

Modeling Stochastic Relative Preferences*

Petra M. Geraats[†]

University of Cambridge

June 2005

Abstract

Stochastic relative preferences are prevalent in the literature, but it appears that modeling them is not trivial. This paper establishes that common stochastic specifications alter average relative preferences, which could induce spurious effects. A simple solution is presented that provides an unbiased specification that parameterizes pure white noise shocks to relative preferences. The importance of the results is illustrated by some instructive examples from consumer choice, monetary policy and micro-founded business cycle models.

Keywords: stochastic preferences, preference uncertainty

JEL codes: D1, D8, E2, E5

*Acknowledgments: I thank participants of the Econometric Society European Meetings (ESEM) in Madrid, and seminar participants at Humboldt University in Berlin, St Andrews University, Tilburg University and the University of Cambridge, and especially Seppo Honkapohja, Jim Mirrlees, Bob Rowthorn and Harald Uhlig for useful comments, and Chiaki Hara and Andrew Harvey for helpful discussion. Part of this paper was written while I was a visiting scholar at the Research Department of the International Monetary Fund, which I thank for its hospitality. All views expressed in this paper are entirely mine.

[†]Faculty of Economics, University of Cambridge, Cambridge, CB3 9DD, United Kingdom. Email: Petra.Geraats@econ.cam.ac.uk.

1 Introduction

Models with stochastic relative preferences are prevalent in many areas of economics, for instance in the literature on monetary policy and micro-founded business cycle models. Relative preferences refer to the desirability of one good compared to others; they could be stochastic due to genuine disturbances to preferences or Bayesian uncertainty. Modeling stochastic relative preferences is not trivial. This paper shows that common specifications are biased in the sense that they amount to stochastic disturbances that affect average relative preferences. This change in relative preferences could lead to arbitrary and misleading conclusions. A solution is presented that generates pure white noise shocks to relative preferences for a large class of problems. This unbiased stochastic specification prevents spurious results.

Stochastic relative preferences arise naturally when the good that is preferred (e.g. ice cream versus hot chocolate) depends on the state of nature (e.g. the weather). Typically, the states of nature are implicit and stochastic preferences are modeled as a random shock to a utility parameter. The literature has made abundant use of three kinds of stochastic specifications for relative preferences, unaware of the fact that each tends to be biased. For concreteness, consider the simple objective function $U = \alpha_1 u_1(x_1) + \alpha_2 u_2(x_2)$, where $u_i(x_i)$ is the sub-utility function for good x_i .

The first approach is to apply an additive, white noise shock to one of the relative preference parameters, say, α_1 and to normalize the other, α_2 . The problem is that this specification alters average relative preferences (see Proposition 2). In particular, it effectively makes the absolute value of the relative preference weight α_1/α_2 biased towards zero, thereby reducing the weight put on x_1 . This parameterization is often used in monetary policy games with asymmetric information, and Beetsma and Jensen (2003) have pointed out the arbitrary effects on economic outcomes that arise depending on which preference parameter is normalized.

The second approach is to apply additive white noise shocks of equal size but opposite sign to both coefficients α_1 and α_2 . It was first used by Sørensen (1991) to model ‘pure uncertainty effects’ to relative policy preferences. It turns out that this specification is only unbiased for a particular parameter configuration, although it tends to give rise to ‘stochastic neutrality’ such that the expected value of a variable of interest corresponds to the deterministic outcome (see Section 3).

The third approach is to apply a multiplicative, lognormal shock κ to either α_1 or α_2 , with $E[\ln \kappa] = 0$. This is common in micro-founded business cycle models, including new open economy macroeconomics. This specification also tends to alter average

relative preferences; it makes the absolute value of the marginal rate of substitution biased towards one (see Proposition 5). In addition, although the parameterizations $\kappa\alpha_1$ and $\kappa\alpha_2$ lead to the same stochastic distribution of the marginal rate of substitution, it is problematic that their welfare effects are generally different. In the case of optimization under preference uncertainty, they could even change the qualitative effect on economic outcomes.

The main contribution of this paper is to present a solution to these problems and provide a way to model pure white noise shocks to relative preferences. Unfortunately, it is not fruitful to use the marginal rate of substitution to define a neutral or unbiased stochastic specification that has no average effect on relative preferences. The reason is that the two (reciprocal) definitions of the marginal rate of substitution yield different average results due to Jensen's inequality. Instead, this paper focuses on a more fundamental geometric measure of relative preferences that does not have this drawback and allows for a very natural definition of white noise shocks to relative preferences.

The results in this paper are relevant to the literature on micro-founded business cycle models, which regularly assumes relative preference shocks (e.g. Hall 1997, Obstfeld and Rogoff 2000). Although the use of biased stochastic specifications may be innocuous for calibration exercises, they are problematic when analyzing the economic effects of uncertainty. The change in average relative preferences induced by a biased specification is likely to affect economic outcomes and possibly also welfare effects, which could be incorrectly attributed to the presence of risk. The findings of this paper are also of particular interest to the literature on transparency of monetary policy, where the use of stochastic relative preferences to model monetary uncertainty has generated conclusions that hinge on the use of a biased specification (see Geraats 2002).

The formal analysis of stochastic relative preferences is in Section 2, with the key result, an unbiased specification of the marginal rate of substitution, in Proposition 3. In addition, Proposition 4 presents a simple specification for the stochastic preference parameters α_i that is unbiased in the deterministic optimum and holds for all (interior) stochastic optima when the marginal rate of transformation is independent of the state of preferences. Conveniently, this specification could also be applied to optimization under preference uncertainty. Section 3 explains how the results generalize and apply to heterogeneous preferences or an Arrow-Debreu endowment economy with no aggregate preference uncertainty, which features a constant aggregate demand for all states of nature. Subsequently, Section 4 illustrates some spurious effects resulting from biased specifications in three simple examples related to consumer choice, monetary policy and micro-founded business cycle models. Finally, Section 5 summarizes the main

findings.

2 Stochastic Relative Preferences

Let $U(\mathbf{x}; s)$ denote a stochastic utility function, where $\mathbf{x} = [x_1, \dots, x_K]'$ with $\mathbf{x} \in \mathcal{X}$ is a vector of goods or decision choices x_i in the choice set $\mathcal{X} \subset \mathbf{R}^K$, and s the state of nature, $s \in \mathcal{S}$.¹ The relative preference for any x_i and x_j ($i, j \in \{1, \dots, K\}$, $j \neq i$) in state s is described by the stochastic marginal rate of substitution (MRS), which can be expressed in two equivalent ways:²

$$MRS_{i,j}(\mathbf{x}; s) \equiv \frac{\partial U(\mathbf{x}; s) / \partial x_i}{\partial U(\mathbf{x}; s) / \partial x_j} \quad (1)$$

$$MRS_{j,i}(\mathbf{x}; s) \equiv \frac{\partial U(\mathbf{x}; s) / \partial x_j}{\partial U(\mathbf{x}; s) / \partial x_i} \quad (2)$$

These two definitions of the MRS are intrinsically related. Cursorily, they are simply reciprocals as $MRS_{i,j}(\mathbf{x}_0; s) = 1/MRS_{j,i}(\mathbf{x}_0; s)$ for any $\mathbf{x}_0 \in \mathcal{X}$. More fundamentally, both correspond to the same tangent hyperplane of the indifference surface at \mathbf{x}_0 , and represent different ways of measuring its angle with respect to the x_i and x_j axes in the two-dimensional subspace spanned by x_i and x_j . To be precise, let $\gamma_{i,j}(\mathbf{x}_0; s)$ denote the angle within the subspace $x_i \times x_j \subset \mathbf{R}^2$ between the hyperplane tangent to the indifference surface $U(\mathbf{x}; s) = U(\mathbf{x}_0; s)$ at point \mathbf{x}_0 for state s , and the hyperplane determined by $\{\mathbf{x} \in \mathbf{R}^K | x_j = 0\}$; and similarly, $\gamma_{j,i}(\mathbf{x}_0; s)$ the angle with respect to the hyperplane determined by $\{\mathbf{x} \in \mathbf{R}^K | x_i = 0\}$. Stochastic shocks to these two related angles affect the marginal rates of substitution in a specific fashion.

Proposition 1 *Without loss of generality, suppose that relative preferences are described by*

$$\gamma_{i,j}(\mathbf{x}; s) = \bar{\gamma}_{i,j}(\mathbf{x}) + \zeta_{i,j}(s; \mathbf{x}) \quad (3)$$

where $\bar{\gamma}_{i,j}(\mathbf{x})$ is deterministic and $\zeta_{i,j}(s; \mathbf{x})$ is stochastic with the conditional probability (mass or density) function $f_\zeta(\zeta_{i,j} | \mathbf{x})$. Then, the marginal rate of substitution for $U(\mathbf{x}; s)$ associated with such stochastic relative preferences satisfies

$$MRS_{i,j}(\mathbf{x}; s) = \frac{\tan \bar{\gamma}_{i,j}(\mathbf{x}) + \eta_{i,j}(s; \mathbf{x})}{\cot \bar{\gamma}_{i,j}(\mathbf{x}) - \eta_{i,j}(s; \mathbf{x})} \cot \bar{\gamma}_{i,j}(\mathbf{x}) \quad (4)$$

¹The existence of a random utility representation is presumed. For a survey of axiomatic approaches to stochastic utility, see Fishburn (1998).

²This holds for any \mathbf{x} where local nonsatiation applies. For any bliss points $\mathbf{x} = \mathbf{b}$ where $\partial U(\mathbf{b}; s) / \partial x_i = 0$, $MRS_{j,i}(\mathbf{x}; s)$ is not well-defined but $MRS_{i,j}(\mathbf{x}; s) = 0$ for all $s \in \mathcal{S}$, which means that relative preferences are not stochastic.




Figure 1: The geometry of relative preferences.

$$MRS_{j,i}(\mathbf{x}; s) = \frac{\tan \bar{\gamma}_{j,i}(\mathbf{x}) + \eta_{j,i}(s; \mathbf{x})}{\cot \bar{\gamma}_{j,i}(\mathbf{x}) - \eta_{j,i}(s; \mathbf{x})} \cot \bar{\gamma}_{j,i}(\mathbf{x}) \quad (5)$$

where $\eta_{i,j}(s; \mathbf{x}) \equiv \tan \zeta_{i,j}(s; \mathbf{x})$, $\eta_{j,i}(s; \mathbf{x}) = -\eta_{i,j}(s; \mathbf{x})$, $s \in \mathcal{S}$, $\mathbf{x} \in \mathcal{X}$ and $i, j \in \{1, \dots, K\}$, $i \neq j$.

The formal proof, which appears in the Appendix, proceeds in two steps. First, a geometric argument is used to show that

$$MRS_{i,j}(\mathbf{x}; s) = \tan \gamma_{i,j}(\mathbf{x}; s) \quad (6)$$

$$MRS_{j,i}(\mathbf{x}; s) = \cot \gamma_{i,j}(\mathbf{x}; s) \quad (7)$$

This is illustrated in Figure 1 for an indifference curve $\bar{U} = U(\mathbf{x}; s)$ in (x_i, x_j) space. The MRS corresponds to the slope of the tangent line to the indifference curve and can be described in two ways, as $-dx_j/dx_i$ or $-dx_i/dx_j$, which yields $MRS_{i,j}(\mathbf{x}; s) = \tan \gamma_{i,j}(\mathbf{x}; s)$ and $MRS_{j,i}(\mathbf{x}; s) = \tan \gamma_{j,i}(\mathbf{x}; s)$, respectively. Using the fact that $\tan \gamma_{j,i}(\mathbf{x}; s) = \cot \gamma_{i,j}(\mathbf{x}; s)$ gives the result. The second step of the proof consists of substituting (3) and using trigonometric properties to obtain (4) and (5).

Proposition 1 shows that $MRS_{i,j}$ and $MRS_{j,i}$ are isomorphic. This is a necessary condition for any proper specification of a stochastic MRS. In the special (possibly hypothetical) state s_0 (with $\Pr\{s_0\} \geq 0$) in which there is no stochastic disturbance and $\zeta_{i,j}(s_0; \mathbf{x}) = 0$, the deterministic expressions for the MRS are obtained: $MRS_{i,j}(\mathbf{x}; s_0) = \tan \bar{\gamma}_{i,j}(\mathbf{x})$ and $MRS_{j,i}(\mathbf{x}; s_0) = \cot \bar{\gamma}_{i,j}(\mathbf{x})$.

The fact that instead of $MRS_{i,j}$ and $MRS_{j,i}$, $\gamma_{i,j}$ and $\gamma_{j,i}$ can be used to describe stochastic relative preferences is the key insight to obtaining a specification for white noise shocks to relative preferences. It leads to a natural definition of a neutral benchmark in which the stochastic variation only changes second moments without affecting first moments of relative preferences.³

Definition 1 A specification of stochastic relative preferences between x_i and x_j ($i, j \in \{1, \dots, K\}, i \neq j$) is *unbiased* for $\mathbf{x} \in \mathcal{X}$ if $\mathbb{E}[\zeta_{i,j}|\mathbf{x}] = \mathbb{E}[\zeta_{j,i}|\mathbf{x}] = 0$.

So, an unbiased specification of stochastic relative preferences amounts to a white noise disturbance $\zeta_{i,j}(s; \mathbf{x})$, which means a mean-preserving spread in $\gamma_{i,j}(\mathbf{x}; s)$. Of course, this is equivalent to a mean-preserving spread in $\gamma_{j,i}(\mathbf{x}; s)$. In contrast, introducing a mean-preserving spread in $MRS_{i,j}(\mathbf{x}; s)$ is generally *not* equivalent to a mean-preserving spread in $MRS_{j,i}(\mathbf{x}; s)$ due to Jensen's inequality. But this problem is resolved by focusing on the fundamental measures of relative preferences, $\gamma_{i,j}(\mathbf{x}; s)$ and $\gamma_{j,i}(\mathbf{x}; s)$. With unbiased stochastic relative preferences, observers in (x_i, x_j) space and (x_j, x_i) space both agree that the position of the tangent hyperplane to the indifference surface at \mathbf{x} has not changed on average ($\Delta \mathbb{E}[\gamma_{i,j}|\mathbf{x}] = \Delta \mathbb{E}[\gamma_{j,i}|\mathbf{x}] = 0$).

Any biased specification of stochastic relative preferences can be decomposed into a deterministic change in $\bar{\gamma}_{i,j}(\mathbf{x})$ and a white noise disturbance. This means that the outcome of a biased stochastic specification is distorted by the fact that it effectively incorporates a change in relative preferences $\bar{\gamma}_{i,j}$ instead of merely generating a second moment effect. This makes it desirable to use an unbiased specification to analyze the effects of stochastic relative preferences.

2.1 Biased Specifications

It appears that commonly used stochastic specifications are biased. Most parameterizations of stochastic relative preferences in the literature employ a utility function of the form

$$U(\mathbf{x}; s) = h\left(\sum_{i=1}^K \alpha_i(s) u_i(x_i)\right) \quad (8)$$

³Note that a specification that is unbiased for some $\mathbf{x} \in \mathcal{X}$ may not be unbiased for all $\mathbf{x} \in \mathcal{X}$.

where $h(\cdot)$ is a monotonic function so $h'(\cdot) \neq 0$, and $u_i(\cdot)$ is a differentiable sub-utility function. The relative preference parameters $\alpha_i(s)$ are stochastic and depend on the state of nature $s \in \mathcal{S}$. The fact that $u_i(x_i)$ is independent of the state of nature implies that *absolute* preferences reflected by bliss points are deterministic, so that only *relative* preferences are stochastic. The MRS equals

$$MRS_{i,j}(\mathbf{x}; s) = \frac{\alpha_i(s) u'_i(x_i)}{\alpha_j(s) u'_j(x_j)} \text{ and } MRS_{j,i}(\mathbf{x}; s) = \frac{\alpha_j(s) u'_j(x_j)}{\alpha_i(s) u'_i(x_i)} \quad (9)$$

The deterministic case indicated by s_0 is denoted by $\alpha_i(s_0) = \bar{\alpha}_i$ for all $i \in \{1, \dots, K\}$. Equating (6) and (9) for state s_0 yields

$$\tan \bar{\gamma}_{i,j}(\mathbf{x}) = \frac{\bar{\alpha}_i u'_i(x_i)}{\bar{\alpha}_j u'_j(x_j)} \quad (10)$$

This shows how the deterministic angle $\bar{\gamma}_{i,j}$ corresponds to the deterministic relative preference parameters $\bar{\alpha}_i$ and $\bar{\alpha}_j$.

The following result establishes that an additive, white noise shock to one of the relative preference parameters affects average relative preferences.

Proposition 2 *A utility function (8) for which $\alpha_i(s) = \bar{\alpha}_i + \xi(s)$ and $\alpha_j(s) = \bar{\alpha}_j \neq 0$ for all $j \neq i$, $i, j \in \{1, \dots, K\}$, where $E[\xi] = 0$ and $\text{sgn}(\alpha_i(s)) = \text{sgn}(\bar{\alpha}_i) \neq 0$ for $s \in \mathcal{S}$, is a biased specification of stochastic relative preferences between x_i and x_j for all $\mathbf{x} \in \mathcal{X}$.*

The proof appears in the Appendix.⁴ It first computes the relative preference shock $\zeta_{i,j}$ implied by this specification and subsequently shows that $E[\xi] = 0$ implies that $E[\zeta_{i,j}|\mathbf{x}] \neq 0$ so that the specification is biased. The Proposition implies that an unbiased specification of stochastic relative preferences would require a biased shock ξ to the preference parameter α_i . Although $\zeta_{i,j} = 0$ amounts to $\xi = 0$, $E[\zeta_{i,j}|\mathbf{x}] = 0$ does not correspond to $E[\xi] = 0$ due to nonlinearities.

For concreteness, consider the case in which $\alpha_i(s)$, $\bar{\alpha}_j$, $u'_i(x_i)$ and $u'_j(x_j)$ are all strictly positive. Then, $\zeta_{i,j}$ is monotonically increasing and concave in ξ , so that ξ is convex in $\zeta_{i,j}$.⁵ This means that an unbiased specification with $E[\zeta_{i,j}|\mathbf{x}] = 0$ requires $E[\xi] > 0$. In addition, $E[\xi] = 0$ implies $E[\zeta_{i,j}|\mathbf{x}] < 0$ so that $E[\gamma_{i,j}|\mathbf{x}] < \bar{\gamma}_{i,j}(\mathbf{x})$. As a result, the white noise shock to α_i effectively lowers the average relative preference weight on x_i and gives rise to a bias.

⁴Incidentally, it is a corollary of Proposition 4 that an additive, symmetric (possibly normal) white noise shock ξ to $\bar{\alpha}_i$ is biased. Nevertheless, the proof in the Appendix provides some useful additional insights.

⁵See (32) and (33) in the Appendix. Table 1 also follows from (33).

Table 1: Direction of bias for the specification in Proposition 2

Effect on $E[\gamma_{i,j} \mathbf{x}]$	$\alpha_i(s) \bar{\alpha}_j > 0$	$\alpha_i(s) \bar{\alpha}_j < 0$
$u'_i(x_i) u'_j(x_j) > 0$	–	+
$u'_i(x_i) u'_j(x_j) < 0$	+	–

Intuitively, a white noise shock to α_i does not affect the expected value of $MRS_{i,j}(\mathbf{x}; s)$, which is linear in α_i . This suggests no average change in relative preferences. However, it increases the expected value of $MRS_{j,i}(\mathbf{x}; s)$, which is convex in α_i . This suggests an average increase in the relative preference weight α_j/α_i . Combining both results suggests a white noise shock to α_i reduces the average preference weight on x_i .

More generally, the direction of the bias induced by the specification in Proposition 2 depends on the sign of $\alpha_i(s) \bar{\alpha}_j$ and $u'_i(x_i) u'_j(x_j)$. The effect on $E[\gamma_{i,j}|\mathbf{x}]$ is presented in Table 1. It appears that the effect depends on the sign of $\gamma_{i,j}$ in such a way that the specification makes the absolute value of $E[\gamma_{i,j}|\mathbf{x}]$ biased towards zero. As a result, the white noise shock to α_i effectively diminishes the relative preference weight put on x_i .

Another common specification of stochastic relative preferences is to apply a multiplicative shock κ to one of the relative preference parameters, where $\ln \kappa$ is white noise. It turns out that this parameterization is also typically biased, with one exception. When the absolute value of the MRS is equal to one in the deterministic optimum, this specification is unbiased. This result is shown in Proposition 5 in the next subsection after more general unbiased specifications are presented.

2.2 Unbiased Specifications

It appears that it is typically not possible to obtain an unbiased specification for $U(\mathbf{x}; s)$ for all $\mathbf{x} \in \mathcal{X}$ with preference parameters $\alpha_i(s)$ that are independent of \mathbf{x} . However, this paper provides two approaches that ensure unbiasedness for constrained optimization problems involving (8). The first approach, presented in Proposition 3, is the most general and formulates an unbiased specification of the MRS for any \mathbf{x} . The second approach, given by Proposition 4, provides a simple specification of the utility function $U(\mathbf{x}; s)$ that is only unbiased at the deterministic optimum $\bar{\mathbf{x}}$, but valid for all stochastic optima $\mathbf{x}(s)$, and applies to the large class of economic problems that feature a unique interior solution and a marginal rate of transformation that is independent of the preference shocks (in equilibrium).

The first result provides an unbiased stochastic specification of the MRS that follows from Proposition 1 in a quite straightforward way and corresponds to symmetric white noise to the angle $\gamma_{i,j}$.

Proposition 3 *For the utility function (8), the marginal rate of substitution*

$$MRS_{i,j}(\mathbf{x}; s) = \frac{1 + \xi_{i,j}(s; \mathbf{x}) \bar{\alpha}_i u'_i(x_i)}{1 + \xi_{j,i}(s; \mathbf{x}) \bar{\alpha}_j u'_j(x_j)} \quad (11)$$

provides an unbiased specification of stochastic relative preferences between x_i and x_j for all $\mathbf{x} \in \mathcal{X}$, where $\xi_{i,j}(s; \mathbf{x}) \equiv \frac{\bar{\alpha}_j u'_j(x_j)}{\bar{\alpha}_i u'_i(x_i)} \eta_{i,j}(s; \mathbf{x})$, $\eta_{j,i}(s; \mathbf{x}) = -\eta_{i,j}(s; \mathbf{x})$ with the conditional probability function $f_\eta(\eta_{i,j}|\mathbf{x}) = f_\eta(-\eta_{i,j}|\mathbf{x})$ and $i, j \in \{1, \dots, K\}$, $i \neq j$.

Proof. Substitute (10) into (4) and rearrange to get

$$MRS_{i,j}(\mathbf{x}; s) = \frac{\frac{\bar{\alpha}_i u'_i(x_i)}{\bar{\alpha}_j u'_j(x_j)} + \eta_{i,j}(s; \mathbf{x}) \bar{\alpha}_j u'_j(x_j)}{\frac{\bar{\alpha}_j u'_j(x_j)}{\bar{\alpha}_i u'_i(x_i)} - \eta_{i,j}(s; \mathbf{x}) \bar{\alpha}_i u'_i(x_i)} = \frac{1 + \frac{\bar{\alpha}_j u'_j(x_j)}{\bar{\alpha}_i u'_i(x_i)} \eta_{i,j}(s; \mathbf{x}) \bar{\alpha}_i u'_i(x_i)}{1 - \frac{\bar{\alpha}_i u'_i(x_i)}{\bar{\alpha}_j u'_j(x_j)} \eta_{i,j}(s; \mathbf{x}) \bar{\alpha}_j u'_j(x_j)}$$

Using $\eta_{j,i}(s; \mathbf{x}) = -\eta_{i,j}(s; \mathbf{x})$ and the definition of $\xi_{i,j}(s; \mathbf{x})$ gives (11). Finally, using $\eta_{i,j}(s; \mathbf{x}) = \tan \zeta_{i,j}(s; \mathbf{x})$ and $f_\eta(\eta_{i,j}|\mathbf{x}) = f_\eta(-\eta_{i,j}|\mathbf{x})$ it follows that $f_\zeta(\zeta_{i,j}|\mathbf{x}) = \left| 1 + (\tan \zeta_{i,j})^2 \right| f_\eta(\tan \zeta_{i,j}|\mathbf{x}) = f_\zeta(-\zeta_{i,j}|\mathbf{x})$, which implies that $E[\zeta_{i,j}|\mathbf{x}] = 0$ so that the specification is unbiased.⁶ ■

This unbiased specification of the MRS for stochastic relative preferences effectively amounts to a multiplicative adjustment of both preference parameters $\bar{\alpha}_i$ and $\bar{\alpha}_j$ which depends on the deterministic MRS, $\frac{\bar{\alpha}_i u'_i(x_i)}{\bar{\alpha}_j u'_j(x_j)}$. The preference shocks satisfy $E[\xi_{i,j}|\mathbf{x}] = E[\xi_{j,i}|\mathbf{x}] = 0$ and $\text{Cov}\{\xi_{i,j}, \xi_{j,i}|\mathbf{x}\} < 0$, and they exhibit heteroskedasticity. The fact that these properties are conditional on \mathbf{x} is critical when considering several \mathbf{x} for the same state s . Also note that the expressions for $MRS_{i,j}(\mathbf{x}; s)$ and $MRS_{j,i}(\mathbf{x}; s)$ are isomorphic, which is a necessary condition for any proper stochastic MRS.

The specification in Proposition 3 is not only unbiased, but it also reflects a symmetric distribution of relative preference shocks $\zeta_{i,j}$. In principle, it would be possible to obtain unbiased specifications that are based on skewed distributions. However, the assumption of symmetry is appealing because it is simpler and has the feature that the

⁶This shows that $f_\eta(\eta_{i,j}|\mathbf{x}) = f_\eta(-\eta_{i,j}|\mathbf{x})$ is only a sufficient condition. Other conditional distributions for η could also yield $E(\zeta_{i,j}|\mathbf{x}) = 0$.

MRS associated with the median preference shock $\zeta_{i,j} = \eta_{i,j} = 0$ is equal to the deterministic value, $MRS_{i,j}(\mathbf{x}; s_0)$. In the presence of preference heterogeneity according to (11), where $\eta_{i,j}(s; \mathbf{x})$ captures the idiosyncratic preferences of agent s , this means that the MRS of the deterministic case s_0 corresponds to the relative preferences of the median agent s_m : $MRS_{i,j}(\mathbf{x}; s_0) = MRS_{i,j}(\mathbf{x}; s_m)$.

It may be tempting to consider a utility function with relative preference parameters $\alpha_i(s) = (1 + \xi_{i,j}(s; \mathbf{x})) \bar{\alpha}_i$ and $\alpha_j(s) = (1 + \xi_{j,i}(s; \mathbf{x})) \bar{\alpha}_j$, but note that such a specification would generally not lead to the unbiased MRS in (11) because of the dependence of the preference shocks on \mathbf{x} . However, there is an unbiased utility specification that applies to a common situation. The constrained optimization problem under consideration tends to yield a unique interior solution $\mathbf{x}(s)$ that equates the MRS to the marginal rate of transformation (MRT). For many economic problems, the MRT is independent of the relative preference shocks (in equilibrium). In such cases, there is a simple utility specification of stochastic relative preferences that is unbiased at $\bar{\mathbf{x}}$ and holds for $\mathbf{x}(s)$.

Proposition 4 *Suppose a constrained optimization problem based on the utility function (8) has a unique interior solution $\mathbf{x}(s)$ satisfying*

$$MRS_{i,j}(\mathbf{x}(s); s) = \lambda \quad (12)$$

for all states of nature $s \in \mathcal{S}$. Then,

$$\alpha_i(s) = [1 + \lambda \xi(s)] \bar{\alpha}_i \quad (13)$$

$$\alpha_j(s) = \left[1 - \frac{1}{\lambda} \xi(s)\right] \bar{\alpha}_j \quad (14)$$

$$\alpha_k(s) = \left[1 - \frac{1}{\lambda} \xi(s)\right] \bar{\alpha}_k \quad (15)$$

where $f_\xi(\xi) = f_\xi(-\xi)$, provides a specification of stochastic relative preferences between x_i and x_j ($i, j \in \{1, \dots, K\}, i \neq j$) that is unbiased at the deterministic optimum $\bar{\mathbf{x}}$ for which $\alpha_n(s) = \bar{\alpha}_n$ for all $n \in \{1, \dots, K\}$, holds for the stochastic optima $\mathbf{x}(s)$ for all $s \in \mathcal{S}$, and does not affect the relative preferences between x_j and x_k for all $k \in \{1, \dots, K\} \setminus \{i, j\}$.

Proof. Substituting the unbiased MRS (11) into (12) and simplifying gives

$$MRS_{i,j}(\mathbf{x}(s); s) = \frac{\frac{\bar{\alpha}_i u'_i(x_i)}{\bar{\alpha}_j u'_j(x_j)} + \eta_{i,j}(s; \mathbf{x}(s))}{1 - \frac{\bar{\alpha}_i u'_i(x_i)}{\bar{\alpha}_j u'_j(x_j)} \eta_{i,j}(s; \mathbf{x}(s))} = \lambda$$

Denoting $\xi(s) \equiv \eta_{i,j}(s; \mathbf{x}(s))$ and rearranging yields the following Euler equation:

$$[1 + \lambda\xi(s)] \bar{\alpha}_i u'_i(x_i) = [\lambda - \xi(s)] \bar{\alpha}_j u'_j(x_j) \quad (16)$$

This can be rewritten as

$$\lambda = \frac{[1 + \lambda\xi(s)] \bar{\alpha}_i u'_i(x_i)}{\left[1 - \frac{1}{\lambda}\xi(s)\right] \bar{\alpha}_j u'_j(x_j)} = MRS_{i,j}(\mathbf{x}(s); s)$$

where the latter MRS is constructed using (13) and (14). In addition, $\xi(s) = \eta_{i,j}(s; \mathbf{x}(s)) = 0$ yields the deterministic case with $MRS_{i,j}(\bar{\mathbf{x}}; s)$, so the specification is unbiased at $\bar{\mathbf{x}}$. Using (14) and (15), $MRS_{j,k}(\mathbf{x}(s); s)$ is constant for all s , so that the relative preferences between x_j and x_k are not affected. ■

The proof yields the stochastic Euler equation (16), which provides a convenient way to characterize the stochastic optima $\mathbf{x}(s)$ associated with this unbiased specification. It also shows how the results in Propositions 3 and 4 can be reconciled. An alternative proof, which does not rely on the result in Proposition 3 but directly establishes unbiasedness, appears in the Appendix.

The unbiased specification in Proposition 4 is less general than in Proposition 3, but it is applicable to any problem which has a unique interior solution and for which the MRT is independent of the preference shock (in equilibrium). So, it holds whenever the MRT stems from a technological or structural constraint that is linear in \mathbf{x} but otherwise independent of the preference shock. For instance, a policymaker could make a decision about the policy variables $\mathbf{x}(s)$ for each state of nature s subject to a linear constraint that reflects the structure of the economy. In addition, the specification can be used to model the case of preference uncertainty in which an agent decides about \mathbf{x} under uncertainty about its own preferences or without knowing the realization of the state of nature s .

The unbiased specification of Proposition 4 could also be described by

$$\alpha_n(s) = (1 + \xi_n(s)) \bar{\alpha}_n \quad (17)$$

for every $n \in \{1, \dots, K\}$, where $\xi_i(s) = \lambda\xi(s)$, $\xi_j(s) = -\frac{1}{\lambda}\xi(s)$ and $\xi_k(s) = \xi_j(s)$, with $E[\xi] = 0$ and $\text{Var}[\xi] = \sigma_\xi^2 > 0$. As a result, the preference shocks have the property that $E[\xi_i] = E[\xi_j] = 0$, $\text{Var}[\xi_i] = \lambda^2\sigma_\xi^2$, $\text{Var}[\xi_j] = \frac{1}{\lambda^2}\sigma_\xi^2$, and $\text{Cov}\{\xi_i, \xi_j\} = -\sigma_\xi^2$.

This shows how unbiasedness could be achieved by applying additive, correlated, heteroskedastic shocks $\xi_n \bar{\alpha}_n$ to $\bar{\alpha}_n$. In addition, it indicates that an additive white noise shock to only one of the preference parameters $\bar{\alpha}_i$ produces a bias, as was formally

shown in Proposition 2. Another common stochastic specification is a multiplicative shock κ to α_i where $\ln \kappa$ is white noise. As mentioned above, such a specification is typically biased, except for the special case in which the absolute value of the MRS equals one in the deterministic optimum.

Proposition 5 *If the deterministic optimum $\bar{\mathbf{x}}$ satisfies $|MRS_{i,j}(\bar{\mathbf{x}})| = 1$, then there exists a $\kappa(s)$ with $E[\ln \kappa] = 0$ such that the utility function $U(\mathbf{x}; s)$ in (8) with $\alpha_i(s) = \kappa(s)\bar{\alpha}_i$ and $\alpha_j(s) = \bar{\alpha}_j \neq 0$ for all $j \neq i, i, j \in \{1, \dots, K\}$, is a specification of stochastic relative preferences between x_i and x_j that is unbiased at $\bar{\mathbf{x}}$.*

The proof appears in the Appendix. It uses a Taylor series expansion of the log of the unbiased MRS in (11) around the deterministic case to show that the specification with $\alpha_i = \kappa\bar{\alpha}_i$ and $\alpha_j = \bar{\alpha}_j$ is observationally equivalent for a particular κ for which $\ln \kappa$ is white noise if the absolute value of the deterministic MRS equals one. In that case, $\ln MRS_{i,j}(\bar{\mathbf{x}}; s)$ and $\ln MRS_{j,i}(\bar{\mathbf{x}}; s)$ are (symmetric) white noise for both specifications.

Without the restriction on the deterministic MRS, the proof shows that there still exists a κ that produces unbiasedness, but now $\ln \kappa$ is no longer white noise. Instead, it has the property that $E[\ln \kappa]$ has the same sign as $|MRS_{i,j}(\bar{\mathbf{x}})| - 1$. So, in general a specification with $\ln \kappa$ white noise displays a systematic bias. In particular, for $|MRS_{i,j}(\bar{\mathbf{x}})| > 1$ unbiasedness requires $E[\ln \kappa] > 0$, so the specification with $E[\ln \kappa] = 0$ effectively reduces the relative preference weight on x_i in absolute value and lowers $|MRS_{i,j}(\bar{\mathbf{x}})|$. Stated differently, for $|MRS_{i,j}(\bar{\mathbf{x}})| \neq 1$, a log white noise shock makes $|MRS_{i,j}(\bar{\mathbf{x}})|$ biased towards 1. This suggests that one should be careful with the interpretation of results based on the commonly used multiplicative, lognormal specification with $\alpha_i = \kappa\bar{\alpha}_i$, where $\ln \kappa \sim N(0, \sigma_\kappa^2)$.⁷

3 Discussion

The unbiased specifications in Propositions 3 and 4 are based on the commonly used utility function (8), but it is straightforward to extend the results to other utility functions. Of course, the expressions for the MRS in (9) should be updated, as well as (10), but once they are replaced, Proposition 1 still holds. To generalize the specification of the unbiased MRS in Proposition 3, simply substitute $\tan \bar{\gamma}_{i,j}(\mathbf{x})$ for $\frac{\bar{\alpha}_i u'_i(x_i)}{\bar{\alpha}_j u'_j(x_j)}$. Then, one can proceed along the steps in the proof of Proposition 4 to obtain an unbiased

⁷Note that Proposition 5 does not guarantee that this lognormal specification is unbiased at $\bar{\mathbf{x}}$ for $|MRS_{i,j}(\bar{\mathbf{x}})| = 1$, although numerical examples indicate it is nearly so.

stochastic Euler equation or even utility specification for the case in which the MRT is state-independent (in equilibrium).

In practice, a stochastic specification may be used without an explicit deterministic benchmark. When the specification involves a continuous, parametric distribution of a utility parameter α , it is natural to take as deterministic benchmark the limit as the variance $\text{Var}[\alpha]$ goes to zero. When the distribution is nonparametric or the states of nature are discrete, the mean or median of α provides a useful benchmark to assess to what extent the stochastic specification distorts average relative preferences.

It is important to realize that an unbiased specification of stochastic relative preferences typically does not have a stochastically neutral effect on the economic outcome \mathbf{x} in the sense that the stochastic optima $\mathbf{x}(s)$ merely add white noise to the deterministic optimum $\bar{\mathbf{x}}$.⁸

Definition 2 A stochastic optimization problem exhibits *stochastic neutrality* in \mathbf{x} if the corresponding stochastic optima $\mathbf{x}(s)$ satisfy $E[\mathbf{x}] = \bar{\mathbf{x}}$, where $\bar{\mathbf{x}}$ denotes the deterministic optimum.

Conceptually, the main distinction between unbiasedness of a preference specification and stochastic neutrality is that the former reflects a pure uncertainty effect on preferences, whereas the latter features a pure uncertainty effect on outcomes. Unbiasedness of a preference specification does not imply stochastic neutrality in \mathbf{x} because the optimum $\mathbf{x}(s)$ results from a combination of preferences $U(\mathbf{x}; s)$ and constraints, $H(\mathbf{x}; s) = 0$, which typically involves a nonlinear interaction. As a result, an unbiased specification of stochastic preferences coincides with stochastic neutrality only in some special cases.⁹ Similarly, stochastic neutrality in \mathbf{x} could hold despite a biased preference specification.

The presence of stochastic neutrality in \mathbf{x} could be analytically convenient for optimization problems in which the constraint depends on $E[\mathbf{x}]$, because the latter is not affected by preference variability. An example would be modeling a policymaker who decides about inflation and unemployment and faces a Phillips curve that depends on expected inflation. Furthermore, stochastic neutrality in \mathbf{x} is desirable when the focus of the analysis is on \mathbf{x} and preference shocks are merely an auxiliary assumption to capture asymmetric information about preferences. Nevertheless, researchers should

⁸Stochastic neutrality is not implied by risk neutral preferences. In addition, it is different from certainty equivalence, in which case $\mathbf{x}(s) = \bar{\mathbf{x}}$ for all s . So, certainty equivalence implies stochastic neutrality, but not vice versa.

⁹One example is the case in which $U(\mathbf{x}; s) = \left(1 + \frac{\beta_1}{\beta_2}\xi(s)\right)x_1^2 + \left(1 - \frac{\beta_2}{\beta_1}\xi(s)\right)x_2^2$ and $H(\mathbf{x}; s) = \beta_0 + \beta_1x_1 + \beta_2x_2$, with $f_\xi(\xi) = f_\xi(-\xi)$.

be aware of the fact that stochastic neutrality typically involves a change in average relative preferences that could induce spurious welfare effects. This does not arise with an unbiased specification.

The present setup for modeling stochastic preferences of a representative agent could also be applied to preference heterogeneity among multiple agents. In particular, use $s = n$ to index agents instead of states of nature. Then, stochastic neutrality in \mathbf{x} amounts to an average optimal outcome of $\bar{\mathbf{x}}$. For instance, consider an endowment economy in which N consumers have heterogeneous preferences about \mathbf{x} . Then, stochastic neutrality in \mathbf{x} means that aggregate demand equals $N\bar{\mathbf{x}}$ and therefore that relative prices for \mathbf{x} are the same regardless of the degree of preference heterogeneity. This result can be extended to an endowment economy with heterogeneous and stochastic preferences. In that case, stochastic neutrality in \mathbf{x} for each state of nature implies that for each state aggregate demand equals $N\bar{\mathbf{x}}$ so that relative prices and the marginal rates of transformation are constant in equilibrium. This means that Proposition 4 also applies to such an endowment economy with no aggregate preference uncertainty.¹⁰

Definition 3 An endowment economy inhabited by N consumers with heterogeneous and stochastic preferences features *no aggregate preference uncertainty* if aggregate demand, $\sum_{n=1}^N \mathbf{x}_n(s)$, is constant for all states of nature $s \in \mathcal{S}$, where $\mathbf{x}_n(s)$ denotes the individual demand function for consumer n in state s .

So, an endowment economy with no aggregate preference uncertainty and with the stochastic specification in Proposition 4 features preferences that are unbiased in the deterministic equilibrium. Except for these special cases, the concepts of unbiasedness, stochastic neutrality and no aggregate preference uncertainty are generally not related. The next section provides several examples that illustrate these concepts further.

4 Examples

The instructive examples in this section focus on consumer choice, monetary policy and micro-founded business cycle models.

¹⁰This is exactly analogous to an Arrow-Debreu endowment economy with no aggregate (endowment) uncertainty, in which case aggregate supply, $\sum_{n=1}^N \omega_n(s)$, is constant for all states of nature s , where $\omega_n(s)$ denotes the endowment of consumer n in state s .

4.1 Consumer Choice

Consider a simple consumer optimization problem under perfect competition with two goods x_1 and x_2 . The representative consumer maximizes the utility function

$$U(x_1, x_2) = \alpha_1 \ln x_1 + \alpha_2 \ln x_2 \quad (18)$$

subject to the budget constraint

$$p_1 x_1 + p_2 x_2 = m \quad (19)$$

where p_i is the price of good i , m denotes nominal assets (in terms of a fictitious numéraire good), and $\alpha_i > 0$, with $i \in \{1, 2\}$. The marginal rates of substitution equal

$$MRS_{1,2}(\mathbf{x}) = \frac{\alpha_1 x_2}{\alpha_2 x_1} \text{ and } MRS_{2,1}(\mathbf{x}) = \frac{\alpha_2 x_1}{\alpha_1 x_2}.$$

Maximizing (18) subject to (19) yields the optimal demand relationships

$$x_1 = \frac{\alpha_1}{\alpha_1 + \alpha_2} \frac{m}{p_1} \text{ and } x_2 = \frac{\alpha_2}{\alpha_1 + \alpha_2} \frac{m}{p_2}. \quad (20)$$

Two cases are considered: First, partial equilibrium, in which p_1 and p_2 are exogenous and deterministic; subsequently, general equilibrium in which the stochastic preferences of the representative consumer affect the equilibrium price.

4.1.1 Partial Equilibrium

There are several ways in which stochastic relative preferences between x_1 and x_2 could be modeled.

First, suppose that $\alpha_1 = \bar{\alpha}_1 + \xi$ and $\alpha_2 = 1$, where ξ is white noise. This has no average effect on $MRS_{1,2}(\mathbf{x})$, which is linear in α_1 , but it increases the expected value of $MRS_{2,1}(\mathbf{x})$, which depends inversely on α_1 . This asymmetry is peculiar, because $MRS_{1,2}(\mathbf{x})$ and $MRS_{2,1}(\mathbf{x})$ are two equivalent ways of measuring the slope of the indifference curve in (x_1, x_2) and (x_2, x_1) space, respectively, and they are simply mirror images of each other. Furthermore, this specification of stochastic preferences leads to a decrease in the expected demand for good x_1 (and an increase for x_2) since x_1 is concave (and x_2 is convex) in α_1 .

Now, suppose that $\alpha_1 = 1$ and $\alpha_2 = \bar{\alpha}_2 + \xi$. Then, the expected value of $MRS_{1,2}(\mathbf{x})$ increases, but there is no average effect on $MRS_{2,1}(\mathbf{x})$. In addition, there is an increase in the expected demand for x_1 (and a decrease for x_2). Clearly, the normalization of α_i is not innocuous in the presence of stochastic relative preferences. More precisely, a

white noise shock to α_i effectively reduces the average relative preference that the consumer attaches to x_i , which is consistent with the conclusion obtained from Proposition 2.

An unbiased specification for the MRS is given by Proposition 3:

$$MRS_{1,2}(\mathbf{x}) = \frac{1 + \frac{\bar{\alpha}_2 x_1}{\bar{\alpha}_1 x_2} \eta \bar{\alpha}_1 x_2}{1 - \frac{\bar{\alpha}_1 x_2}{\bar{\alpha}_2 x_1} \eta \bar{\alpha}_2 x_1} \quad (21)$$

where η is symmetric white noise conditional on \mathbf{x} , so $f_\eta(\eta|\mathbf{x}) = f_\eta(-\eta|\mathbf{x})$. To illustrate unbiasedness, consider a simple numerical example. Suppose there are two equally likely states of nature $s \in \{1, 2\}$ with $\eta(1) = +\eta$ and $\eta(2) = -\eta$, where $0 < \eta < 1$. Assume that $\bar{\alpha}_1 = \bar{\alpha}_2 = 1$ and focus on $x_1 = x_2 = 1$. Then, $\bar{\gamma}_{1,2}(\mathbf{x}) = \arctan 1 = \frac{1}{4}\pi$. In addition, $\gamma_{1,2}(\mathbf{x}; 1) = \arctan \frac{1+\eta}{1-\eta}$ and $\gamma_{1,2}(\mathbf{x}; 2) = \arctan \frac{1-\eta}{1+\eta}$, so $E[\gamma_{1,2}|\mathbf{x}] = \frac{1}{4}\pi = \bar{\gamma}_{1,2}(\mathbf{x})$ and the specification is unbiased.¹¹

To derive the demand function associated with the unbiased MRS, substitute (21) into the first order condition for optimization, $MRS_{1,2}(\mathbf{x}) = \frac{p_1}{p_2}$, and rearrange to get the stochastic Euler equation

$$\left(1 + \frac{p_1}{p_2} \eta\right) \frac{\bar{\alpha}_1}{x_1} = \left(\frac{p_1}{p_2} - \eta\right) \frac{\bar{\alpha}_2}{x_2} \quad (22)$$

Solve this for x_2 and substitute it into (19) to find

$$x_1 = \frac{\left(1 + \frac{p_1}{p_2} \eta\right) \bar{\alpha}_1}{\left(1 + \frac{p_1}{p_2} \eta\right) \bar{\alpha}_1 + \left(1 - \frac{p_2}{p_1} \eta\right) \bar{\alpha}_2 p_1} m \quad (23)$$

Now consider the stochastic specification in Proposition 4 with $\alpha_1 = \left(1 + \frac{p_1}{p_2} \xi\right) \bar{\alpha}_1$ and $\alpha_2 = \left(1 - \frac{p_2}{p_1} \xi\right) \bar{\alpha}_2$, where ξ is symmetric white noise, so $f_\xi(\xi) = f_\xi(-\xi)$. Then,

$$MRS_{1,2}(\mathbf{x}) = \frac{1 + \frac{p_1}{p_2} \xi \bar{\alpha}_1 x_2}{1 - \frac{p_2}{p_1} \xi \bar{\alpha}_2 x_1}$$

This MRS is only unbiased at the deterministic optimum $\bar{\mathbf{x}}$.¹² However, it applies to all (interior) stochastic optima \mathbf{x} that satisfy $MRS_{1,2}(\mathbf{x}) = \frac{p_1}{p_2}$, and for $\xi = \eta$ it yields the same stochastic Euler equation (22) and demand function (23).

¹¹Use the trigonometric identities $\arctan z + \operatorname{arccot} z = \pi/2$ and $\operatorname{arccot} z = \arctan(1/z)$.

¹²Note that this MRS cannot be used to check whether $E[\gamma_{1,2}|\bar{\mathbf{x}}] = E[\arctan MRS_{1,2}|\bar{\mathbf{x}}] = \bar{\gamma}_{1,2}(\bar{\mathbf{x}})$, because evaluating it at $\bar{\mathbf{x}}$ for any state different from s_0 violates the condition that $MRS_{1,2}(\mathbf{x}(s); s) = \frac{p_1}{p_2}$.

Expected demand $E[x_1]$ for these unbiased stochastic specifications is typically different from the deterministic outcome $\bar{x}_1 = \frac{\bar{\alpha}_1}{\bar{\alpha}_1 + \bar{\alpha}_2} \frac{m}{p_1}$, so the effect on \mathbf{x} is not stochastically neutral. Inspecting (23) reveals that only the special case in which $\bar{\alpha}_1/\bar{\alpha}_2 = p_2^2/p_1^2$, so that $x_1 = \frac{\alpha_1}{\bar{\alpha}_1 + \bar{\alpha}_2} \frac{m}{p_1}$, features both unbiasedness and stochastic neutrality in \mathbf{x} . A simple parameterization that is stochastically neutral in \mathbf{x} is $\alpha_1 = \bar{\alpha}_1 + \tilde{\xi}$ and $\alpha_2 = 1 - \bar{\alpha}_1 - \tilde{\xi}$, where $\tilde{\xi}$ is symmetric white noise, but this specification is biased unless $p_1/p_2 = 1$.¹³

4.1.2 General Equilibrium

Now consider consumer choice in general equilibrium where ω_i is the endowment of good i for all states of nature s , and $p_i(s)$ is the equilibrium price of good i in state s .¹⁴ As a result, the consumer's nominal assets are stochastic and equal

$$m(s) = p_1(s)\omega_1 + p_2(s)\omega_2 \quad (24)$$

Furthermore, goods market equilibrium requires that for all states of nature s ,

$$x_1(s) = \omega_1 \quad \text{and} \quad x_2(s) = \omega_2. \quad (25)$$

Using (25), (20) and (24) gives the relative prices in equilibrium:

$$\frac{p_1}{p_2} = \frac{\alpha_1 \omega_2}{\alpha_2 \omega_1} \quad \text{and} \quad \frac{p_2}{p_1} = \frac{\alpha_2 \omega_1}{\alpha_1 \omega_2}.$$

Suppose that $\alpha_1 = \bar{\alpha}_1 + \xi$ and $\alpha_2 = 1$, where ξ is white noise. Then, the expected relative price of x_1 is not affected, despite the lower expected demand for x_1 , but there is an average increase in the relative price of x_2 . In contrast, when $\alpha_1 = 1$ and $\alpha_2 = \bar{\alpha}_2 + \xi$, the expected value of the relative price of x_1 increases, consistent with the higher expected demand for x_1 , but there is no average effect on the relative price of x_2 . Again, the results depend on the normalization of the relative preference parameter.

Next, consider the unbiased specification of Proposition 3 or 4. Substituting (24) into (23) and using (25) gives

$$\frac{p_1}{p_2} = \frac{\bar{\alpha}_1 \omega_2 + \bar{\alpha}_2 \omega_1 \xi}{\bar{\alpha}_2 \omega_1 - \bar{\alpha}_1 \omega_2 \xi}$$

¹³More generally, for $u_k(x_k) = \ln x_k$ and $H(\mathbf{x}; s) = \beta_0 + \sum_k \beta_k x_k$, there is stochastic neutrality in \mathbf{x} if $\alpha_i = (1 + \xi) \bar{\alpha}_i$, $\alpha_j = (1 - \delta \xi) \bar{\alpha}_j$ and $\alpha_k = (1 - \delta \xi) \bar{\alpha}_k$, where $\delta = \bar{\alpha}_i / \sum_{k \neq i} \bar{\alpha}_k$.

¹⁴An alternative approach to equilibrium analysis in an endowment economy with random preferences is to consider a deterministic price vector such that expected excess demand equals zero and excess demand per capita converges in probability to zero. (Hildenbrand 1971)

Now the effect of the stochastic preference shock ξ on the relative prices p_1/p_2 and p_2/p_1 is entirely isomorphic.

Instead of a representative consumer with stochastic preferences, suppose that this economy is inhabited by N consumers with equal endowments but heterogenous preferences. Each consumer can be characterized by the preference parameter $\xi = \eta$, which has the relative frequency function $g(\xi)$. The level of aggregate demand and therefore the equilibrium relative price depend on the consumer heterogeneity. But, when the specification is stochastically neutral in \mathbf{x} , aggregate demand and the equilibrium price are independent of the degree of preference heterogeneity. In particular, consider the special case in which the preference specification is stochastically neutral and unbiased, so demand $x_{1,n}$ by consumer n is given by (23) with $\bar{\alpha}_1/\bar{\alpha}_2 = (p_2/p_1)^2$ and $g(\xi) = g(-\xi)$. Then, substitute (24) and impose the equilibrium condition, $\frac{1}{N} \sum_{n=1}^N x_{1,n} = \omega_1$, to get¹⁵

$$\frac{p_1}{p_2} = \left(\frac{\omega_2}{\omega_1} \right)^{\frac{1}{3}} \quad (26)$$

So, in this special case with $\bar{\alpha}_1/\bar{\alpha}_2 = (\omega_1/\omega_2)^{2/3}$ the economy with heterogenous preferences is observationally equivalent to an economy with a single, representative consumer, which by virtue of the unbiased specification corresponds to the median consumer, regardless of the degree of preference heterogeneity.

Finally, introduce stochastic preferences into this endowment economy with heterogenous consumers. Let $\xi_n(s)$ denote the stochastic preference parameter for consumer n in state of nature s . In particular, take the unbiased stochastic neutrality case with $\bar{\alpha}_1/\bar{\alpha}_2 = (\omega_1/\omega_2)^{2/3}$, where $\xi_n(s)$ has the relative frequency (or frequency density) function $g(\xi_n(s)) = g(-\xi_n(s))$ for each state s and probability function $f_n(\xi_n(s)) = f_n(-\xi_n(s))$ for each consumer n . Then, $\frac{1}{N} \sum_{n=1}^N x_{1,n}(s) = \omega_1$ for all s , so there is no aggregate preference uncertainty. The equilibrium in this economy with heterogeneous and stochastic preferences is now identical to a deterministic representative agent economy with relative price (26), regardless of the degree of consumer heterogeneity and preference variability. This convenient property of both no aggregate preference uncertainty and unbiased stochastic relative preferences for each consumer holds for the preference parameters $\alpha_{1,n}(s) = \left(1 + \left(\frac{\omega_2}{\omega_1} \right)^{\frac{1}{3}} \xi_n(s) \right) \left(\frac{\omega_1}{\omega_2} \right)^{\frac{2}{3}} \bar{\alpha}_2$ and $\alpha_{2,n}(s) = \left(1 - \left(\frac{\omega_1}{\omega_2} \right)^{\frac{1}{3}} \xi_n(s) \right) \bar{\alpha}_2$, which follow from Proposition 4.

This example from consumer choice has illustrated how a biased specification of

¹⁵In case of a continuum of consumers, use the equilibrium condition $\int x_1(\xi) g(\xi) d\xi = \omega_1$, where $g(\xi)$ is the frequency density function, to get the same result.

stochastic relative preferences could alter qualitative conclusions and it has shown useful applications of the unbiased specification proposed in this paper.

4.2 Monetary Policy

Consider a simple monetary policy game in which the central bank maximizes

$$U(p, y) = -\frac{1}{2}\alpha_p p^2 - \frac{1}{2}\alpha_y y^2 \quad (27)$$

where p denotes the (log) aggregate price level, y (log) aggregate output, α_p the preference parameter for price stabilization and α_y the preference parameter for output stabilization ($\alpha_p, \alpha_y > 0$). The structure of the economy is described by the aggregate supply relation

$$y = \theta(p - w) \quad (28)$$

where w is the (log) nominal wage and θ the sensitivity of output to the real wage ($\theta > 0$). Maximizing (27) subject to (28) yields the optimal price and output levels:

$$p = \frac{\alpha_y \theta^2}{\alpha_p + \alpha_y \theta^2} w \text{ and } y = -\frac{\alpha_p \theta}{\alpha_p + \alpha_y \theta^2} w$$

The relative preferences of the monetary policymakers are assumed to be stochastic, for instance to capture asymmetric information about the central bank's objectives.

First, suppose that $\alpha_p = \bar{\alpha}_p + \xi$ and $\alpha_y = 1$, where ξ is white noise. Then, the stochastic variation in preferences causes the expected price level and output to rise because p and y are convex in α_p .

Now, suppose that $\alpha_p = 1$ and $\alpha_y = \bar{\alpha}_y + \xi$. Then, the expected price level and output drop because p and y are concave in α_y . Again, changing the normalization of the relative preference parameter drastically alters the results. In particular, a white noise shock to α_p effectively reduces the central bank's relative preference for price stability, which is associated with less conservativeness, whereas a white noise shock to α_y has the opposite effect and essentially makes the central bank more conservative. These findings are consistent with the general result obtained from Proposition 2.

These distortions to the degree of conservativeness can be avoided by applying specific white noise shocks to both α_p and α_y . Using the unbiased stochastic specification in Proposition 4, $\alpha_p = (1 - \theta\xi) \bar{\alpha}_p$ and $\alpha_y = \left(1 + \frac{1}{\theta}\xi\right) \bar{\alpha}_y$, where ξ is symmetric white noise. This specification guarantees unbiasedness at the deterministic optimum, but the expected price and output level are generally affected. Only the special case in which $\bar{\alpha}_p = \bar{\alpha}_y$ also gives stochastic neutrality in p and y .

Finally, Sørensen (1991) and Beetsma and Jensen (2003) focus on the case in which $\theta = 1$ and use the specification $\alpha_p = \bar{\alpha}_p + \tilde{\xi}$ and $\alpha_y = 1 - \bar{\alpha}_p - \tilde{\xi}$, where $\tilde{\xi}$ is white noise. This means that $p = (1 - \bar{\alpha}_p - \tilde{\xi}) w$ and $y = -(\bar{\alpha}_p + \tilde{\xi}) w$, so stochastic neutrality prevails.¹⁶ However, this preference specification is typically biased; unbiasedness at the deterministic optimum holds for $\bar{\alpha}_p = \bar{\alpha}_y$ and symmetric white noise $\tilde{\xi}$. Nevertheless, it could be sensible to use a specification that is stochastically neutral in p and y . When stochastic preferences are merely used to capture asymmetric information about the central bank's behavior, it may be desirable to employ a specification that does not directly distort variables of interest such as $E[p]$ and $E[y]$. In addition, stochastic neutrality is convenient when preset nominal wages and rational expectations imply $w = E[p]$, because this is not affected by the degree of preference variability.

This example has shown the pitfalls of applying a white noise shock to either α_p or α_y in a monetary policy game. The problem with this biased specification is that changing the normalization of the relative preference parameter could completely reverse results. Geraats (2002) discusses such spurious effects in the literature on central bank transparency.

4.3 Microfounded Business Cycle Models

Consider the following simple, static Robinson Crusoe economy. The representative agent maximizes the utility function

$$U(C, L) = \frac{\alpha_C}{1 - \rho} C^{1 - \rho} - \frac{\alpha_L}{\nu} L^\nu$$

subject to the budget constraint

$$C = wL \tag{29}$$

where C is consumption, L labor supply, w the real wage, $\alpha_C > 0$, $\alpha_L > 0$, $\rho > 0$ and $\nu \geq 1$. Assume a linear production technology, $Y = AL$, where Y is output and A labor productivity ($A > 0$), and a competitive labor market so that $w = A$. Then (29) also corresponds to equilibrium in the goods market: $C = Y$. Optimal labor supply and consumption equal

$$L = \left(\frac{\alpha_C}{\alpha_L} A^{1 - \rho} \right)^{\frac{1}{\rho + \nu - 1}} \quad \text{and} \quad C = \left(\frac{\alpha_C}{\alpha_L} A^\nu \right)^{\frac{1}{\rho + \nu - 1}}$$

In the literature on dynamic stochastic general equilibrium models, stochastic relative preferences are typically modelled by lognormal shocks to α_C or α_L , so $\alpha_C = \kappa \bar{\alpha}_C$

¹⁶More generally, $\alpha_p = \bar{\alpha}_p + \theta \tilde{\xi}$ and $\alpha_y = \bar{\alpha}_y - \frac{1}{\theta} \tilde{\xi}$ ensures stochastic neutrality in p and y in this monetary policy game.

and $\alpha_L = \bar{\alpha}_L$ (e.g. Hall 1997), or $\alpha_C = \bar{\alpha}_C$ and $\alpha_L = \kappa \bar{\alpha}_L$ (e.g. Obstfeld and Rogoff 2000), where $\ln \kappa \sim N(0, \sigma_\kappa^2)$. For both specifications, $\ln \frac{\alpha_C}{\alpha_L} \sim N\left(\ln \frac{\bar{\alpha}_C}{\bar{\alpha}_L}, \sigma_\kappa^2\right)$ which implies they are stochastically neutral in $\ln L$ and $\ln C$. In fact, the two specifications are observationally equivalent in the sense that both generate the same probability distributions for L and C . However, this equivalence does not extend to welfare effects.

A useful measure for welfare analysis in the presence of preference shocks is the percentage change in deterministic consumption \bar{C} that would bring about an equivalent change in expected utility $\Delta E[U]$.¹⁷ Using this measure, it is straightforward to show that the welfare effect of stochastic preferences is typically different for the two lognormal specifications.¹⁸ Intuitively, although the outcomes L and C have the same (lognormal) distribution, the stochastic interaction of C and L with the preference parameters α_C and α_L is different for the two lognormal specifications.

Furthermore, for optimization under preference uncertainty, in which case the agent decides about L and C before knowing the realization of the preference shocks, the two lognormal specifications can easily generate qualitatively different results. In particular, introducing preference uncertainty with a lognormal shock to α_C leads to an increase in labor supply L , whereas a lognormal shock to α_L gives a decline in L . Not surprisingly, the magnitude of the welfare effect of preference uncertainty is also different for the two lognormal specifications.¹⁹

So, the lognormal specification is problematic. In the case of optimization with known preference shocks, the two lognormal specifications yield the same stochastic outcomes, but different welfare effects. For optimization under preference uncertainty, the outcomes also differ and could even affect qualitative conclusions. In addition, the lognormal specifications are typically biased as indicated by Proposition 5, except for $A = 1$ when they are virtually unbiased at the deterministic optimum. For $A \neq 1$, lognormal preference shocks make the absolute value of the MRS biased towards one.

¹⁷More precisely, it is the percent deviation of \tilde{C} from \bar{C} , where \tilde{C} satisfies $\Delta E[U] = \frac{1}{1-\rho} E[\alpha_C] \left(\tilde{C}^{1-\rho} - \bar{C}^{1-\rho}\right)$.

¹⁸In particular, for $\ln \alpha_C \sim N(0, \sigma_\kappa^2)$, $\tilde{C} = \left(\bar{C}^{1-\rho} + \left(e^{\frac{1}{2}\left(\frac{\rho-1}{\rho+\nu-1}\right)^2 \sigma_\kappa^2} - 1\right) e^{-\frac{1}{2}\sigma_\kappa^2} B\right)^{\frac{1}{1-\rho}}$, whereas for $\ln \alpha_L \sim N(0, \sigma_\kappa^2)$, $\tilde{C} = \left(\bar{C}^{1-\rho} + \left(e^{\frac{1}{2}\left(\frac{\rho-1}{\rho+\nu-1}\right)^2 \sigma_\kappa^2} - 1\right) B\right)^{\frac{1}{1-\rho}}$, where $B \equiv \frac{\nu+\rho-1}{\nu} \left(\frac{\bar{\alpha}_C}{\bar{\alpha}_L} A^\nu\right)^{\frac{1-\rho}{\rho+\nu-1}}$, $\bar{C} = \left(\frac{\bar{\alpha}_C}{\bar{\alpha}_L} A^\nu\right)^{\frac{1}{\rho+\nu-1}}$, and using $E[\kappa^\varepsilon] = e^{\frac{1}{2}\varepsilon^2 \sigma_\kappa^2}$.

¹⁹For $\ln \alpha_C \sim N(0, \sigma_\kappa^2)$, $L = \left(\frac{\bar{\alpha}_C}{\bar{\alpha}_L} A^{1-\rho}\right)^{\frac{1}{\rho+\nu-1}} e^{\frac{1}{2}\frac{1}{\rho+\nu-1}\sigma_\kappa^2}$ and $\tilde{C} = \left(\bar{C}^{1-\rho} + \left(e^{-\frac{1}{2}\frac{\rho-1}{\rho+\nu-1}\sigma_\kappa^2} - e^{-\frac{1}{2}\sigma_\kappa^2}\right) B\right)^{\frac{1}{1-\rho}}$, whereas for $\ln \alpha_L \sim N(0, \sigma_\kappa^2)$, $L = \left(\frac{\bar{\alpha}_C}{\bar{\alpha}_L} A^{1-\rho}\right)^{\frac{1}{\rho+\nu-1}} e^{-\frac{1}{2}\frac{1}{\rho+\nu-1}\sigma_\kappa^2}$ and $\tilde{C} = \left(\bar{C}^{1-\rho} + \left(e^{\frac{1}{2}\frac{\rho-1}{\rho+\nu-1}\sigma_\kappa^2} - 1\right) B\right)^{\frac{1}{1-\rho}}$.

The unbiased specification of Proposition 4 amounts to $\alpha_C = \left(1 - \frac{1}{A}\xi\right) \bar{\alpha}_C$ and $\alpha_L = (1 + A\xi) \bar{\alpha}_L$, where ξ is symmetric white noise with variance σ_ξ^2 . Using a Taylor series expansion of $\ln \frac{\alpha_C}{\alpha_L}$ around $\xi = 0$ (similar to the proof of Proposition 5), it follows that $E\left[\ln \frac{\alpha_C}{\alpha_L}\right] = \ln \frac{\bar{\alpha}_C}{\bar{\alpha}_L}$ for $A = 1$. As a result, the unbiased specification is only stochastically neutral in $\ln L$ and $\ln C$ if $A = 1$.

Microfounded business cycle models often use calibration and it is important to choose a sensible value for σ_ξ^2 . To set σ_ξ^2 such that the variances of the preference shocks κ and ξ are comparable, focus on the distribution of the MRS at the deterministic optimum. For both lognormal specifications, $\ln\left(-MRS_{C,L}(\bar{C}, \bar{L})\right)$ has a normal distribution with variance σ_κ^2 . For the unbiased specification, a first-order Taylor approximation of the unbiased MRS around the deterministic case $\xi = 0$ (as in the proof of Proposition 5) shows that $\ln\left(-MRS_{C,L}(\bar{C}, \bar{L})\right)$ is approximately normal with variance $\left(A + \frac{1}{A}\right)^2 \sigma_\xi^2$ for $\xi \sim N(0, \sigma_\xi^2)$. So, the specifications yield comparable variances when $\sigma_\xi^2 = \frac{A^2}{(A^2+1)^2} \sigma_\kappa^2$.

Regarding σ_κ^2 , most studies simply pick some plausible value; for instance Obstfeld and Rogoff (2002) take $\sigma_\kappa^2 = 2\%$. A notable exception is Hall (1997), who establishes econometrically that preference shocks are quantitatively important for employment fluctuations. Using US data for 1947-1993 and assuming $\bar{\alpha}_C = \bar{\alpha}_L = 1$, $\rho = 1$ and $\nu = 2.7$, Hall (1997, Table 1) finds that at high frequencies, the atemporal effect of preference shocks on the log of hours of work, $\frac{1}{\nu} \ln \kappa$, has a standard deviation of 1.97%, so $\sigma_\kappa = 5.32\%$.

Given the fact that relative preference shocks appear to be empirically significant, a proper justification of the stochastic specification is warranted. The lognormal specifications common in microfounded business cycle models could affect qualitative conclusions and welfare effects, depending on which preference parameter is normalized. The unbiased stochastic specification presented in this paper does not suffer from this problem. It has the feature that it preserves average relative preferences and that the average and median preference shock correspond to the deterministic case. It also means that shocks to relative preferences no longer have a similar effect as lognormal productivity shocks.

5 Conclusion

This paper shows that the specifications for stochastic relative preferences commonly used in the literature are problematic because they distort the first moment of relative preferences instead of merely affecting the second moment. For instance, an additive,

white noise shock to the relative preference parameter α_i reduces the relative preference weight put on x_i , and a multiplicative, lognormal shock makes the absolute value of the MRS biased towards one. In each case, the corresponding change in average relative preferences could induce spurious effects. This can be avoided by using an unbiased specification that generates pure white noise shocks to relative preferences, which is presented in this paper.

Researchers using biased specifications of stochastic relative preferences should be aware of the fact that these are equivalent to a white noise shock plus a deterministic change in relative preferences. Clearly, the latter should be properly motivated because it generally affects (even deterministic) optimal outcomes as well as welfare results. In addition, such biased specifications could lead to erroneous findings when analyzing the effect of stochastic preferences on economic outcomes. In particular, the use of an unbiased specification is critical when performing comparative statics with respect to risk or uncertainty.

It should be noted that an unbiased specification of relative preferences does not imply that the optimal stochastic outcomes simply add white noise to the deterministic optimum, although such stochastic neutrality could be analytically convenient. Instead, the interaction between preferences and constraints is typically nonlinear so that white noise relative preference shocks tend to affect optimal outcomes on average, but this is a genuine effect due to the stochastic nature of preferences.

The relevance of the results is illustrated by three examples based on consumer choice, monetary policy and microfounded business cycle models. They show how the biased stochastic specifications for relative preferences that are prevalent in the literature could generate misleading conclusions. These pitfalls are easily avoided by properly modeling stochastic relative preferences.

A Appendix

This appendix contains the proofs of Propositions 1, 2 and 5, and an alternative proof to Proposition 4.

Proof of Proposition 1:

The proof proceeds in two steps: (i) a geometric argument is used to show that (6) and (7) hold for all states of nature $s \in \mathcal{S}$, for every $\mathbf{x} \in \mathcal{X}$ and every $i, j \in \{1, \dots, K\}$, $i \neq j$; (ii) trigonometric properties subsequently imply (4) and (5).

(i) Image two agents, A and B, observing the indifference sets associated with $U(\mathbf{x}; s)$ in hyperspace. Both analyze the relative preference between goods x_i and x_j at point \mathbf{x}_0 for state s , focusing on the same subspace spanned by $x_i \times x_j \subset \mathcal{X}$, but viewed from different perspectives. Agent A observes it in (x_i, x_j) space and finds that the tangency plane to the indifference contour satisfies $-\frac{dx_j}{dx_i} = \tan \gamma_{i,j}(\mathbf{x}_0; s)$. Agent B observes the tangency plane in (x_j, x_i) space and finds that $-\frac{dx_i}{dx_j} = \tan \gamma_{j,i}(\mathbf{x}_0; s)$. Obviously, $\gamma_{i,j}$ and $\gamma_{j,i}$ are closely related. In particular, $\gamma_{i,j}(\mathbf{x}_0; s) + \gamma_{j,i}(\mathbf{x}_0; s) = \frac{1}{2}\pi$ and $\tan \gamma_{j,i}(\mathbf{x}_0; s) = \cot \gamma_{i,j}(\mathbf{x}_0; s)$. Using the fact that $-\frac{dx_j}{dx_i} = MRS_{i,j}(\mathbf{x}_0; s)$ and $-\frac{dx_i}{dx_j} = MRS_{j,i}(\mathbf{x}_0; s)$ gives

$$\begin{aligned} MRS_{i,j}(\mathbf{x}_0; s) &= \tan \gamma_{i,j}(\mathbf{x}_0; s) \\ MRS_{j,i}(\mathbf{x}_0; s) &= \cot \gamma_{i,j}(\mathbf{x}_0; s) \end{aligned}$$

which holds for any $\mathbf{x}_0 \in \mathcal{X}$ and $s \in \mathcal{S}$.

(ii) Now, substituting (3) gives $MRS_{i,j}(\mathbf{x}; s) = \tan(\bar{\gamma}_{i,j}(\mathbf{x}) + \zeta_{i,j}(s; \mathbf{x}))$. Using trigonometric identities it is straightforward to show that²⁰

$$\begin{aligned} \tan(\bar{\gamma}_{i,j}(\mathbf{x}) + \zeta_{i,j}(s; \mathbf{x})) &= \frac{\tan \bar{\gamma}_{i,j}(\mathbf{x}) + \tan \zeta_{i,j}(s; \mathbf{x})}{1 - \tan \bar{\gamma}_{i,j}(\mathbf{x}) \tan \zeta_{i,j}(s; \mathbf{x})} \\ &= \frac{\tan \bar{\gamma}_{i,j}(\mathbf{x}) + \eta_{i,j}(s; \mathbf{x})}{\cot \bar{\gamma}_{i,j}(\mathbf{x}) - \eta_{i,j}(s; \mathbf{x})} \cot \bar{\gamma}_{i,j}(\mathbf{x}) \end{aligned}$$

where $\eta_{i,j}(s; \mathbf{x}) = \tan \zeta_{i,j}(s; \mathbf{x})$. This yields (4). Similarly, $MRS_{j,i}(\mathbf{x}; s) = \cot \gamma_{i,j}(\mathbf{x}; s) = \tan \gamma_{j,i}(\mathbf{x}; s)$, which gives (5). Finally, $\bar{\gamma}_{j,i}(\mathbf{x}) + \zeta_{j,i}(s; \mathbf{x}) = \gamma_{j,i}(\mathbf{x}; s) = \frac{1}{2}\pi - \gamma_{i,j}(\mathbf{x}; s) = \frac{1}{2}\pi - \bar{\gamma}_{i,j}(\mathbf{x}) - \zeta_{i,j}(s; \mathbf{x})$ so that $\zeta_{j,i}(s; \mathbf{x}) = -\zeta_{i,j}(s; \mathbf{x})$ and $\eta_{j,i}(s; \mathbf{x}) = \tan \zeta_{j,i}(s; \mathbf{x}) = -\tan \zeta_{i,j}(s; \mathbf{x}) = -\eta_{i,j}(s; \mathbf{x})$. ■

²⁰In particular, use the fact that $\tan a = \frac{\sin a}{\cos a}$, $\sin(a+b) = \sin a \cos b + \cos a \sin b$, $\cos(a+b) = \cos a \cos b - \sin a \sin b$ and $\cot a = 1/\tan a$.

Proof of Proposition 2:

The proof that the specification $U(\mathbf{x}; s)$ is biased proceeds in two steps: (i) Proposition 1 is used to compute the relative preference shock $\zeta_{i,j}(s; \mathbf{x})$ implied by $U(\mathbf{x}; s)$; (ii) it is shown that $E[\zeta_{i,j}|\mathbf{x}] \neq 0$.

(i) The MRS between good x_i and x_j for $j \in \{1, \dots, K\}$ with $j \neq i$, equals

$$MRS_{i,j}(\mathbf{x}; s) = \frac{\alpha_i(s) u'_i(x_i)}{\bar{\alpha}_j u'_j(x_j)} \quad (30)$$

Using (4) and (30) to solve for $\eta_{i,j}(s; \mathbf{x})$ gives

$$\eta_{i,j}(s; \mathbf{x}) = \frac{\alpha_i(s) u'_i(x_i) - \bar{\alpha}_j u'_j(x_j) \tan \bar{\gamma}_{i,j}(\mathbf{x})}{\alpha_i(s) u'_i(x_i) \tan \bar{\gamma}_{i,j}(\mathbf{x}) + \bar{\alpha}_j u'_j(x_j)}$$

Substituting (10) and rearranging produces

$$\eta_{i,j}(s; \mathbf{x}) = \frac{\xi(s) u'_i(x_i) \bar{\alpha}_j u'_j(x_j)}{\bar{\alpha}_i^2 (u'_i(x_i))^2 + \xi(s) \bar{\alpha}_i (u'_i(x_i))^2 + \bar{\alpha}_j^2 (u'_j(x_j))^2} \quad (31)$$

using $\xi(s) = \alpha_i(s) - \bar{\alpha}_i$. Then, the implied relative preference shock equals $\zeta_{i,j}(s; \mathbf{x}) = \arctan \eta_{i,j}(s; \mathbf{x})$.

(ii) Substituting (31) and differentiating gives after some simplification²¹

$$\frac{d\zeta_{i,j}(s; \mathbf{x})}{d\xi(s)} = \frac{u'_i(x_i) \bar{\alpha}_j u'_j(x_j)}{(\bar{\alpha}_i + \xi(s))^2 (u'_i(x_i))^2 + \bar{\alpha}_j^2 (u'_j(x_j))^2} \quad (32)$$

$$\frac{d^2\zeta_{i,j}(s; \mathbf{x})}{d(\xi(s))^2} = -\frac{2(\bar{\alpha}_i + \xi(s)) (u'_i(x_i))^3 \bar{\alpha}_j u'_j(x_j)}{\left((\bar{\alpha}_i + \xi(s))^2 (u'_i(x_i))^2 + \bar{\alpha}_j^2 (u'_j(x_j))^2 \right)^2} \quad (33)$$

Note that any deterministic bliss points $\mathbf{x} = \mathbf{b}$ with $u'_j(b_j) = 0$ or $u'_i(b_i) = 0$ can be ignored because they yield $\zeta_{i,j}(s; \mathbf{x}) = 0$ for all $s \in \mathcal{S}$ so that relative preferences are not stochastic. Using $\text{sgn}(\alpha_i(s)) = \text{sgn}(\bar{\alpha}_i) \neq 0$ and $\bar{\alpha}_j \neq 0$ it follows that $\frac{d^2\zeta_{i,j}(s; \mathbf{x})}{d(\xi(s))^2} \neq 0$. So, $\zeta_{i,j}(s; \mathbf{x})$ is either convex or concave in $\xi(s)$. Hence, using $E[\xi] = \xi(s_0) = 0$ and Jensen's inequality, $E[\zeta_{i,j}|\mathbf{x}] \neq \zeta_{i,j}(s_0; \mathbf{x}) = 0$. Therefore, $U(\mathbf{x}; s)$ is a biased specification of stochastic relative preferences between x_i and x_j for all $\mathbf{x} \in \mathcal{X}$. ■

²¹Use the fact that $d \arctan x/dx = 1/(1+x^2)$.

Alternative Proof of Proposition 4:

This proof is similar to the one for Proposition 2. Using (1) and (4) to solve for $\eta_{i,j}(s; \mathbf{x})$ gives

$$\eta_{i,j}(s; \mathbf{x}) = \frac{\alpha_i(s) u'_i(x_i) - \alpha_j(s) u'_j(x_j) \tan \bar{\gamma}_{i,j}(\mathbf{x})}{\alpha_i(s) u'_i(x_i) \tan \bar{\gamma}_{i,j}(\mathbf{x}) + \alpha_j(s) u'_j(x_j)}$$

Substituting (10), (13) and (14) gives

$$\eta_{i,j}(s; \mathbf{x}) = \frac{(\lambda^2 + 1) \xi(s) \bar{\alpha}_i u'_i(x_i) \bar{\alpha}_j u'_j(x_j)}{\lambda (1 + \lambda \xi(s)) \bar{\alpha}_i^2 (u'_i(x_i))^2 + (\lambda - \xi(s)) \bar{\alpha}_j^2 (u'_j(x_j))^2} \quad (34)$$

Substituting (1), (13) and (14) into $\lambda = MRS_{i,j}(\mathbf{x}; s)$ gives after rearranging

$$\bar{\alpha}_j u'_j(x_j) = \frac{1 + \lambda \xi(s)}{\lambda - \xi(s)} \bar{\alpha}_i u'_i(x_i)$$

Substitute this into (34) and simplify to get

$$\eta_{i,j}(s; \mathbf{x}) = \frac{(\lambda^2 + 1) \xi(s)}{\lambda (\lambda - \xi(s)) + (1 + \lambda \xi(s))} = \xi(s)$$

It follows from $f_\xi(\xi) = f_\xi(-\xi)$ that $f_\eta(\eta_{i,j}|\mathbf{x}) = f_\eta(-\eta_{i,j}|\mathbf{x})$, so $f_\zeta(\zeta_{i,j}|\mathbf{x}) = f_\zeta(-\zeta_{i,j}|\mathbf{x})$. Note that $\xi = 0$ corresponds to $\eta_{i,j} = \zeta_{i,j} = 0$ and amounts to the deterministic case. As a result, the specification of stochastic relative preferences between x_i and x_j is unbiased at $\bar{\mathbf{x}}$. Finally, substituting (14) and (15) into (1) yields $MRS_{j,k}(\mathbf{x}; s) = \bar{\alpha}_j u'_j(x_j) / \bar{\alpha}_k u'_k(x_k)$, so that indeed the relative preferences between x_j and x_k are not affected. ■

Proof of Proposition 5:

For ease of notation, assume $MRS_{i,j}(\bar{\mathbf{x}}; s) > 0$. The proof for the general case is cumbersome to write out but exactly similar.

For the multiplicative specification in the Proposition,

$$\ln MRS_{i,j}(\bar{\mathbf{x}}; s) = \ln \frac{\bar{\alpha}_i u'_i(\bar{x}_i)}{\bar{\alpha}_j u'_j(\bar{x}_j)} + \ln \kappa$$

For the unbiased MRS specification in (11), a Taylor series expansion around the deterministic case $\eta_{i,j} = 0$ yields²²

$$\begin{aligned} \ln MRS_{i,j}(\bar{\mathbf{x}}; s) &= \ln \frac{\bar{\alpha}_i u'_i(\bar{x}_i)}{\bar{\alpha}_j u'_j(\bar{x}_j)} + \ln \left(1 + \frac{1}{\bar{\lambda}} \eta\right) - \ln (1 - \bar{\lambda} \eta) \\ &= \ln \bar{\lambda} + \left(\bar{\lambda} + \frac{1}{\bar{\lambda}}\right) \eta + \frac{1}{2} \left(\bar{\lambda}^2 - \frac{1}{\bar{\lambda}^2}\right) \eta^2 + \frac{1}{3} \left(\bar{\lambda}^3 + \frac{1}{\bar{\lambda}^3}\right) \eta^3 + \frac{1}{4} \left(\bar{\lambda}^4 - \frac{1}{\bar{\lambda}^4}\right) \eta^4 + \dots \\ &= \ln \bar{\lambda} + \sum_{j=0}^{\infty} \frac{1}{1+2j} \left(\bar{\lambda}^{1+2j} + \frac{1}{\bar{\lambda}^{1+2j}}\right) \eta^{1+2j} + \sum_{j=1}^{\infty} \frac{1}{2j} \left(\bar{\lambda}^{2j} - \frac{1}{\bar{\lambda}^{2j}}\right) \eta^{2j} \end{aligned}$$

where $\bar{\lambda} \equiv \frac{\bar{\alpha}_i u'_i(\bar{x}_i)}{\bar{\alpha}_j u'_j(\bar{x}_j)}$ and $\eta = \eta_{i,j}(s; \bar{\mathbf{x}})$. For $|MRS_{i,j}(\bar{\mathbf{x}})| = |\bar{\lambda}| = 1$, the second summation term drops out. So, the two specifications are observationally equivalent for $\ln \kappa = \sum_{j=0}^{\infty} \frac{1}{1+2j} \left(\bar{\lambda}^{1+2j} + \frac{1}{\bar{\lambda}^{1+2j}}\right) \eta^{1+2j}$. As a result, the specification in the Proposition with this κ is unbiased at $\bar{\mathbf{x}}$. Note that $f_{\eta}(\eta|\mathbf{x}) = f_{\eta}(-\eta|\mathbf{x})$ implies $E[\eta^{1+2j}|\mathbf{x}] = 0$, so $E[\ln \kappa] = 0$. For $|MRS_{i,j}(\bar{\mathbf{x}})| \neq 1$, an adjustment is required so that $E[\ln \kappa] = \sum_{j=1}^{\infty} \frac{1}{2j} \left(\bar{\lambda}^{2j} - \frac{1}{\bar{\lambda}^{2j}}\right) E[\eta^{2j}]$, which means that $\text{sgn } E[\ln \kappa] = \text{sgn}(|\bar{\lambda}| - 1)$. ■

²²Note that the unbiased specification (17) cannot be used to construct the MRS for this Taylor series expansion because evaluating it at $\bar{\mathbf{x}}$ for any state different from s_0 would violate the condition that $MRS_{i,j}(\mathbf{x}(s); s) = \lambda$, so that this specification is not valid at $\bar{\mathbf{x}}$ for $\eta_{i,j} \neq 0$.

References

- Beetsma, R. M. and Jensen, H. (2003), ‘Comment on “Why money talks and wealth whispers: Monetary uncertainty and mystique’, *Journal of Money, Credit and Banking* **35**(1), 129–136.
- Fishburn, P. C. (1998), Stochastic utility, in S. Barberà, P. J. Hammond and C. Seidl, eds, ‘Handbook of Utility Theory, Volume 1: Principles’, Kluwer Academic Publishers, chapter 7, pp. 273–319.
- Geraats, P. M. (2002), ‘Central bank transparency’, *Economic Journal* **112**, F532–F565.
- Hall, R. E. (1997), ‘Macroeconomic fluctuations and the allocation of time’, *Journal of Labor Economics* **15**(1), S223–S250.
- Hildenbrand, W. (1971), ‘Random preferences and equilibrium analysis’, *Journal of Economic Theory* **3**(4), 414–429.
- Obstfeld, M. and Rogoff, K. (2000), ‘New directions for stochastic open economy models’, *Journal of International Economics* **50**(1), 117–153.
- Obstfeld, M. and Rogoff, K. (2002), ‘Global implications of self-oriented national monetary rules’, *Quarterly Journal of Economics* **117**(2), 503–535.
- Sørensen, J. R. (1991), ‘Political uncertainty and macroeconomic performance’, *Economics Letters* **37**(4), 377–381.