol{\iota}_n$, where $\otimes$ denotes the Kronecker product. Since parameters $[\boldsymbol{\alpha}, \boldsymbol{\gamma}]$ are the only elements of $\boldsymbol{g}_N(\cdot)$ unknown to the econometrician, the GMM estimates $[\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\gamma}}]$ are chosen so as to render $\boldsymbol{g}_N(\cdot)$ as close to zero as possible by minimizing a quadratic form in $\boldsymbol{g}_N(\cdot)$ with respect to the unknown parameters:

$$(3.3) \qquad [\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\gamma}}] = argmax_{[\boldsymbol{\alpha},\boldsymbol{\gamma}]} [\boldsymbol{g}_N(\boldsymbol{\alpha}, \boldsymbol{\gamma})' \, \boldsymbol{V}_N \, \boldsymbol{g}_N(\boldsymbol{\alpha}, \boldsymbol{\gamma})].$$

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

In (3.3), $\boldsymbol{V}_N$ is a positive definite weighting matrix which converges in probability to a positive definite matrix $\boldsymbol{V}_0$. 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 2003), 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 $\boldsymbol{V}_N = (\mathbf{1}_k \otimes \boldsymbol{\iota}'\boldsymbol{\iota})^{-1}$, where $\mathbf{1}_k$ is the $(k \times k)$ identity matrix, k is the number of regression equations (three in the present application), and $\boldsymbol{\iota} = [\iota_1, ..., \iota_N]'$ (Cliff 2003). In addition, to ensure that $\hat{\gamma}_0$ does not violate the constraint $(\hat{\gamma}_0 + W) > 0$ associated with HARA utility (3), for estimation purposes, we impose the restriction that $\hat{\gamma}_0 > 0.001 - min(W_{1,n},\ n = 1,\ ...,\ N)$, where $W_{1,n} \equiv p_n y_n - r_{A,n} x^*_{A,n} - r_{B,n} x^*_{B,n} - W_{0,n}$.

III.1. Maximum-Likelihood Estimation

Despite the aforementioned shortcomings of ML estimation for the present purposes, ML is of interest because it is the procedure typically used in agricultural economics. Hence, to mimic the strategy followed by such literature, estimation is also performed using ML under the assumption that the errors defined by (10) and (11) are jointly normally distributed.

IV. Results and Discussion

GMM estimation results for the flexible utility specification under the baseline initial wealth distribution (6) are summarized in tables 1 through 3. The tables contain the median and the 2.5% and 97.5% quantiles (within parentheses) 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. To save space, outcomes for the medium-variance scenario are

Table 1. GMM parameter estimates for flexible utility estimation specification, assuming baseline initial wealth distribution, 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	−19.2	4.8	2.909	0.211	0.608
		(−20.3, −18.9)	(2.3, 19.4)	(2.40, 3.61)	(0.18, 0.23)	(0.51, 0.70)
DRRA[b]	500	−19.1	3.9	2.842	0.215	0.628
		(−19.3, −18.9)	(2.8, 5.9)	(2.50, 3.10)	(0.20, 0.22)	(0.58, 0.68)
DRRA[b]	1,000	−19.0	3.7	2.847	0.216	0.625
		(−19.2, −18.9)	(2.9, 5.4)	(2.62, 3.06)	(0.20, 0.22)	(0.58, 0.66)
CRRA[b]	100	4.0	194	3.112	0.232	0.572
		(−20.4, 38227)	(20.8, 81743)	(2.96, 3.31)	(0.21, 0.28)	(0.52, 0.60)
CRRA[b]	500	11.5	92.0	3.064	0.222	0.583
		(−14.8, 22006)	(26.6, 24280)	(2.99, 3.15)	(0.21, 0.24)	(0.56, 0.60)
CRRA[b]	1,000	10.4	73.3	3.055	0.219	0.585
		(−9.1, 18113)	(29.5, 16305)	(3.00, 3.11)	(0.21, 0.23)	(0.57, 0.60)
IRRA[b]	100	30.9	194	3.097	0.228	0.576
		(−20.0, 33558)	(23.7, 72700)	(2.94, 3.30)	(0.20, 0.28)	(0.53, 0.61)
IRRA[b]	500	40.2	66.6	3.002	0.210	0.598
		(−14.4, 80342)	(15.8, 53750)	(2.94, 3.08)	(0.20, 0.22)	(0.58, 0.61)
IRRA[b]	1,000	40.6	54.2	2.993	0.209	0.599
		(−2.7, 87001)	(18.5, 50828)	(2.95, 3.04)	(0.20, 0.22)	(0.59, 0.61)

[a]For each parameter, the table reports the median and the 2.5% and 97.5% quantiles (within parentheses) from 20000, 4000, and 2000 estimates for sample sizes of 100, 500, and 1000, respectively. Two million simulated observations about decision variables were used to construct 20000 (4000, 2000) samples with 100 (500, 1000) observations per sample, which were then employed to obtain the reported parameter estimates.

[b]DRRA, CRRA, and IRRA risk structures correspond to [γ_0, γ_1] equal to [−18.4, 1], [0, 3] and [43, 6], respectively.

Table 2. GMM parameter estimates for flexible utility estimation specification, assuming baseline initial wealth distribution, 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	−16.4 (−19.7, 17864)	4.4 (1.0, 3718)	2.896 (2.50, 3.31)	0.213 (0.19, 0.25)	0.604 (0.51, 0.68)
DRRA[b]	500	−15.6 (−19.0, 8.8)	4.2 (2.0, 11.7)	2.875 (2.70, 3.05)	0.212 (0.20, 0.23)	0.607 (0.56, 0.64)
DRRA[b]	1,000	−15.3 (−18.7, −3.6)	4.2 (2.3, 8.0)	2.871 (2.75, 2.99)	0.212 (0.20, 0.22)	0.606 (0.58, 0.63)
CRRA[b]	100	18.4 (−13.9, 30617)	15.5 (3.5, 7464)	2.908 (2.51, 3.30)	0.217 (0.19, 0.25)	0.612 (0.53, 0.69)
CRRA[b]	500	17.8 (−0.5, 80.3)	13.4 (7.4, 27.5)	2.857 (2.68, 3.04)	0.217 (0.20, 0.23)	0.614 (0.58, 0.65)
CRRA[b]	1,000	17.9 (3.8, 51.6)	13.1 (8.8, 21.0)	2.855 (2.73, 2.98)	0.217 (0.21, 0.23)	0.612 (0.58, 0.64)
IRRA[b]	100	42.7 (−1.3, 52997)	27.7 (6.5, 11036)	2.911 (2.51, 3.30)	0.217 (0.19, 0.25)	0.613 (0.53, 0.69)
IRRA[b]	500	42.5 (41.4, 18365)	21.1 (10.6, 3127)	2.860 (2.64, 3.03)	0.216 (0.20, 0.23)	0.613 (0.58, 0.65)
IRRA[b]	1,000	42.6 (41.8, 15549)	14.8 (10.8, 2325)	2.787 (2.63, 2.96)	0.214 (0.21, 0.22)	0.620 (0.59, 0.65)

[a]For each parameter, the table reports the median and the 2.5% and 97.5% quantiles (within parentheses) from 20000, 4000, and 2000 estimates for sample sizes of 100, 500, and 1000, respectively. Two million simulated observations about decision variables were used to construct 20000 (4000, 2000) samples with 100 (500, 1000) observations per sample, which were then employed to obtain the reported parameter estimates.

[b]DRRA, CRRA, and IRRA risk structures correspond to [γ_0, γ_1] equal to [−18.4, 1], [0, 3] and [43, 6], respectively.

Table 3. GMM parameter estimates for flexible utility estimation specification, assuming baseline initial wealth distribution, 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	−19.0 (−19.8, 15936)	4.6 (1.6, 6120)	2.946 (2.53, 3.35)	0.211 (0.19, 0.24)	0.598 (0.52, 0.67)
DRRA[b]	500	−17.3 (−19.1, 13.6)	4.8 (2.5, 19.4)	2.950 (2.62, 3.08)	0.208 (0.20, 0.22)	0.602 (0.57, 0.65)
DRRA[b]	1,000	−16.3 (−19.0, −2.7)	5.5 (2.8, 13.0)	2.954 (2.73, 3.04)	0.208 (0.20, 0.22)	0.602 (0.58, 0.63)
CRRA[b]	100	16.1 (−19.2, 37748)	21.4 (5.1, 13654)	2.982 (2.70, 3.28)	0.214 (0.20, 0.24)	0.598 (0.55, 0.64)
CRRA[b]	500	17.3 (−1.1, 131)	17.9 (10.2, 51.3)	2.946 (2.82, 3.08)	0.212 (0.20, 0.22)	0.604 (0.58, 0.62)
CRRA[b]	1,000	17.0 (2.7, 66.4)	17.1 (11.6, 31.2)	2.946 (2.86, 3.04)	0.212 (0.21, 0.22)	0.603 (0.59, 0.62)
IRRA[b]	100	42.0 (−9.8, 61318)	36.9 (8.0, 17638)	2.982 (2.71, 3.27)	0.214 (0.20, 0.24)	0.598 (0.55, 0.64)
IRRA[b]	500	42.4 (40.3, 31389)	25.1 (11.8, 5845)	2.941 (2.82, 3.04)	0.210 (0.20, 0.22)	0.604 (0.58, 0.62)
IRRA[b]	1,000	42.4 (41.2, 27094)	16.2 (12.0, 5047)	2.916 (2.81, 3.00)	0.206 (0.20, 0.22)	0.608 (0.59, 0.62)

[a]For each parameter, the table reports the median and the 2.5% and 97.5% quantiles (within parentheses) from 20000, 4000, and 2000 estimates for sample sizes of 100, 500, and 1000, respectively. Two million simulated observations about decision variables were used to construct 20000 (4000, 2000) samples with 100 (500, 1000) observations per sample, which were then employed to obtain the reported parameter estimates.

[b]DRRA, CRRA, and IRRA risk structures correspond to [γ_0, γ_1] equal to [−18.4, 1], [0, 3] and [43, 6], respectively.

not reported, as they typically lie between the small- and large-variance results.

The clearest pattern in tables 1 through 3 is that, as expected, the precision of the parameter estimates, as measured by the width of the 95% CIs, increases with the sample size. Regarding technology parameters, α_A and α_0 are the ones estimated with the greatest and the least precision, respectively. That is, the precision of the technology estimates varies inversely with the true parameter values. Median estimates are typically closest to the true values in the mixed-variance scenario. However, there seems to be little association between the precision of $[\hat{\alpha}_0, \hat{\alpha}_A, \hat{\alpha}_B]$ and the variance of the output and price shocks or the risk preferences. For example, consider the 95% CIs for $[\alpha_0, \alpha_A, \alpha_B]$ corresponding to 100-observation samples. With DRRA preferences, the 95% CIs for α_0 and α_B (α_A) are narrower (wider) under the large-variance scenario than under the small-variance scenario. In contrast, with IRRA preferences, the 95% CIs for α_0 and α_B (α_A) are narrower (wider) when the variance is small than when the variance is large.

Overall, the figures reported in tables 1 through 3 show that technology parameters can be accurately recovered, even with samples comprising only 100 observations. Interestingly, however, all of the median estimates of α_A reported in tables 1 through 3 exceed the true parameter value, and it is often the case that the 95% CIs are very close to, but fail to contain, the true value of α_A. This happens in all of the scenarios reported in tables 1 through 3 for 500- and 1000-observation samples, and also for 100-observation samples and CRRA and IRRA preferences in the small-variance scenario. The true parameter value also lies outside the corresponding 95% CI in the case of α_0 when preferences are CRRA and IRRA and samples are of size 1000 under the large-variance scenario.

Unlike technology estimates, the estimates of the preference parameters bear little resemblance to the true values and are very imprecise (except for $\hat{\gamma}_0$ under DRRA preferences in the small-variance scenario). 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, 4.0 (100-observation samples in the small-variance scenario) and 13.1 (1000-observation samples in the large-variance

scenario), compared to true values of $\gamma_0 = 0$ and $\gamma_1 = 3$. Prominently, none of the 95% CIs for γ_1 includes the true value. Further, the distance between the true γ_1 value and the closest (i.e., lower) bound of the 95% CIs is large and increasing with the sample size. For example, in the mixed-variance scenario, the 2.5% quantiles for $\hat{\gamma}_1$ corresponding to 100-, 500-, and 1000-observation samples are 5.1, 10.2, and 11.6, respectively. For γ_0, the true value lies outside the 95% CIs for samples of 1000 observations in the large- and mixed-variance scenarios. When the variance is small or samples have 100 observations, the 97.5% quantiles for $\hat{\gamma}_0$ and $\hat{\gamma}_1$ are too large to be credible. It is also worth noting that the distributions of both $\hat{\gamma}_0$ and $\hat{\gamma}_1$ are substantially skewed to the right.

As for CRRA preferences, the true values of γ_1 for DRRA and IRRA are grossly overestimated and lie outside the respective 95% CIs. Also similar to the CRRA case, the gap between the true values of γ_1 and the bound of the 95% CI closest to them increases with the number of observations in the sample, except for IRRA preferences in the small-variance scenario. In addition, the upper bound for the 95% CIs corresponding to γ_1 under IRRA is unreasonably large, and the same is true under DRRA for 100-observation samples in the large- and mixed-variance scenarios. In all instances, the distribution of $\hat{\gamma}_1$ depicts a noticeable skew to the right.

Under DRRA and IRRA preferences, distributions of $\hat{\gamma}_0$ are also clearly right-skewed. Median estimates of γ_0 are relatively close to the respective true values. However, there are obvious problems with $\hat{\gamma}_0$, as well. For IRRA preferences, the lower bound of the 95% CI is negative in the small-variance scenario and in the 100-observation-sample cases of the large- and mixed-variance scenarios. The problem with this finding is that a negative $\hat{\gamma}_0$ means that the econometrician would erroneously conclude that preferences are DRRA instead of IRRA (see discussion of expression (4) above). Upper bounds for the 95% CIs are also problematic under IRRA because they are excessively large.

For DRRA preferences, the 95% CIs for γ_0 are narrower than for CRRA or IRRA. However, the 97.5% quantiles of $\hat{\gamma}_0$ under DRRA are positive for samples of 100 and 500

observations in the large- and mixed-variance scenarios. A positive $\hat{\gamma}_0$ would incorrectly infer that the corresponding preferences are IRRA, rather than DRRA. Further, the 95% CIs for DRRA preferences lie entirely below the true value of $\gamma_0 = -18.4$ in the small-variance scenario.

The CI for γ_0 when the variance is small and preferences are DRRA is unusually narrow relative to any of the other 95% CIs for utility parameters. The 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 compute optimal input choices. Such constraint binds in 22.7% (14.3%, 8.5%) of the small- (medium-, large-) variance simulated observations for DRRA, versus only 4% (0.2%, 0%) and 0.2% (0%, 0%) of the observations for CRRA and IRRA, respectively. As samples contain larger percentages of observations where the cost constraint is binding, the estimation constraint $\hat{\gamma}_0 > 0.001 - min(W_{1n}, n = 1, ..., N)$ is more likely to bind, as well. Importantly, the fact that the cost constraint is relevant only for DRRA preferences in the small-variance scenario means that it cannot be blamed for the poor quality of the preference estimates shown in tables 1 through 3.

The mixed-variance scenario reported in table 3 contains additional information potentially useful for identifying 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 exactly knows the level of uncertainty embedded in the decision maker's optimal choices, which is highly unrealistic. Therefore, it seems safe to conclude that the structure of risk aversion is very unlikely to be rendered identifiable by the presence of heteroskedasticity.

Table 4 reports results for the medium-variance scenario, assuming a wide distribution of initial wealth. Estimates of the technology parameters appear to be largely unaffected. The wide initial distribution also seems to have a neutral impact on the quality of the IRRA utility estimates. In contrast, utility estimates for DRRA preferences are clearly worse compared to the

Table 4. GMM parameter estimates for flexible utility estimation specification, assuming initial wealth uniformly distributed with lower bound 18.9 and upper bound 500, medium-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	7.6	34.9	2.967	0.212	0.600
		(–26.4, 62451)	(2.3, 12471)	(2.67, 3.26)	(0.19, 0.24)	(0.55, 0.64)
DRRA[b]	500	19.0	22.7	2.932	0.209	0.603
		(–19.9, 34755)	(4.0, 5228)	(2.77, 3.23)	(0.20, 0.22)	(0.56, 0.62)
DRRA[b]	1,000	19.4	20.2	2.926	0.208	0.603
		(–19.5, 28029)	(4.7, 3415)	(2.81, 3.23)	(0.20, 0.22)	(0.56, 0.62)
CRRA[b]	100	58.4	62.3	2.971	0.215	0.601
		(–25.0, 58270)	(8.4, 10984)	(2.68, 3.24)	(0.20, 0.24)	(0.56, 0.64)
CRRA[b]	500	71.8	44.5	2.921	0.212	0.606
		(–1.0, 1217)	(18.2, 218)	(2.78, 3.05)	(0.20, 0.22)	(0.58, 0.62)
CRRA[b]	1,000	70.3	40.8	2.914	0.212	0.606
		(13.7, 599)	(19.8, 109)	(2.81, 3.01)	(0.20, 0.22)	(0.59, 0.62)
IRRA[b]	100	42.2	61.4	2.971	0.213	0.602
		(–19.9, 55608)	(9.5, 9261)	(2.69, 3.24)	(0.20, 0.24)	(0.56, 0.64)
IRRA[b]	500	42.7	35.0	2.906	0.209	0.608
		(41.2, 56880)	(12.5, 6494)	(2.76, 3.01)	(0.20, 0.22)	(0.59, 0.63)
IRRA[b]	1,000	42.8	29.8	2.890	0.209	0.608
		(42.2, 31199)	(13.9, 3536)	(2.80, 2.97)	(0.20, 0.22)	(0.60, 0.62)

[a]For each parameter, the table reports the median and the 2.5% and 97.5% quantiles (within parentheses) from 20000, 4000, and 2000 estimates for sample sizes of 100, 500, and 1000, respectively. Two million simulated observations about decision variables were used to construct 20000 (4000, 2000) samples with 100 (500, 1000) observations per sample, which were then employed to obtain the reported parameter estimates.

[b]DRRA, CRRA, and IRRA risk structures correspond to [γ_0, γ_1] equal to [−18.4, 1], [0, 3] and [43, 6], respectively.

baseline distribution of initial wealth, as the median $\hat{\gamma}_0$ is positive (implying IRRA preferences), the median $\hat{\gamma}_1$ is far greater than the true value $\gamma_1 = 1$, the 95% CI for γ_1 does not contain its true value, and the 97.5% quantiles for $\hat{\gamma}_0$ and $\hat{\gamma}_1$ are unrealistically large. The wide initial wealth distribution also has a negative effect on the quality of $\hat{\gamma}_0$ under CRRA preferences, because the median lies much farther apart from the true value than under the baseline initial wealth distribution. In summary, if anything, a wider initial wealth distribution seems to negatively affect the estimation of the flexible utility parameters.

ML results for the medium-variance scenario are shown in table 5. ML estimates of technology parameters seem slightly better than their GMM counterparts, as the ML medians tend to be closer to the true values and the ML 95% CIs are often narrower. In regard to utility parameters, for $\hat{\gamma}_1$ the ML medians are closer to the respective true values than the GMM medians, but the opposite is true for $\hat{\gamma}_0$ medians under CRRA and IRRA preferences. ML estimates of both γ_0 and γ_1 are clearly better than the GMM estimates in terms of 97.5% quantiles, because the former are usually much smaller. Overall, however, ML estimates of the flexible utility parameters are quite poor. This assessment is reinforced by the facts that the estimates reported in table 5 correspond only to samples where the numerical optimization involved in ML estimation converged, and that convergence was achieved for only 60% to 80% of the samples, depending on the scenario under analysis.

IV.1. Estimation under Restricted Utility Specifications

As noted earlier, Černý (2004) 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 an argument applies to agricultural production under uncertainty, observed input choices would be consistent with an infinite number of $\boldsymbol{\gamma}$ 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 incur difficulties due to lack of identification.

Table 5. Maximum-likelihood parameter estimates for flexible utility estimation specification, assuming baseline initial wealth distribution, medium-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	−17.8	3.7	2.871	0.206	0.615
		(−19.3, 29.7)	(1.1, 19.4)	(2.55, 3.13)	(0.19, 0.23)	(0.58, 0.67)
DRRA[b]	500	−18.4	2.2	2.934	0.201	0.601
		(−18.9, −11.3)	(1.3, 5.0)	(2.81, 3.05)	(0.19, 0.21)	(0.58, 0.62)
DRRA[b]	1,000	−18.5	2.0	2.949	0.200	0.598
		(−18.9, −15.8)	(1.2, 3.3)	(2.86, 3.04)	(0.20, 0.20)	(0.58, 0.61)
CRRA[b]	100	19.9	8.9	2.896	0.205	0.610
		(−18.3, 148)	(1.5, 36.4)	(2.54, 3.20)	(0.19, 0.23)	(0.57, 0.67)
CRRA[b]	500	35.9	8.5	2.951	0.201	0.600
		(−8.6, 130)	(2.7, 20.6)	(2.80, 3.10)	(0.19, 0.21)	(0.58, 0.62)
CRRA[b]	1,000	40.4	8.6	2.958	0.200	0.599
		(−0.1, 124)	(3.5, 18.9)	(2.85, 3.07)	(0.20, 0.21)	(0.58, 0.62)
IRRA[b]	100	32.3	10.1	2.901	0.204	0.609
		(−17.4, 174)	(1.8, 40.0)	(2.54, 3.20)	(0.19, 0.22)	(0.56, 0.67)
IRRA[b]	500	53.5	10.5	2.957	0.201	0.599
		(−1.0, 168)	(3.7, 22.8)	(2.81, 3.10)	(0.19, 0.21)	(0.58, 0.62)
IRRA[b]	1,000	67.0	11.4	2.960	0.200	0.598
		(12.4, 159)	(5.2, 21.1)	(2.85, 3.08)	(0.20, 0.21)	(0.58, 0.61)

[a]For each parameter, the table reports the median and the 2.5% and 97.5% quantiles (within parentheses) of the "converged" estimates obtained from 20000, 4000, and 2000 samples of 100, 500, and 1000 observations per sample, respectively. For DRRA (CRRA, IRRA) sample sizes [100, 500, 1000], convergence was achieved for [66.1%, 72.1%, 78.6%] ([66.0%, 76.0%, 82.5%], [61.2%, 63.1%, 65.5%]) of the samples. Two million simulated observations about decision variables were used to construct 20000 (4000, 2000) samples with 100 (500, 1000) observations per sample, which were then employed to obtain the reported parameter estimates.

[b]DRRA, CRRA, and IRRA risk structures correspond to [γ_0, γ_1] equal to [−18.4, 1], [0, 3] and [43, 6], respectively.

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 estimating 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 6, 7, and 8, which correspond respectively to tables 1, 2, and 3 under unrestricted estimation. In the interest of space, the restricted-utility counterparts of tables 4 and 5 are omitted, as they provide few additional insights.

Compared to tables 1 through 3, tables 6 through 8 do not exhibit substantial differences regarding the estimates of technology parameters. Overall, however, the quality of the technology estimates tends to be better for the restricted specification, in terms of both the width of the 95% CIs and the difference between the medians and the respective true values.

Inspection of the estimates of the coefficient of relative risk aversion shown in tables 6 through 8 uncovers a number of interesting outcomes. First, estimates of γ_1 for the small-variance scenario are noticeably worse than the estimates corresponding to the large- and mixed-variance scenarios. In the small-variance scenario, $\hat{\gamma}_1$ medians are substantially further apart from the true value of the coefficient of local relative risk aversion (i.e., 3), and the 95% CIs are much wider and more skewed to the right. Greater variance increases the overall noise in the 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.) This suggests that a considerable amount of variability in the price and output shocks is required to adequately identify the level of relative

Table 6. GMM parameter estimates for restricted utility estimation specification ($\gamma_0 = 0$), assuming baseline initial wealth distribution, 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	120	3.116	0.232	0.572
		(26.3, 2443)	(2.95, 3.31)	(0.20, 0.28)	(0.53, 0.61)
DRRA[b]	500	44.1	3.010	0.210	0.595
		(14.3, 134)	(2.95, 3.12)	(0.20, 0.24)	(0.57, 0.61)
DRRA[b]	1,000	36.7	3.004	0.209	0.597
		(15.1, 89.6)	(2.96, 3.08)	(0.20, 0.22)	(0.58, 0.61)
CRRA[b]	100	94.7	3.101	0.228	0.576
		(15.4, 1157)	(2.94, 3.31)	(0.20, 0.28)	(0.53, 0.61)
CRRA[b]	500	30.1	3.011	0.210	0.596
		(10.2, 101)	(2.95, 3.11)	(0.20, 0.23)	(0.57, 0.61)
CRRA[b]	1,000	23.6	2.997	0.208	0.599
		(10.6, 52.2)	(2.96, 3.04)	(0.20, 0.22)	(0.59, 0.61)
IRRA[b]	100	84.0	3.088	0.225	0.579
		(12.8, 953)	(2.94, 3.30)	(0.20, 0.28)	(0.53, 0.61)
IRRA[b]	500	27.2	3.006	0.210	0.598
		(10.6, 88.1)	(2.95, 3.10)	(0.20, 0.23)	(0.57, 0.61)
IRRA[b]	1,000	21.3	2.994	0.208	0.600
		(10.2, 47.6)	(2.96, 3.04)	(0.20, 0.22)	(0.59, 0.61)

[a]For each parameter, the table reports the median and the 2.5% and 97.5% quantiles (within parentheses) from 20000, 4000, and 2000 estimates for sample sizes of 100, 500, and 1000, respectively. Two million simulated observations about decision variables were used to construct 20000 (4000, 2000) samples with 100 (500, 1000) observations per sample, which were then employed to obtain the reported parameter estimates.

[b]DRRA, CRRA, and IRRA risk structures correspond to [γ_0, γ_1] equal to [−18.4, 1], [0, 3] and [43, 6], respectively.

Table 7. GMM parameter estimates for restricted utility estimation specification ($\gamma_0 = 0$), assuming baseline initial wealth distribution, 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	8.4	2.893	0.213	0.615
		(2.7, 19.6)	(2.51, 3.27)	(0.19, 0.25)	(0.54, 0.68)
DRRA[b]	500	7.7	2.837	0.212	0.617
		(4.4, 11.6)	(2.67, 3.00)	(0.20, 0.23)	(0.58, 0.65)
DRRA[b]	1,000	7.8	2.834	0.212	0.616
		(4.7, 10.2)	(2.72, 2.96)	(0.20, 0.22)	(0.59, 0.64)
CRRA[b]	100	8.8	2.922	0.213	0.609
		(3.5, 20.2)	(2.53, 3.30)	(0.19, 0.25)	(0.53, 0.68)
CRRA[b]	500	8.3	2.884	0.213	0.608
		(5.3, 12.1)	(2.71, 3.05)	(0.20, 0.23)	(0.57, 0.64)
CRRA[b]	1,000	8.3	2.879	0.214	0.606
		(5.4, 10.8)	(2.74, 3.00)	(0.20, 0.22)	(0.58, 0.63)
IRRA[b]	100	8.1	2.920	0.212	0.609
		(3.0, 19.4)	(2.53, 3.30)	(0.19, 0.24)	(0.53, 0.68)
IRRA[b]	500	7.2	2.876	0.211	0.608
		(5.0, 11.2)	(2.69, 3.05)	(0.20, 0.22)	(0.57, 0.64)
IRRA[b]	1,000	7.1	2.865	0.211	0.608
		(5.8, 8.7)	(2.74, 3.00)	(0.20, 0.22)	(0.58, 0.64)

[a]For each parameter, the table reports the median and the 2.5% and 97.5% quantiles (within parentheses) from 20000, 4000, and 2000 estimates for sample sizes of 100, 500, and 1000, respectively. Two million simulated observations about decision variables were used to construct 20000 (4000, 2000) samples with 100 (500, 1000) observations per sample, which were then employed to obtain the reported parameter estimates.

[b]DRRA, CRRA, and IRRA risk structures correspond to [γ_0, γ_1] equal to [−18.4, 1], [0, 3] and [43, 6], respectively.

Table 8. GMM parameter estimates for restricted utility estimation specification ($\gamma_0 = 0$), assuming baseline initial wealth distribution, 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	13.0	2.968	0.211	0.602
		(4.4, 37.3)	(2.70, 3.24)	(0.19, 0.24)	(0.55, 0.65)
DRRA[b]	500	7.7	2.905	0.203	0.612
		(3.9, 14.8)	(2.80, 3.02)	(0.20, 0.22)	(0.59, 0.63)
DRRA[b]	1,000	6.5	2.887	0.202	0.615
		(4.1, 12.7)	(2.82, 2.98)	(0.20, 0.21)	(0.60, 0.63)
CRRA[b]	100	12.3	2.982	0.213	0.598
		(4.7, 35.4)	(2.71, 3.26)	(0.20, 0.24)	(0.55, 0.64)
CRRA[b]	500	10.5	2.948	0.210	0.604
		(5.8, 16.7)	(2.84, 3.07)	(0.20, 0.22)	(0.58, 0.62)
CRRA[b]	1,000	8.5	2.936	0.207	0.605
		(5.6, 14.1)	(2.86, 3.02)	(0.20, 0.22)	(0.59, 0.62)
IRRA[b]	100	11.5	2.980	0.212	0.599
		(4.3, 34.2)	(2.71, 3.26)	(0.20, 0.24)	(0.55, 0.64)
IRRA[b]	500	9.1	2.954	0.208	0.603
		(6.5, 16.1)	(2.83, 3.06)	(0.20, 0.22)	(0.58, 0.62)
IRRA[b]	1,000	8.3	2.952	0.207	0.603
		(7.1, 13.4)	(2.85, 3.02)	(0.20, 0.21)	(0.59, 0.62)

[a]For each parameter, the table reports the median and the 2.5% and 97.5% quantiles (within parentheses) from 20000, 4000, and 2000 estimates for sample sizes of 100, 500, and 1000, respectively. Two million simulated observations about decision variables were used to construct 20000 (4000, 2000) samples with 100 (500, 1000) observations per sample, which were then employed to obtain the reported parameter estimates.

[b]DRRA, CRRA, and IRRA risk structures correspond to [γ_0, γ_1] equal to [−18.4, 1], [0, 3] and [43, 6], respectively.

risk aversion, let alone the structure of risk aversion.

Second, the quality of the γ_1 estimates exhibits a clear improvement as the sample size increases. For 100-observation samples, the 95% CIs are much wider and the difference between the $\hat{\gamma}_1$ medians and $\hat{\gamma}_1 \cong 3$ is greater than for larger samples. These findings cast doubt on the ability to accurately estimate risk preferences (even if restricted) from small samples.

A third interesting feature of the $\hat{\gamma}_1$ distributions is that the medians systematically overestimate the true value of the coefficient of local relative risk aversion (i.e., 3). Further, the latter is smaller than the lower bound of the 95% CIs in all but two of the cases reported in tables 6 through 8 (the exceptions being 100-observation samples for DRRA and IRRA in the large-variance scenario). The magnitude of the positive bias in the medians of $\hat{\gamma}_1$ declines with the number of observations in the sample, but is still substantial, even when samples have as many as 1000 observations.

Finally, the distribution of $\hat{\gamma}_1$ for the large-variance scenario is very similar across DRRA, CRRA, and IRRA preferences. Since all utility functions are parameterized to have $R(W_0) = 3$, this finding is consistent with Černý's (2004) claim that optimal decisions for individuals with different preferences but the same value of $R(W_0)$ are similar, except for very large and/or very skewed risks. For the small- and mixed-variance scenario, the distributions of $\hat{\gamma}_1$ for CRRA and IRRA preferences look similar to each other and relatively different from the one for DRRA preferences. The most plausible explanation for this finding is the cost constraint discussed earlier, because such constraint is more often binding when preferences are DRRA and the variance is small.

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 8 with $\hat{\gamma}_1$s in tables 6 and 7). This finding lends further support to the conclusion that heteroskedasticty is highly unlikely to provide the additional 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 often fail to include the true values. Further, there are scenarios for which the 95% CIs imply increasing relative risk aversion (IRRA) when the true preferences are characterized by decreasing relative risk aversion (DRRA), and *vice versa*.

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 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 the estimates of the single CRRA parameter are positively biased in all scenarios analyzed, and are much poorer when the underlying risk is small or when samples have 100 observations (even if the true utility is CRRA). The latter finding is relevant, as studies that have simultaneously estimated technology and the structure of risk aversion have often relied on real-world samples comprising fewer than 100 observations. The restricted utility specification lends support to Černý's (2004) claim that optimal decisions for individuals with different preferences, but the same value of local risk aversion measured at the initial wealth level, are similar, except for very large and/or very skewed risks.

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 purposefully 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 slightly 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. In this regard, it would be useful to investigate what additional data requirements and assumptions are needed to obtain reliable estimates of flexible

functional forms simultaneously with technology estimates. However, our findings strongly suggest that such gains are likely to be unwarranted in many 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 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 wasting scarce resources by gathering and analyzing data that are highly unlikely to yield the information researchers seek.

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

The initial wealth pdf (6) is estimated using the balanced panel employed by Hart and Lence (2004), which contains annual initial wealth observations for 317 Iowa farms over the period 1991 through 1998. Since monetary data in the simulation model are scaled 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 scaled initial wealth values. In the scaling ratio, the numerator 7 equals the approximate median costs (i.e., *median*(***r 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 *median*(***r 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 (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*(·) 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 scaled 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 $R^{(j)} = mean[R(W_0)|\gamma_0^{(j)} = 0.5 - \underline{W}_0^{(j)}, \gamma_1 = 1]$.

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 [γ_0, γ_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 (2004) 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 [$e_{y,0.5}$, $e_{y,1.5}$, ..., $e_{y,98.5}$, $e_{y,99.5}$], where $e_{y,q}$ is the qth quantile of standardized yields. The probability distribution of $\tilde{e}_y$ for the medium-variance scenario consists of [$e_{y,0.5}$, $e_{y,1.5}$, ..., $e_{y,98.5}$, $e_{y,99.5}$], with probabilities [$\pi^M_{y,0.5}, \pi^M_{y,1.5}, ..., \pi^M_{y,98.5}, \pi^M_{y,99.5}$] = [0.01, 0.01, ..., 0.01, 0.01].

For the low- and high-variance scenarios, Prelec's (1998) probability weighting function is used to assign the probabilities $\pi^L_{y,q}$ and $\pi^H_{y,q}$ corresponding to $e_{y,q}$. For the low-variance case, the distribution of $\tilde{e}_y$ is given by $e_{y,q}$ with probability $\pi^L_{y,q} \equiv \pi(q + 0.5; \phi_1^L, \phi_2^L) - \pi(q - 0.5; \phi_1^L, \phi_2^L)$ 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, ..., 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 $\pi^H_{y,q} \equiv \pi(q + 0.5; \phi_1^H, \phi_2^H) - \pi(q - 0.5; \phi_1^H, \phi_2^H)$ for [ϕ_1^H, ϕ_2^H] = [0.93, 1.73]. Compared to the medium-variance probabilities $\pi^M_{y,q}$, the low-variance (high-variance) probabilities shift weight from the extremes (middle) of vector [$e_{y,0.5}$, $e_{y,1.5}$, ..., $e_{y,98.5}$, $e_{y,99.5}$] to its middle (extremes), so as to reduce (increase) the standard deviation by 50%, while maintaining the mean unchanged at 1.[15]

[15]Note that $\phi^M_{y,q} = \pi(q + 0.5; \phi_1^M, \phi_2^M) - \pi(q - 0.5; \phi_1^M, \phi_2^M)$ for [ϕ_1^M, ϕ_2^M] = [1, 1].

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

Expression (7) is an approximation based on the regression estimates reported in table A.1.

Table A1. Price regressions for corn.

Data	Regression	Regression Std. Error	R^2
U.S. aggregate data, 1970-2005	$ln(p_t) = 47 - 0.0248\ t + 0.48\ ln(p_{t-1}) - 0.85\ ln(e_{y,t}) + e_{p,t}$ (11) (0.0061) (0.13) (0.26)	0.156	0.914
Farm-level yield data, 1991-1998	$ln(p_t) = 36.4 - 0.0197\ t + 0.303\ ln(p_{t-1}) - 0.248\ ln(e_{y,f,t}) + e_{p,f,t}$ (2.1) (0.0011) (0.024) (0.011)	0.135	0.266

Note: Variable p_t is the U.S. season-average corn 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 $lne_{y,t}$ is the ordinary least-squares residual of the regression of $ln(yield_t)$ 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). 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 2004 dataset). Each regression comprises 35 observation for the U.S. aggregate data, and 3,200 observations for the farm-level yield data.

References

Abdulkadri, A. O., M. R. Langemeier, and A. M. Featherstone. 2003. "Estimating Risk Aversion Coefficients for Dryland Wheat, Irrigated Corn and Dairy Producers in Kansas. *Applied Economics* 35:825-834.

Altuğ , S., and P. Labadie. 1994. *Dynamic Choice and Asset Markets*. San Diego, CA: Academic Press.

Antle, J. M. 1987. "Econometric Estimation of Producers' Risk Attitudes." *American Journal of Agricultural Economics* 69:509-522.

Antle, J. M. 1989. "Nonstructural Risk Attitude Estimation." *American Journal of Agricultural Economics* 71:774-784.

Bar-Shira, Z., R. E. Just, and D. Zilberman. 1997. "Estimation of Farmers' Rrisk Attitude: An Eeconometric Approach. *Agricultural Economics* 17:211-222.

Binswanger, H. P. 1980. "Attitudes toward Risk: Experimental Measurement In Rural India." *American Journal of Agricultural Economics* 62:395-407.

Brink, L., and B. McCarl. 1978. "The Tradeoff Between Expected Return and Risk Among Cornbelt Farmers." *American Journal of Agricultural Economics* 60:258-263.

Černý, A. 2004. *Mathematical Techniques in Finance: Tools for Incomplete Markets*. Princeton, NJ: Princeton University Press.

Chavas, J. P., and M. T. Holt. 1996. "Economic Behavior Under Uncertainty: A Joint Analysis of Risk Preferences and Technology." *Review of Economics and Statistics* 78:329-335.

Cliff, M. T. 2003. "GMM and MINZ Program Libraries for MATLAB." Unpublished manuscript, Krannert Graduate School of Management, Purdue University, Indiana. Available at *http://mcliff.cob.vt.edu/progs/gmmdoc.pdf* (accessed May 20, 2007).

Coyle, B. T. 1999. "Risk Aversion and Yield Uncertainty in Duality Models of Production: A Mean-Variance Approach." *American Journal of Agricultural Economics* 81:553-567.

Davidson, R., and J. G. Mackinnon. 1993. *Estimation and Inference in Econometrics*. New York: Oxford University Press.

Economic Research Service, U.S. Department of Agriculture. 2008. "Briefing Rooms: Farm Household Economics and Well Being: Assets, Debt, and Wealth." Available at *http://www.ers.usda.gov/briefing/WellBeing/Gallery/Table1FamilyNetWorth.htm* (accessed June 17, 2008).

Geweke, J. 1988. "Antithetic Acceleration of Monte Carlo Integration in Bayesian Inference." *Journal of Econometrics* 38: 73-90.

Gollier, C. 2001. *The Economics of Risk and Time.* Cambridge, MA: The MIT Press.

Hansen, L. P. 1982. "Large Sample Properties of Generalized Method of Moment Estimators." *Econometrica* 50:1029-1054.

Hansen, L. P., and K. J. Singleton. 1982. "Generalized Instrumental Variables Estimation of Nonlinear Rational Exopectations Models." *Econometrica* 50:1269-1286.

Hart, C. E., and S. H. Lence. 2004. "Financial Constraints and Farm Investment: A Bayesian Examination." *Journal of Business and Economic Statistics* 22:51-63.

Harwood, J., R. Heifner, K. Coble, J. Perry, and A. Somwaru. 1999. "Managing Risk in Farming: Concepts, Research, and Analysis." Economic Research Service, Department of Agriculture, Agricultural Economic Report N0. 774.

Isik, M., and M. Khanna. 2003. "Stochastic Technology, Risk Preferences, and Adoption of Site-Specific Technologies." *American Journal of Agricultural Economics* 85:305-317.

Jagannathan, R., G. Skoulakis, and Z. Wang. 2002. "Generalized Method of Moments: Applications in Finance." *Journal of Business and Economic Statistics* 20:470-481.

Just, R. E., and R. D. Pope, Editors. 2002. *A Comprehensive Assessment of the Role of Risk in U.S. Agriculture*. Boston, MA: Kluwer Academic Publishers.

Just, R. E, and R. D. Pope. 1978. "Stochastic Specification of Production Functions and Economic Implications." *Journal of Econometrics* 7:67-89.

Kallberg, J. G., and W. T. Ziemba. 1983. "Comparison of Alternative Utility Functions in Portfolio Selection Problems." *Management Science* 29:1257-1276.

Kocherlakota, N. 1996. "The Equity Premium: It's Still A Puzzle." *Journal of Economic*

Literature 34:42-71.

Koundouri, P., and C. Nauges. 2005. "On Production Function Estimation with Selectivity and Risk Considerations." *Journal of Agricultural and Resource Economics* 30:597-608.

Kumbhakar, S. C. 2002a. "Risk Preference and Productivity Measurement under Output Price Uncertainty." *Empirical Economics* 27:461-472.

Kumbhakar, S. C. 2002b. "Specification and Estimation of Production Risk, Risk Preferences and Technical Efficiency." *American Journal of Agricultural Economics* 84:386-390.

Kumbhakar, S. C. 2001. "Risk Preferences under Price Uncertainties and Production Risk." *Communications in Statistics – Theory and Methods* 30:1715-1735.

Kumbhakar, S. C., and R. Tveterås. 2003. "Risk Preferences, Production Risk and Firm Heterogeneity." *Scandinavian Journal of Economics* 105:275-293.

Love, H. A., and S. T. Buccola. 1991. "Joint Risk Preference-Technology Estimation with a Primal System." *American Journal of Agricultural Economics* 73:765-774.

Meyer, D. J., and J. Meyer. 2005. "Relative Risk Aversion: What Do We Know?" *Journal of Risk and Uncertainty* 31:243-262.

Miranda, M., and P. Fackler. 2002. *Applied Computational Economics and Finance*. Cambridge, MA: MIT Press.

Ogaki, M. 1993. "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.

Pennacchi, G. 2007. *Theory of Asset Pricing*. Boston: Pearson-Addison Wesley.

Prelec, D. 1998. "The Probability Weighting Function." *Econometrica* 66:497-528.

Risk Management Agency. 2008. *Actuarial Data Master Dataset*.

Robison, L. J. 1982. "An Appraisal of Expected Utility Hypothesis Tests Constructed from Responses to Hypothetical Questions and Experimental Choices." *American Journal of Agricultural Economics* 64:367-375.

Saha, A. 1997. "Risk Preference Estimation in the Nonlinear Mean Standard Deviation

Approach." *Economic Inquiry* 35:770-782.

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

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