

Risk, Ambiguity, and State-Preference Theory

Robert Nau
Fuqua School of Business, Duke University *†‡

August 30, 2010

Abstract

The state-preference framework for modeling choice under uncertainty, in which objects of choice are allocations of wealth or commodities across states of the world, is a natural one for modeling “smooth” ambiguity-averse preferences. It does not require reference to objective probabilities, personalistic consequences, or counterfactual acts, and it allows for state-dependence of utility and unobservable background risk. The decision maker’s local beliefs are encoded in her risk neutral probabilities (her relative marginal rates of substitution between states) and her local risk preferences are encoded in the matrix of derivatives of the risk neutral probabilities. This matrix plays a central but generally unappreciated role in the modeling of risk attitudes in the state-preference framework. It can be computed by inverting a bordered Slutsky matrix and vice versa, it generalizes the Arrow-Pratt measure for approximating local risk premia, and its structure reveals whether the decision maker’s risk preferences are ambiguity-averse as well as risk averse. Two versions of the smooth-ambiguity model are analyzed—the source-dependent risk aversion model and the second-order uncertainty (KMM) model—and it is shown that in both cases the overall premium for local uncertainty can be decomposed as the sum of a risk premium and an ambiguity premium.

Published in *Economic Theory* (2011) 49: 1, 437-46
© Springer-Verlag 2011

*JEL classification D81. Key words: Ellberg’s paradox; risk aversion; ambiguity aversion; state-preference theory; state-dependent utility.

†Corresponding author’s address: Fuqua School of Business, Duke University, Durham, NC, 27708-0120, USA. robert.nau@duke.edu

‡I am very grateful to Mark Machina for many helpful discussions and also to Edi Karni for comments on an earlier draft. All errors and omissions and polemics are entirely my own.

1 Introduction

The urn experiments devised by Ellsberg (1961) expose a subtle restriction imposed by the axioms of subjective expected utility, namely that the additivity property they confer on utility functions requires the decision maker to exhibit the same degree of risk aversion toward every kind of risk, well-defined or ambiguous. Over the last few decades, a reappraisal of the normative as well as descriptive validity of this requirement has motivated the development of an array of generalized models that permit the decision maker to display aversion to ambiguity as well as risk, or, equivalently, to exhibit differing degrees of aversion to different sources of risk. These models fall into three categories in terms of the representations through which they model decisions under ambiguity: smooth preferences, kinked preferences, or incomplete preferences.

The incomplete-preference approach, which relaxes the completeness axiom rather than the independence axiom of SEU theory, is historically the oldest. It originated in the fields of statistics and philosophy (Koopman 1940, Smith 1961, Levi 1970) in the context of Bayesian inference with imprecise prior probabilities, which has become its own subfield since the publication of Walley's (1989) book. Even de Finetti's "Fundamental Theorem of Probability" (1974) is stated in terms of lower and upper bounds on subjective probabilities. A theory of expected utility without the completeness axiom was first proposed by Aumann (1962), and relaxation of the completeness axiom in the context of subjective expected utility, leading to joint imprecision in probabilities and utilities, has been explored by Seidenfeld et al. (1990, 1995), Nau (2006a), Ok et al. (2008), and Galaabaatar and Karni (2010). A model of preferences that are incomplete due to indeterminacy of probabilities was introduced into microeconomics by Bewley (1986), who referred to it as "Knightian" uncertainty¹. Its implications for phenomena such as general equilibrium have been elaborated by Rigotti and Shannon (2005) and others. A decision maker with incomplete preferences does not have fixed indifference curves, and she may exhibit indecision or inertia when called upon to make a choice under ambiguous conditions. This model provides a natural explanation for incompleteness of markets—indeed, an agent with incomplete preferences is a microcosm of an incomplete market.

The kinked-preference approach originated in the work of Schmeidler (1989) and Gilboa and Schmeidler (1989), who proposed models in which the decision maker's beliefs are represented by a Choquet capacity or else by a set of priors accompanied by a decision rule (maximization of minimum expected utility) that completes the preference relation despite the imprecision of beliefs. The preferences of such a decision maker do have an indifference curve representation, which has the characteristic property that the curves have kinks at points where the same outcome is received in two or more states, at which she exhibits first-order risk aversion rather than second-order risk aversion toward some sources

¹The term is based on a distinction between risk and uncertainty proposed by Knight (1921), but that was prior to the development of subjective probability theory and its integration with expected utility theory in microeconomics and finance, so it is not entirely clear which yet-to-be-developed theory he was criticizing.

of uncertainty. Many extensions and applications of such models have been explored in recent decades; for example, a recursive form of the multiple priors model was axiomatized by Epstein and Schneider (2003) and has been applied to intertemporal choice in asset markets.

The smooth-preference approach has been developed more recently (Nau 2001a & 2006b, Klibanoff et al. 2005, Chew and Sagi 2008, Ergin and Gul 2009, Seo 2009), and it explains ambiguity aversion as a form of second-order risk aversion in which indifference curves bend more sharply in some directions than others—namely, directions that correspond to bets on more ambiguous events—in a way that is incompatible with the independence axiom. It will be discussed in much more depth later in this paper.

Most of the models in the three general categories mentioned above adhere to the spirit of the SEU model insofar as they still try to separate beliefs from tastes. That is, they typically have one set of parameters for beliefs about events and another set of parameters for evaluations of outcomes. A recent review of all of them is given by Etner et al. (2010).

An alternative framework for modeling choice under uncertainty descends from a literature that predates the theory of subjective probability, namely 19th-Century consumer theory. This is the *state-preference* framework of Arrow and Debreu (Arrow 1953, Debreu 1959), in which the consumer's objects of choice are bundles of commodities (money and goods and services) that may be state-contingent and also time-contingent. When preferences among such objects satisfy the usual ordering axioms—completeness, transitivity and continuity—they are representable by ordinal utility functions. When an independence condition is also imposed, namely that preferences between two state-contingent commodity bundles are independent of elements they have in common, the utility function has an additive structure and is cardinally measurable and can be given a subjective expected utility interpretation.

The state-preference paradigm is less general than that of SEU theory insofar as it requires the outcomes to have some sort of interpersonal observability and transferability, unlike consequences that may be private experiences of the decision maker. But by the same token it is more concrete: it does not require the contemplation of an idiosyncratic small world within which outcomes and events are defined in a uniquely personal way. The objects of choice have parameters whose meaning is understood in the same way by everyone, and preferences among them can be expressed in a way that makes them common knowledge, namely, through public offers to make purchases or sales at quoted prices. This model also does not depend sensitively on the independence axiom (Savage's P2) nor on the additional axioms of SEU theory that are conventionally used to separate probabilities from utilities and to force preferences to be state-independent (Savage's P3 and P4). From the perspective of consumer theory, additive utility functions are not realistic because they rule out complements and substitutes, and even if additivity is assumed, little additional modeling power is gained by requiring their terms to be uniquely decomposable into products of probabilities for states and utilities for consumption in those states. Rather, all that is gained is a labeling of the components of the state-

preference model that aligns with the labeling of the components of a subjective expected utility model. Insofar as state-preference theory requires little modification when the independence axiom is dropped, it provides a tractable setting in which to model ambiguity aversion and other violations of independence, and it leads to a simple model of smooth preferences that can be used to explain Ellsberg's paradox.²

State-preference theory is closer in spirit to asset pricing theory than to subjective expected utility theory in that its natural parameters of belief are not probabilities that are measures of belief alone. Rather, they are *risk neutral probabilities*, which are the decision maker's marginal betting rates on events (a.k.a. state prices). The risk neutral probability of an event E is the price π_E at which the decision maker would indifferently buy or sell small numbers of lottery tickets (a.k.a. Arrow securities) that yield a payoff of \$1 if E occurs and \$0 otherwise. If the decision maker is truly risk neutral with state-independent utility for money, her risk neutral probabilities can be directly interpreted as subjective probabilities, which is the definition proposed by de Finetti (1937, 1974). But if she has subjective expected utility preferences and is risk averse, her risk neutral probabilities must be interpreted instead as products of subjective probabilities and relative marginal utilities for money, where the latter may depend on unobservable background risk such as pre-existing financial and psychological stakes in events as well as on intrinsic state-dependence of utility (Karni et al. 1983, Karni 1985, Schervish et al. 1990, Nau 2001b). It is therefore generally impossible to uniquely separate probabilities from utilities based on observations of the decision maker's preferences among bets (or any other concrete acts, for that matter), and if she has non-expected-utility preferences—in particular, preferences that exhibit aversion to ambiguity—then her beliefs may not be representable by subjective probabilities at all, even in principle. Yet even under such conditions, beliefs can still be usefully characterized by risk neutral probabilities that vary with observable changes in wealth, as will be discussed here.

Another advantage of working within the state-preference framework is that it allows preferences to be modeled in such a way that the most fundamental principle of rationality is that of *no-arbitrage*, i.e., avoiding a sure loss. This is an axiom that is hard to challenge on either normative or descriptive grounds, and it applies to groups as well as to individuals, so it strongly unifies the theory of rational choice across the domains of personal decisions, games of strategy, and competitive markets. Indeed, the term "risk neutral probabilities" originated in the finance literature, where it refers to a probability distribution on states of the world with respect to which the price of every asset is equal to its expected payoff and whose existence is necessary and sufficient for the prices to be arbitrage-free.³ As such, it can be regarded as the subjective probability distribution of a "risk neutral representative investor" who prices all the assets according to their expected payoffs, but it is not the subjective probability dis-

²Also the Allais paradox (Nau 2001a).

³This result, which follows from a separating-hyperplane argument, is the first fundamental theorem of asset pricing and also the fundamental theorem of subjective probability.

tribution of the "average" investor. The average investor is risk averse and holds positive amounts of primary risky securities such as stocks and bonds. If she has subjective-expected-utility preferences, then her personal risk neutral distribution, which determines the marginal prices she is willing to pay for assets, is the product of her true subjective probability distribution and her state-dependent relative marginal utility for money. Her marginal utility for money is typically greater in low-price states, where she is poorer, so her personal risk neutral distribution is left-shifted relative to her true subjective probability distribution. In equilibrium, the personal risk neutral distributions of all investors must agree, because universal agreement on asset prices is necessary for no-arbitrage. Their true subjective probability distributions are generally heterogenous and unobservable, although on average they must believe that riskier assets have greater expected returns. However, the investors need not have subjective-expected-utility preferences at all—there is no market principle that compels it—and if they are ambiguity-averse as well as risk-averse, they do not.

This paper extends earlier work by Nau (2001a, 2003, 2006b) on the representation of smooth ambiguity-averse preferences by source-dependent utility functions within the state-preference framework. Along the way it explores the connection between the modeling of risk and ambiguity aversion in terms of risk-neutral probabilities and the modeling of income and substitution effects in consumer theory via the Slutsky matrix. A behavioral test for local ambiguity aversion that is a state-dependent generalization of Ellsberg's 2-urn experiment is applied to examples involving correlated background risk. The alternative version of the smooth-ambiguity model developed by Klibanoff et al. (2005) is also discussed—translated into the state-preference setting—and it is shown that the two versions give rise to the essentially the same formula for computing local premia for risk and ambiguity, generalizing Pratt's (1964) formula, although they respond somewhat differently to the presence of background risk.

2 Risk neutral probabilities and their derivatives

Let $\mathcal{X} \subseteq \mathbb{R}^n$ be a set of *acts* that are assignments of money to states of the world. An act is an observable share of the decision maker's wealth (e.g., a portfolio of investments), measured with respect to an arbitrary, possibly-state-dependent status quo. (She might have other assets that are undeclared or non-monetary.) Assume that the decision maker has preferences among acts that are complete, transitive, strictly monotonic, and sufficiently smooth to be represented by a twice-differentiable, strictly-increasing, utility function U . The decision maker is [strictly] *risk averse*, according to Yaari's (1969) generalized definition, if her preferences are [strictly] *payoff-convex*, which means that if $\mathbf{x} \succsim \mathbf{z}$ and $\mathbf{y} \succsim \mathbf{z}$ [and $\mathbf{x} \neq \mathbf{y}$], then $\alpha\mathbf{x} + (1 - \alpha)\mathbf{y} \succsim [\succ] \mathbf{z}$ for $\alpha \in (0, 1)$. In this case her utility function is [strictly] quasi-concave, i.e., the level set consisting of all \mathbf{x} such that $U(\mathbf{x}) \succsim U(\mathbf{y})$ is a [strictly] convex set for every \mathbf{y} . This definition of

risk aversion applies locally (in-the-small) as well as globally (in-the-large), and the decision maker is locally "more risk averse," in a sense to be made precise later, if her local preferences are "more convex," as measured by the curvature of her indifference curves in the vicinity of the status quo. Throughout the remainder of this paper, it will be assumed for simplicity that the decision maker is strictly risk averse. The question is whether she is also ambiguity averse in the sense of being more risk averse toward more ambiguous sources of risk. The assumptions given so far do not rule out the possibility that the decision maker has non-expected-utility preferences that exhibit this property.

Stronger assumptions can be imposed on preferences to restrict the form of the utility function. Let $\mathbf{x}_{-i}a$ denote the act that yields the amount a in state i and agrees with \mathbf{x} in all other states. The decision maker's preferences satisfy the axiom of *coordinate independence*⁴ if $\mathbf{x}_{-i}a \succsim \mathbf{y}_{-i}a \iff \mathbf{x}_{-i}b \succsim \mathbf{y}_{-i}b$ for all $\mathbf{x}, \mathbf{y}, a$, and b . As is well known, when the coordinate independence axiom is added to the ones just mentioned, and $n \geq 3$,⁵ the utility function has an additive representation:

$$U(\mathbf{x}) = \sum_{i=1}^n u_i(x_i), \quad (1)$$

where the evaluation functions $\{u_i\}$ are unique up to positive affine transformations with a common scale factor (Wakker 1989). This can be viewed as a model of state-dependent expected utility without separation of probabilities and utilities, because it is possible to think of $u_i(x_i)$ as being the product of a subjective probability p_i and a state-dependent utility function $v_i(x_i)$, but insofar as v_i could have a state-dependent scale factor, this decomposition is not unique, so the probabilities are indeterminate.

If the decision maker additionally satisfies an axiom of *tradeoff-consistency*, which requires that if $\mathbf{x}_{-i}a \succsim \mathbf{y}_{-i}b$ and $\mathbf{x}_{-i}c \succsim \mathbf{y}_{-i}d$ and $\mathbf{z}_{-i}a \succsim \mathbf{w}_{-i}b$, then $\mathbf{z}_{-i}c \succsim \mathbf{w}_{-i}d$, it follows that u_i can be decomposed as $u_i(x_i) = p_i v(x_i)$ where $v(x)$ is a state-independent utility function, unique up to a positive affine transformation, and \mathbf{p} is a unique probability distribution. \mathbf{p} behaves just like a measure of pure belief under these assumptions, so it is conventionally interpreted as one, and a similar convention is followed in models of choice under uncertainty that use the Savage or Anscombe-Aumann frameworks rather than the state-preference framework. However, even under these conditions, it is generally impossible to be sure that \mathbf{p} correctly measures the decision maker's beliefs, because it still could be distorted by state-dependent utility scale factors.

With or without the stronger assumptions, the first-order properties of the decision maker's local preferences in the vicinity of act \mathbf{x} are represented by the probability distribution $\boldsymbol{\pi}(\mathbf{x})$ whose i^{th} element is the relative marginal utility of money in state i , which is simply the gradient of U at \mathbf{x} normalized so that

⁴This condition, also known as *preferential independence*, is a weaker form of Savage's *sure-thing principle*.

⁵The $n = 2$ case can be covered by strengthening coordinate independence to *generalized triple cancellation* or by imposing an additional requirement known as the *hexagon condition*.

its elements sum to 1, i.e.,

$$\pi_i(\mathbf{x}) := \frac{\partial U(\mathbf{x})/\partial x_i}{\sum_{j=1}^n \partial U(\mathbf{x})/\partial x_j}. \quad (2)$$

The gradient of U at \mathbf{x} will be denoted here by $DU(\mathbf{x})$, in terms of which $\boldsymbol{\pi}(\mathbf{x}) = DU(\mathbf{x})/(\mathbf{1} \cdot DU(\mathbf{x}))$ where $\mathbf{1}$ denotes a vector of 1's. By the strict monotonicity assumption, $\boldsymbol{\pi}(\mathbf{x}) > 0$ for all \mathbf{x} .

$\boldsymbol{\pi}(\mathbf{x})$ is called the decision maker's *risk neutral probability distribution* at \mathbf{x} , because it has the property that she is willing to accept small bets whose expected value is positive with respect to it, as though she is risk neutral with subjective probability distribution $\boldsymbol{\pi}(\mathbf{x})$. If \mathbf{f} is a vector that satisfies $\boldsymbol{\pi}(\mathbf{x}) \cdot \mathbf{f} > 0$, then there exists some $\epsilon > 0$ such that $U(\mathbf{x} + \delta \mathbf{f}) > U(\mathbf{x})$ for all δ between 0 and ϵ , hence she is willing to accept any sufficiently-small bet whose payoff vector is proportional to \mathbf{f} , and on the margin she will accept a bet whose risk neutral expected value is zero.

In the special case where the decision maker has additive utility (1), her risk neutral probability for state i is:

$$\pi_i(\mathbf{x}) = \frac{u'_i(x_i)}{\sum_{j=1}^n u'_j(x_j)}. \quad (3)$$

If the additive utility function is interpreted as the product of a subjective probability distribution and a state-dependent utility function ($u_i(x_i) = p_i v_i(x_i)$), then $\pi_i(\mathbf{x})$ is expressible as

$$\pi_i(\mathbf{x}) = \frac{p_i v'_i(x_i)}{\sum_{j=1}^n p_j v'_j(x_j)}, \quad (4)$$

i.e., it is the renormalized product of the decision maker's subjective probability of state i and her marginal utility for money in state i , where the latter may be intrinsically state-dependent as well dependent on wealth in state i . In the even-more-special case where the decision maker has state-independent expected utility preferences ($v_i(x) \equiv v(x)$), this simplifies to

$$\pi_i(\mathbf{x}) = \frac{p_i v'(x_i)}{\sum_{j=1}^n p_j v'(x_j)}. \quad (5)$$

Second-order properties of the decision maker's local preferences are represented by the matrix of derivatives of $\boldsymbol{\pi}(\mathbf{x})$, which will be denoted as $D\boldsymbol{\pi}(\mathbf{x})$:

$$[D\boldsymbol{\pi}(\mathbf{x})]_{ij} := \partial \pi_i(\mathbf{x})/\partial x_j. \quad (6)$$

Its elements are measured in units of 1/money, e.g., 1/\$, and they are uniquely determined by preferences for a given choice of currency.⁶ This matrix is

⁶A subtlety arises here when the states of the world refer to different possible exchange rates between two currencies. In this case, changing the currency in terms of which risk

an object of fundamental importance in state-preference theory, yet it has not been much discussed in the literature. By definition it determines the marginal change in the decision maker’s risk neutral probabilities that results from a change in state-dependent wealth: $D\boldsymbol{\pi}(\mathbf{x})d\mathbf{x} = d\boldsymbol{\pi}$. It also determines the decision maker’s local degree of risk aversion with respect to small bets, and its structure determines whether the decision maker’s risk preferences are source-dependent, as will be shown later. $D\boldsymbol{\pi}(\mathbf{x})$ can be regarded as a behavioral primitive insofar as it is directly observable—its elements can be estimated from small-but-finite bets the decision maker is willing to accept in the vicinity of the status quo—but it also can be expressed in terms of U as follows:

$$[D\boldsymbol{\pi}(\mathbf{x})]_{ij} = \frac{[D^2U(\mathbf{x})]_{ij} - \pi_i(\mathbf{x}) \sum_{k=1}^n [D^2U(\mathbf{x})]_{kj}}{\sum_{k=1}^n [DU(\mathbf{x})]_k}, \quad (7)$$

where $D^2U(\mathbf{x})$ is the Hessian matrix of U evaluated at \mathbf{x} , i.e., $[D^2U(\mathbf{x})]_{ij} := \partial^2U(\mathbf{x})/\partial x_i\partial x_j$. The Hessian matrix has rank n for a decision maker who is strictly risk averse, and $D\boldsymbol{\pi}$ has rank $n-1$ because its columns sum to zero. The latter is uniquely determined by preferences while the former is not, because U is unique only up to an increasing transformation.

If the decision maker has additive utility, it is possible to express $D\boldsymbol{\pi}(\mathbf{x})$ in terms of a vector-valued Arrow-Pratt⁷ measure of local risk aversion $\mathbf{r}(\mathbf{x})$:

$$r_i(x) := -u_i''(x)/u_i'(x), \quad (8)$$

which is also measured in units of 1/money and uniquely determined by preferences up to the choice of currency. In particular, $D\boldsymbol{\pi}(\mathbf{x})$ satisfies:

$$[D\boldsymbol{\pi}(\mathbf{x})]_{ij} = -\pi_i(\mathbf{x})(1_{i=j} - \pi_j(\mathbf{x}))r_j(x), \quad (9)$$

where $1_{i=j} := 1$ if $i = j$ and $1_{i=j} := 0$ otherwise. If $D\boldsymbol{\pi}(\mathbf{x})$ has this structure for some set of functions $\{r_i(x)\}$, the decision maker’s local indifference curves at \mathbf{x} bend in a way that is consistent with additive utility. In the special cases of state-dependent or state-independent expected utility, the probabilities in (8) can be canceled between the numerator and denominator, yielding $r_i(x) = -v_i''(x)/v_i'(x)$ or $r_i(x) = -v''(x_i)/v'(x_i)$. Notice that even if the decision maker has state-*independent* expected utility preferences, the local Arrow-Pratt measure in state i is generally state-*dependent* due to its dependence on x_i , which is the local background risk.

neutral probabilities are elicited does not merely rescale the elements of $D\boldsymbol{\pi}(\mathbf{x})$ by the same conversion factor, while leaving $\boldsymbol{\pi}(\mathbf{x})$ unaffected. Rather, it changes them both in a state-dependent way. This is not a problem. It does not change the results of any analysis of the decision maker’s preferences among assets, whatever their denomination. In fact, this is one more reason why beliefs ought to be measured in terms of risk neutral probabilities, without any pretense of divorcing them from a financial context. As noted by Schervish et al. (1990), it is impossible to apply de Finetti’s definition of subjective probability in this situation, even under an assumption that the decision maker is risk neutral.

⁷The relevance of the quantity $-u''(x)/u'(x)$ as a measure of local risk aversion was independently discovered by de Finetti (1952), Pratt (1964), and Arrow (1965).

To sum up, in the most general case the properties of local preferences are determined up to second-order effects by the vector $\boldsymbol{\pi}(\mathbf{x})$ and the matrix $D\boldsymbol{\pi}(\mathbf{x})$, and in the various special cases of additive utility they are determined by the vectors $\boldsymbol{\pi}(\mathbf{x})$ and $\mathbf{r}(\mathbf{x})$, both of which vary with \mathbf{x} . All of these objects are uniquely determined by preferences and are directly measurable in terms of acceptable small bets.

$D\boldsymbol{\pi}$ is closely related to the Slutsky matrix of consumer theory. Suppose that there are n commodities (which could be financial assets or other goods), and let $\boldsymbol{\pi}(\mathbf{x})$ denote the consumer's relative marginal rates of substitution at \mathbf{x} , which are risk-neutral probabilities if the commodities are Arrow securities. Then \mathbf{x} is an optimal act for the consumer when $\boldsymbol{\pi}(\mathbf{x})$ coincides with the vector of relative prices, because the first-order condition for an optimum is that relative marginal rates of substitution should equal relative prices. The Slutsky matrix at \mathbf{x} , which will be denoted here as $S(\mathbf{x})$,⁸ determines the marginal change in the consumption vector that is induced by a compensated change in the relative-marginal-price vector. That is, $S(\mathbf{x})$ has the defining property that the optimal act changes from \mathbf{x} to $\mathbf{x} + S(\mathbf{x})d\boldsymbol{\pi}$ when relative prices change from $\boldsymbol{\pi}(\mathbf{x})$ to $\boldsymbol{\pi}(\mathbf{x}) + d\boldsymbol{\pi}$ and the consumer is compensated for the change in prices, or equivalently, if she already owns \mathbf{x} and can reallocate it at the new relative prices. In a financial market, the Slutsky matrix and the matrix of derivatives of the risk neutral probabilities are therefore linked by the following equivalence:

$$S(\mathbf{x})d\boldsymbol{\pi} = d\mathbf{x} \iff D\boldsymbol{\pi}(\mathbf{x})d\mathbf{x} = d\boldsymbol{\pi}. \quad (10)$$

There are several analytic formulas that can be used to compute the Slutsky matrix. The one that is commonly seen nowadays in microeconomics textbooks is stated in terms of the consumer's *Marshallian demand function*, denoted here as $\mathbf{x}^*(\mathbf{p}, W)$, which is the utility-maximizing act that a consumer with disposable wealth W will choose to purchase when the price vector is \mathbf{p} . In other words, $\mathbf{x}^*(\mathbf{p}, W)$ is the solution to $\max_{\mathbf{x} \in X} U(\mathbf{x})$ subject to $\mathbf{p} \cdot \mathbf{x} \leq W$. The first-order condition for an optimum is that $\boldsymbol{\pi}(\mathbf{x}) \propto \mathbf{p}$, i.e., a consumer with strictly convex preferences in possession of \mathbf{x} is at an optimum if and only her relative marginal rates of substitution at \mathbf{x} are proportional to market prices.

If \mathbf{x} is an optimum at prices \mathbf{p} and wealth W , the Slutsky matrix at \mathbf{x} can be determined from the Marshallian demand function as follows:

$$[S(\mathbf{x})]_{ij} := \frac{\partial x_i^*}{\partial p_j} + \frac{\partial x_i^*}{\partial W} x_j^*. \quad (11)$$

The terms on the RHS of this equation are called the price effect and the income effect caused by a compensated change in prices. If the price of commodity j changes by dp_j , ceteris paribus, then the consumer's demand for commodity i changes by $\frac{\partial x_i^*}{\partial p_j} dp_j$. If she is to be compensated for the price change, so that she can afford to retain her original act, she must receive an amount of income equal to $x_j^* dp_j$, which results in a further change $\frac{\partial x_i^*}{\partial W} x_j^* dp_j$ in demand. The

⁸Here "at \mathbf{x} " means "at \mathbf{x} under conditions where it is an optimum."

vector $\mathbf{w}(\mathbf{x})$ that is defined by

$$w_i(\mathbf{x}) := \frac{\partial x_i^*}{\partial W} \quad (12)$$

is the *wealth expansion path* at \mathbf{x} , i.e., the marginal change in the optimal act per unit of change in wealth, holding prices fixed. It is also the direction along which the consumer's relative marginal rates of substitution are unchanged, i.e. it satisfies $D\boldsymbol{\pi}(\mathbf{x}) \cdot \mathbf{w}(\mathbf{x}) = \mathbf{0}$.

A now-obscure fact is that the Slutsky matrix and the wealth expansion vector can also be obtained by inverting the Hessian matrix of the utility function after bordering it with the price vector, which is the derivation used by Slutsky and Hicks:⁹

$$\begin{bmatrix} D^2U(\mathbf{x}) & \mathbf{p} \\ \mathbf{p}^T & 0 \end{bmatrix}^{-1} = \begin{bmatrix} \alpha S(\mathbf{x}) & \mathbf{w}(\mathbf{x}) \\ \mathbf{w}(\mathbf{x})^T & \beta_x \end{bmatrix}. \quad (13)$$

Here, α^{-1} is the marginal utility of income (the ratio of marginal utilities to marginal prices), and β_x is the income derivative of the marginal utility of income (the derivative of the marginal utility of income along the wealth expansion path). Equation (13) can be rewritten as

$$\begin{bmatrix} \alpha D^2U(\mathbf{x}) & \mathbf{p} \\ \mathbf{p}^T & 0 \end{bmatrix}^{-1} = \begin{bmatrix} S(\mathbf{x}) & \mathbf{w}(\mathbf{x}) \\ \mathbf{w}(\mathbf{x})^T & a\beta_x \end{bmatrix}, \quad (14)$$

which cancels the units of utility on both sides. When prices are normalized to sum to 1, $\alpha = (\mathbf{1} \cdot DU(\mathbf{x}))^{-1}$, in which case $\alpha D^2U(\mathbf{x})$ is the same normalized Hessian matrix that is the first term in the formula (7) for $D\boldsymbol{\pi}(\mathbf{x})$.

The Slutsky matrix and the $D\boldsymbol{\pi}$ matrix have inverse properties insofar as $S(\mathbf{x})d\boldsymbol{\pi} = d\mathbf{x}$ if and only if $D\boldsymbol{\pi}(\mathbf{x})d\mathbf{x} = d\boldsymbol{\pi}$, as noted above. However, they are not literally inverses of each other, because both are singular. $S(\mathbf{x})$ is singular because it satisfies $S(\mathbf{x})\boldsymbol{\pi}(\mathbf{x}) = \mathbf{0}$ and $\boldsymbol{\pi}(\mathbf{x}) \cdot S(\mathbf{x}) = \mathbf{0}$, and $D\boldsymbol{\pi}(\mathbf{x})$ is singular because it satisfies $D\boldsymbol{\pi}(\mathbf{x})\mathbf{w}(\mathbf{x}) = \mathbf{0}$ and $\mathbf{1} \cdot D\boldsymbol{\pi}(\mathbf{x}) = \mathbf{0}$. Also, $S(\mathbf{x})$ is symmetric, while $D\boldsymbol{\pi}(\mathbf{x})$ is not. The main result of this section is that the Slutsky matrix and the wealth expansion vector can be obtained by inverting the $D\boldsymbol{\pi}$ matrix after bordering it with the vector of risk neutral probabilities, without any reference to the decision maker's utility function:

Theorem 1:¹⁰ If the decision maker is strictly risk averse, then

$$\begin{bmatrix} D\boldsymbol{\pi}(\mathbf{x}) & \boldsymbol{\pi}(\mathbf{x}) \\ \boldsymbol{\pi}(\mathbf{x})^T & 0 \end{bmatrix}^{-1} = \begin{bmatrix} S(\mathbf{x}) & \mathbf{w}(\mathbf{x}) \\ \mathbf{1}^T & 0 \end{bmatrix}. \quad (15)$$

⁹ This equation is discussed by Stern (1986). The bordered Hessian matrix was used by Pareto as well as Slutsky and Hicks in modeling the utility function of the consumer (Afriat 1980, Doole 1983).

¹⁰ Proofs are given in the appendix.

If the prices are not normalized to sum to 1, the equation is generalized to:

$$\begin{bmatrix} D\boldsymbol{\pi}(\mathbf{x}) & \mathbf{p} \\ \mathbf{p}^T & 0 \end{bmatrix}^{-1} = \begin{bmatrix} S(\mathbf{x}) & \mathbf{w}(\mathbf{x}) \\ ((\mathbf{1} \cdot \mathbf{p})^{-1})\mathbf{1}^T & 0 \end{bmatrix}. \quad (16)$$

The vector $\mathbf{w}(\mathbf{x})$ on the right is now the wealth expansion path with respect to real prices rather than normalized prices. This follows from the fact that if \mathbf{p} is rescaled by some positive constant, the length of \mathbf{w} and the constant in the bottom row of the matrix on the right are rescaled in a reciprocal fashion. The Slutsky matrix is unaffected.

Theorem 1 can be used to rewrite the relation $S(\mathbf{x})d\boldsymbol{\pi} = d\mathbf{x}$ in terms of $D\boldsymbol{\pi}(\mathbf{x})$, yielding the result that the change $d\mathbf{x}$ in the optimal act that follows from an infinitesimal change $d\boldsymbol{\pi}$ in relative marginal prices is determined by:

$$\begin{bmatrix} d\mathbf{x} \\ 0 \end{bmatrix} = \begin{bmatrix} D\boldsymbol{\pi}(\mathbf{x}) & \boldsymbol{\pi}(\mathbf{x}) \\ \boldsymbol{\pi}(\mathbf{x})^T & 0 \end{bmatrix}^{-1} \begin{bmatrix} d\boldsymbol{\pi} \\ 0 \end{bmatrix} \quad (17)$$

However, the same formula does not provide the best first-order approximation to the effect of a *finite* change $\Delta\boldsymbol{\pi}$. This is given instead by:

Corollary 1: The change $\Delta\mathbf{x}$ in the optimal act that results from

a small finite change $\Delta\boldsymbol{\pi}$ in relative marginal prices satisfies:¹¹

$$\begin{bmatrix} \Delta\mathbf{x} \\ 0 \end{bmatrix} \approx \begin{bmatrix} D\boldsymbol{\pi}(\mathbf{x}) & \boldsymbol{\pi}(\mathbf{x}) + \Delta\boldsymbol{\pi} \\ (\boldsymbol{\pi}(\mathbf{x}) + \Delta\boldsymbol{\pi})^T & 0 \end{bmatrix}^{-1} \begin{bmatrix} \Delta\boldsymbol{\pi} \\ 0 \end{bmatrix} \quad (18)$$

3 A generalized measure of aversion to risk and ambiguity

Pratt (1964) defined a risk premium as the amount of expected value that a decision maker with state-independent expected-utility preferences would be willing to give up to dispose of all the risk she faces. In the general state-preference setting, there may no such thing as an objective or subjective probability distribution from which to calculate expected values, and there may be no such thing as a riskless position that is feasible to achieve by selling off risky assets. Nevertheless, it is straightforward to generalize the risk premium concept to deal with these issues by re-defining it in terms of the amount the decision maker would need to be paid to take on a new risk, on top of whatever risk she already faces, and by taking expected values with respect to the local risk neutral probability distribution rather than a probability distribution that is a hypothetical measure of pure belief.

Suppose that a decision maker in possession of an act \mathbf{x} contemplates taking on some additional risk in the form of a small bet whose payoff vector is \mathbf{f} . The

¹¹This is a simplification of Proposition 4(b) in Nau (2003).

marginal price of the bet \mathbf{f} in the vicinity of act \mathbf{x} is its risk neutral expected value $\boldsymbol{\pi}(\mathbf{x}) \cdot \mathbf{f}$, because on the margin a decision maker who already possesses \mathbf{x} would be willing to pay $\boldsymbol{\pi}(\mathbf{x}) \cdot \mathbf{f}$ per unit of \mathbf{f} . The bet is *neutral* in the vicinity of act \mathbf{x} if $\boldsymbol{\pi}(\mathbf{x}) \cdot \mathbf{f} = 0$. The *buying price* for \mathbf{f} in its entirety in the vicinity of act \mathbf{x} , denoted by $B(\mathbf{f}; \mathbf{x})$, is the maximum amount the decision maker would be willing to pay for \mathbf{f} , i.e., it satisfies

$$\mathbf{x} + \mathbf{f} - B(\mathbf{f}; \mathbf{x}) \sim \mathbf{x}, \quad (19)$$

or equivalently

$$U(\mathbf{x} + \mathbf{f} - B(\mathbf{f}; \mathbf{x})) = U(\mathbf{x}). \quad (20)$$

The *overall uncertainty premium*¹² for \mathbf{f} at \mathbf{x} , denoted by $b(\mathbf{f}; \mathbf{x})$, is the difference between the risk neutral expected value of \mathbf{f} and its buying price:

$$b(\mathbf{f}; \mathbf{x}) := \boldsymbol{\pi}(\mathbf{x}) \cdot \mathbf{f} - B(\mathbf{f}; \mathbf{x}) \quad (21)$$

If the decision maker is strictly risk averse, this quantity is strictly positive for all $\mathbf{f} \neq \mathbf{0}$. The overall uncertainty premium is a generalization of Pratt's concept of a risk premium, and it can be locally approximated by a formula that is a multivariate generalization of Pratt's formula.¹³

Theorem 2: The overall uncertainty premium of a neutral bet \mathbf{f} in the vicinity of act \mathbf{x} has the local quadratic approximation:

$$b(\mathbf{f}; \mathbf{x}) \approx \widehat{b}(\mathbf{f}; \mathbf{x}) := -\frac{1}{2} \mathbf{f} \cdot D\boldsymbol{\pi}(\mathbf{x}) \mathbf{f} \quad (22)$$

$$= -\frac{1}{2} \frac{\mathbf{f} \cdot D^2U(\mathbf{x}) \mathbf{f}}{\mathbf{1} \cdot DU(\mathbf{x})} \quad (23)$$

$$= \frac{1}{2} \mathbf{f} \cdot \boldsymbol{\Pi}(\mathbf{x}) \mathbf{R}(\mathbf{x}) \mathbf{f} \quad (24)$$

where $\boldsymbol{\Pi}(\mathbf{x})$ is the diagonal matrix whose diagonal elements are the coordinates of the vector $\boldsymbol{\pi}(\mathbf{x})$ and $\mathbf{R}(\mathbf{x})$ is a matrix-valued measure of local risk aversion whose ij^{th} element is

$$r_{ij}(\mathbf{x}) := -\frac{\partial^2 U(\mathbf{x}) / \partial x_i \partial x_j}{\partial U(\mathbf{x}) / \partial x_i} \quad (25)$$

¹²This was called the total risk premium in Nau 2003 & 2006a, and it was then decomposed into a pure-risk component and an ambiguity component. Here the pure-risk component will be called the risk premium, and total premium for risk and ambiguity will be called the overall uncertainty premium.

¹³A multivariate generalization of the risk premium formula and a corresponding matrix-valued measure of local risk aversion also arise naturally in the setting where the decision maker's utility function is a von Neumann-Morgenstern utility function over vectors of commodities, which has been studied by Duncan (1977) and Karni (1979). There the objects of choice are objective probability distributions over commodity vectors rather than vectors of monetary payoffs across states whose probabilities are subjective or undefined, and the risk premium is vector-valued rather than scalar-valued. In both settings the risk aversion measure takes the form of a matrix of second partial derivatives divided by first partial derivatives of the appropriate utility function, although the applications are otherwise quite different.

Under the models to be discussed in the following two sections, this formula will turn out to be decomposable into the sum of a premium for risk and a premium for ambiguity.

The matrix $\mathbf{R}(\mathbf{x})$ that appears in (24) is a generalization of the Arrow-Pratt risk aversion measure that applies to situations like those considered here in which the decision maker has smooth preferences that may or may not conform to subjective expected utility theory. It is not always uniquely defined by preferences, because $U(\mathbf{x})$ is unique only up to increasing transformations if the coordinate independence axiom is not satisfied, and these transformations also affect $D^2U(\mathbf{x})$ and $\mathbf{R}(\mathbf{x})$, although their effects ultimately cancel out in the evaluation of (23) and (24). The fundamental equation in Theorem 2 is (22), because it refers only to quantities that are uniquely determined and observable.

In the special case where $U(\mathbf{x})$ is additive, which is the case of expected-utility preferences with possibly-state-dependent utility and possibly-significant background risk, the formula for the overall uncertainty premium simplifies to¹⁴

$$\widehat{b}(\mathbf{f}; \mathbf{x}) = \frac{1}{2} \sum_i \pi_i(\mathbf{x}) r_i(x_i) f_i^2, \quad (26)$$

where $r_i(x) = -u_i''(x)/u_i'(x)$ is a state-dependent local Arrow-Pratt risk aversion measure at \mathbf{x} in state i , which is uniquely determined by preferences. In the *very* special case where the decision maker has state-independent expected utility preferences and zero background risk, this can be further simplified to Pratt's formula

$$\widehat{b}(\mathbf{f}; \mathbf{x}) = \frac{1}{2} r(x) \sum_i p_i f_i^2, \quad (27)$$

where p_i is interpretable as a subjective probability of state i (which is uniquely determined under these conditions) and $r(x)$ is the Arrow-Pratt risk aversion measure at riskless wealth x . Notice, however, that in the passage from state-independent preferences and zero background risk to more general settings, *two* things happen. First, the state-independent Arrow-Pratt measure $r(x)$ gets replaced by a state-dependent measure $r(x_i)$ or $r_i(x_i)$ or $r_{ij}(\mathbf{x})$. Second, and less obviously, the decision maker's true subjective probability p_i gets replaced by the local risk-neutral probability $\pi_i(\mathbf{x})$. What this shows is that the correct probability with which to multiply the Arrow-Pratt measure is the risk neutral probability, not the true probability. The true probability only appears in the very special case where it happens to coincide with the risk neutral probability.

4 Ambiguity aversion as source-dependent risk aversion

The behavior typically observed in experiments such as Ellsberg's is most commonly explained by models in which the decision maker's first-order beliefs are

¹⁴This is equation (2.8) in Karni (1985).

to some extent indeterminate, represented by a set of subjective probability distributions (as in models of incomplete preferences or multiple priors) or by a second-order probability distribution (as in the KMM model). However, there is another, and simpler, explanation, namely that the decision maker has determinate subjective probabilities, but she is merely more risk averse toward some bets than toward others, perhaps due to differences in the quality of information on which her beliefs are based or due to differences in her tastes for different kinds of risks. For example, in Ellsberg’s two-urn problem, she might regard red and black as equally likely to be drawn from the urn with unknown composition, on the grounds that her ignorance is symmetric between the two colors. Nevertheless, she might regard a bet on that urn as more risky, and therefore less preferred, than a bet on the urn with known composition. Such preferences are incompatible with the independence axiom of subjective expected utility theory but are otherwise rational insofar as they do not expose the decision maker to any form of exploitation. In fact, they may be quite rational in protecting her from exploitation by others with better information.

In the simplest case, there are two distinct sources of risk represented by two logically independent partitions of events, $\mathcal{A} := \{A_1, \dots, A_J\}$ and $\mathcal{B} := \{B_1, \dots, B_K\}$, and the state space is $\mathcal{A} \times \mathcal{B}$. Events that are measurable with respect to \mathcal{A} and \mathcal{B} will be called \mathcal{A} -events and \mathcal{B} -events, respectively. Suppose that the decision maker’s utility function has the compound form:

$$U(\mathbf{x}) = \sum_j u_j \left(\sum_k v_{jk}(x_{jk}) \right), \quad (28)$$

where j and k refer to A_j and B_k .¹⁵ This functional form can be obtained by a corresponding compound application of the independence axiom (Nau 2006b), in which the following two conditions hold for all $\mathbf{x}, \mathbf{x}^*, \mathbf{y}$, and \mathbf{y}^* . First, $\mathbf{x}_A \mathbf{y}_{-A} \succsim \mathbf{x}_A^* \mathbf{y}_{-A}$ if and only if $\mathbf{x}_A \mathbf{y}_{-A}^* \succsim \mathbf{x}_A^* \mathbf{y}_{-A}^*$ for every \mathcal{A} -event A (“ \mathcal{A} -independence”). Second, for every j , $\mathbf{x}_B \mathbf{y}_{-B} \succsim_j \mathbf{x}_B^* \mathbf{y}_{-B}$ if and only if $\mathbf{x}_B \mathbf{y}_{-B}^* \succsim_j \mathbf{x}_B^* \mathbf{y}_{-B}^*$ for every \mathcal{B} -event B , where \succsim_j denotes conditional preference given event A_j (“conditional \mathcal{B} -independence”).

In general the evaluation functions $\{u_j\}$ and $\{v_{jk}\}$ can be state-dependent in arbitrary ways and cannot be factored into products of probabilities for events and utilities for outcomes. However, if assumptions of tradeoff consistency are imposed, the functional form can be simplified to

$$U(\mathbf{x}) = \sum_j p_j u \left(\sum_k q_{k|j} v(x_{jk}) \right), \quad (29)$$

in which p_j might be interpreted as a subjective probability for A_j and $q_{k|j}$ might be interpreted as a subjective conditional probability for B_k given A_j , and u and v might be interpreted as second-order and first-order Bernoulli

¹⁵In principle there could be more than two qualitatively different sources of risk, in which case a more deeply nested utility function of this form could be used.

utility functions, respectively. (This is the same functional form as Ergin and Gul's model.) A decision maker with this utility function will choose among acts whose payoffs depend only on j as if her risk attitude is represented by $u \circ v$, and she will choose among acts whose payoffs depend only on k as if her risk attitude is represented by v alone. It will be assumed henceforth that the u 's and v 's in (28) and (29) are increasing and concave, in which case $u \circ v$ is more concave than v alone.

Even if the decision maker's utility function can be expressed in the form (29), the possibility of state-dependent utility factors can never be ruled out, so $\{p_j\}$ and $\{q_{k|j}\}$ are not necessarily measures of pure belief. Also, the decision maker's status quo investments may not be observable, in which case her preferences among small bets may depend in unknown ways on the background risk she faces. These issues raise complications for definitions of aversion to risk and ambiguity that depend on using riskless positions as reference points. However, they pose no difficulty for this model: it is possible to determine whether the decision maker is systematically more risk averse toward bets on \mathcal{A} -events than toward bets on \mathcal{B} -events through appropriately calibrated bets, as shown later in this section.

In the state-dependent form of this model (28), the decision maker's local risk neutral distribution satisfies

$$\pi_{jk}(\mathbf{x}) \propto u'_j v'_{jk}, \quad (30)$$

where u'_j is shorthand for $u'_j(\sum_k v_{jk}(x_{jk}))$ and v'_{jk} is shorthand for $v'_{jk}(x_{jk})$, both of which are positive. Let u''_j and v''_{jk} similarly stand for $u''_j(\sum_k v_{jk}(x_{jk}))$ and $v''_{jk}(x_{jk})$, and also let $v'_{j\bullet} := \sum_k v'_{jk}(x_{jk})$. Next, define vectors $\mathbf{s}(\mathbf{x})$ and $\mathbf{t}(\mathbf{x})$ as follows (henceforth suppressing the dependence of all terms on \mathbf{x}):

$$s_j := -(u''_j/u'_j)v'_{j\bullet} \quad (31)$$

$$t_{jk} := -v''_{jk}/v'_{jk} \quad (32)$$

t_{jk} is a first-order state-dependent Arrow-Pratt measure of risk aversion that applies to the actual payoff in event $A_j B_k$, while s_j is a second-order Arrow-Pratt measure that applies to the conditional expected payoff in event A_j , in a manner to be explained below. The multiplicative factor $v'_{j\bullet}$ in the expression for s_j serves to convert its units to 1/money, the same as the units of t_{jk} .

Now let $\pi_j := \sum_k \pi_{jk}$ denote the marginal risk neutral probability of event A_j , let $\pi_{k|j} := \pi_{jk}/\pi_j$ be the conditional risk neutral probability of B_k given A_j , and let $\bar{\mathbf{f}}$ denote the vector whose j^{th} element is the conditional risk neutral expected value of \mathbf{f} given event A_j :

$$\bar{f}_j := \sum_k \pi_{k|j} f_{jk}, \quad (33)$$

which measures risk that is specific to A -events. By applying Theorem 2 of the previous section to this utility function, the following result is obtained:

Theorem 3: Under the two-source utility model, the overall uncertainty premium of a neutral bet \mathbf{f} has the local quadratic approximation:

$$\widehat{b}(\mathbf{f}; \mathbf{x}) = \frac{1}{2} \sum_j \pi_j s_j \bar{f}_j^2 + \frac{1}{2} \sum_{jk} \pi_{jk} t_{jk} f_{jk}^2. \quad (34)$$

The summation on the right in this expression would be the approximate risk premium for a decision maker with state-dependent expected utility preferences represented by $\{v_{jk}\}$. The summation on the left is an additional premium that is required by a decision maker with non-expected-utility preferences who, for whatever reason, especially dislikes her expected payoffs to be dependent on \mathcal{A} -events. It will be henceforth called the approximate local *ambiguity premium* associated with the bet.

The extent to which the decision maker is locally more risk-averse toward \mathcal{A} -events can be measured by a simple test that is a state-dependent generalization of Ellsberg's 2-urn experiment, without eliciting $D\pi(\mathbf{x})$ in all its detail. Let A denote some \mathcal{A} -event and let B denote some \mathcal{B} -event, and let \bar{A} and \bar{B} denote their complements, so the four possible joint realizations of the two events are expressible as AB , $A\bar{B}$, $\bar{A}B$, and $\bar{A}\bar{B}$. Let their risk neutral probabilities be denoted π_{AB} , $\pi_{A\bar{B}}$, etc., and assume they are all non-zero. Then an *A:B-bet of size Δ* , denoted $\mathbf{f}_{A:B}(\Delta)$, is defined as a bet in which the payoff in each joint event is equal to Δ divided by the corresponding risk neutral probability, with a positive sign attached if A occurs and a negative sign if \bar{A} occurs. (Δ itself may be positive or negative.) A *B:A-bet of size Δ* , denoted $\mathbf{f}_{B:A}(\Delta)$, has payoffs of the same magnitudes as $\mathbf{f}_{A:B}(\Delta)$ in every state but with signs that depend on whether B occurs rather than whether A occurs. The payoffs of the two kinds of bets look like this in contingency-table form:

$$\begin{array}{c} \begin{array}{cc} & \begin{array}{c} B \\ \bar{B} \end{array} \\ \begin{array}{c} A \\ \bar{A} \end{array} & \begin{array}{|c|c|} \hline +\Delta/\pi_{AB} & +\Delta/\pi_{A\bar{B}} \\ \hline -\Delta/\pi_{\bar{A}B} & -\Delta/\pi_{\bar{A}\bar{B}} \\ \hline \end{array} \\ & \mathbf{f}_{A:B}(\Delta) \end{array} \qquad \begin{array}{cc} & \begin{array}{c} B \\ \bar{B} \end{array} \\ \begin{array}{c} A \\ \bar{A} \end{array} & \begin{array}{|c|c|} \hline +\Delta/\pi_{AB} & -\Delta/\pi_{A\bar{B}} \\ \hline +\Delta/\pi_{\bar{A}B} & -\Delta/\pi_{\bar{A}\bar{B}} \\ \hline \end{array} \\ & \mathbf{f}_{B:A}(\Delta) \end{array} \end{array}.$$

If the risk neutral probabilities are all equal by virtue of symmetries of the decision maker's beliefs and prior investments, then the payoffs are all equal in magnitude, in which case the *A:B-bet* and the corresponding *B:A-bet* are precisely the choices in Ellsberg's 2-urn problem, merely recentered so that they are all neutral bets in which losses as well as gains are possible. In general the risk neutral probabilities of the four events will vary, but because the payoff magnitudes are scaled in inverse proportion to them, the bets remain neutral and symmetry is preserved among the four cells in the sense that they each contribute $\pm\Delta$ to the risk neutral expected value. By construction, the bets not only have identical risk neutral expected values (namely zero), but they also have identical risk premia if the second-order utility function u_j is linear for all

j (i.e., if $U(\mathbf{x})$ is additive), because in that case $u_j'' = 0$, so the first term in the formula for the overall uncertainty premium drops out, and the result depends only on the squares of the payoffs that appear in the second term. However, if some or all of the second-order utility functions in (28) are strictly concave, then the first term is strictly larger for $\mathbf{f}_{A:B}(\Delta)$ than for $\mathbf{f}_{B:A}(\Delta)$, as formalized in:

Corollary 2: Under the two-source utility model, the quadratic approximation of the local ambiguity premium of an $A:B$ -bet of size Δ is strictly positive while that of the corresponding $B:A$ -bet is zero, hence $\widehat{b}(\mathbf{f}_{A:B}(\Delta); \mathbf{x}) > \widehat{b}(\mathbf{f}_{B:A}(\Delta); \mathbf{x})$ for all \mathbf{x} , if u_j is strictly concave for all j .

This inequality holds for both positive and negative Δ : changing the sign of Δ is merely equivalent to swapping A and B for their complements. As an example, consider a decision problem with state space $\{A_1, A_2, A_3\} \times \{B_1, B_2, B_3\}$, and suppose that the decision maker's beliefs and state-dependent marginal utilities combine to yield the following local risk neutral probabilities in the 9 states:

	B_1	B_2	B_3
A_1	.05	.20	.05
A_2	.20	.10	.05
A_3	.10	.05	.20

Now consider a pair of $A:B$ and $B:A$ bets in which $A = A_1 \cup A_2$ and $B = B_1$. Collecting terms, the risk neutral probabilities for these events are

	B_1	$B_2 \cup B_3$
$A_1 \cup A_2$.25	.40
A_3	.10	.25

and the payoff tables for the bets (in which the the payoffs are inversely proportional to the risk neutral probabilities) are

	B_1	$B_2 \cup B_3$
$A_1 \cup A_2$	$+4\Delta$	$+2.5\Delta$
A_3	-10Δ	-4Δ

 $\mathbf{f}_{A:B}(\Delta)$

	B_1	$B_2 \cup B_3$
$A_1 \cup A_2$	$+4\Delta$	-2.5Δ
A_3	$+10\Delta$	-4Δ

 $\mathbf{f}_{B:A}(\Delta)$

A decision maker who is uniformly more risk averse toward bets on \mathcal{A} -events than toward bets on \mathcal{B} -events, consistent with a utility function of the form (28), will assign a higher overall uncertainty premium to the first of the two bets and will therefore strictly prefer the second one for any small Δ , positive or negative. The same would be true for any other 2×2 partition of the two

sets of events, for example, $A = A_2$ and $B = B_1 \cup B_2$, which yields the following pair of bets, of which the second would again be strictly preferred:

$$\begin{array}{c}
 \begin{array}{cc}
 & \begin{array}{cc} B_1 \cup B_2 & B_3 \end{array} \\
 \begin{array}{c} A_2 \\ A_1 \cup A_3 \end{array} & \begin{array}{|cc|}
 \hline
 +3.33\Delta & +20\Delta \\
 -2.5\Delta & -4\Delta \\
 \hline
 \end{array} \\
 \mathbf{f}_{A:B}(\Delta)
 \end{array}
 \qquad
 \begin{array}{c}
 \begin{array}{cc}
 & \begin{array}{cc} B_1 \cup B_2 & B_3 \end{array} \\
 \begin{array}{c} A_2 \\ A_1 \cup A_3 \end{array} & \begin{array}{|cc|}
 \hline
 +3.33\Delta & -20\Delta \\
 +2.5\Delta & -4\Delta \\
 \hline
 \end{array} \\
 \mathbf{f}_{B:A}(\Delta)
 \end{array}
 \end{array}
 \end{array}$$

A caveat here is that a systematic preference for $B:A$ -bets over $A:B$ -bets may be apparent only for small values of Δ if the situation is one in which the decision maker is relatively much more risk averse than ambiguity averse (i.e., when the ambiguity premium term in (34) is comparatively small) and risk neutral probabilities have extreme values. For example, if event $A\bar{B}$ has a very low risk neutral probability, then the payoffs of $\mathbf{f}_{A:B}(\Delta)$ and $\mathbf{f}_{B:A}(\Delta)$ in event $A\bar{B}$ will be very large in magnitude with opposite signs, as they are in the last pair of payoff tables ($\pm 20\Delta$). A decision maker who is highly risk averse will wish to avoid a large negative payoff and therefore might prefer $\mathbf{f}_{A:B}(\Delta)$ if Δ is positive and large enough to invalidate the quadratic approximation, obscuring the fact that she may be modestly ambiguity-averse as well.

5 Ambiguity aversion as aversion to second-order uncertainty

The functional forms (28) and (29) can be given another interpretation, in which the decision maker has second-order uncertainty about her first-order probabilities and/or utilities, and she is averse to second-order uncertainty as well as to first-order risk. In this interpretation, the \mathcal{B} -events are first-order events that are observable, while the \mathcal{A} -events are second-order events that are credal states of the decision maker, each characterized by its own first-order subjective probability distribution over the first-order events and its own state-dependent and perhaps even credal-state-dependent Bernoulli utility function representing first-order aversion to risk. First-order acts and second-order acts are defined as acts whose payoffs depend only on first-order events or second-order events, respectively. Second-order acts are only imaginary, but they can be envisioned as tools of thought that the decision maker might use to construct her first-order preferences. To facilitate comparison with the two-source model, let the set of second-order states be henceforth denoted by \mathcal{C} and indexed by i , and let the first-order states have a two-way partition $\mathcal{A} \times \mathcal{B}$ with elements indexed by jk . The general state-dependent version of the utility function is then

$$U(\mathbf{x}) = \sum_i u_i \left(\sum_{jk} v_{ijk}(x_{jk}) \right), \tag{35}$$

and the particular case that is obtained when the first- and second-order utility functions are entirely state-independent is

$$U(\mathbf{x}) = \sum_i \mu_i u \left(\sum_{jk} p_{jk|i} v(x_{jk}) \right), \quad (36)$$

where μ_i is the probability of credal state i and $p_{jk|i}$ is the probability of first-order event jk in credal state i . This is the discrete form of the Klibanoff et al. (KMM) model, which they axiomatized in a Savage-type framework¹⁶ with the set of consequences being an interval of real numbers (like amounts of money). Insofar as it has the same functional form as the two-source model discussed in the previous section, it could also be axiomatized in the state-preference framework, although (as in the the KMM model) this would require the explicit articulation of preferences over second-order acts, and here it would also require a corresponding re-interpretation of the compound application of the independence axiom. Putting those issues aside, the results of the previous section concerning risk premia will be recast here, and the risk neutral probabilities and conditional expectations will be indexed by i as well as j and k as if bets on the credal state were possible, but the status quo wealth vector \mathbf{x} and the bet \mathbf{f} will be indexed only by j and k so as to reflect only first-order payoffs, i.e., acts that are concrete.

The local risk neutral distribution determined by the state-independent version of the second-order uncertainty model is a joint distribution over the first-order and second-order states, and it satisfies

$$\pi_{ijk}(\mathbf{x}) \propto (\mu_i u'_i)(p_{jk|i} v'_{jk}), \quad (37)$$

where u'_i is shorthand for $u' \left(\sum_{jk} p_{jk|i} v(x_{jk}) \right)$ and v'_{jk} is shorthand for $v'(x_{jk})$, again suppressing the dependence of all terms on \mathbf{x} . Let u''_i and v''_{jk} similarly stand for $u'' \left(\sum_{jk} p_{jk|i} v_{jk}(x_{jk}) \right)$ and $v''(x_{jk})$, and also let $v'_{i\bullet} := \sum_{jk} p_{jk|i} v'(x_{jk})$ be the expected marginal utility of money in credal state i .¹⁷ Next, define vectors $\mathbf{s}(\mathbf{x})$ and $\mathbf{t}(\mathbf{x})$ exactly as before:

$$s_i := - (u''_i / u'_i) v'_{i\bullet} \quad (38)$$

$$t_{jk} := - v''_{jk} / v'_{jk} \quad (39)$$

Again, t_{jk} is interpretable as a first-order state-dependent Arrow-Pratt measure for actual payoffs, but now s_i is a second-order state-dependent Arrow-Pratt measure for expected payoffs in credal states. t_{jk} is not also subscripted by i (unlike the analogous dependence in the two-source model) because neither v nor \mathbf{x} depends on the credal state when only first-order acts are considered.

¹⁶Some questions about the axiomatization of second-order beliefs in this framework have been raised by Epstein (2010), with a response from Klibanoff et al. (2010).

¹⁷The first-order marginal utility of money, v'_{jk} , is independent of the credal state, but its expected value varies with i insofar as the probabilities of first-order events depend on i .

Now let $\pi_i := \sum_{jk} \pi_{ijk}$ denote the risk neutral probability of credal state C_i , let $\pi_{jk|i} := \pi_{ijk}/\pi_i$ denote the conditional risk neutral probability of $A_j B_k$ given C_i , and let $\pi_{jk} := \sum_i \pi_{ijk}$ denote the unconditional risk neutral probability of $A_j B_k$. The conditional risk neutral expected payoff of \mathbf{f} in credal state i is

$$\bar{f}_i := \sum_{jk} \pi_{jk|i} f_{jk}, \quad (40)$$

in terms of which we have the following analog of Theorem 3:

Theorem 4: Under the second-order uncertainty model, the overall uncertainty premium of a neutral bet \mathbf{f} has the local quadratic approximation:

$$\hat{b}(\mathbf{f}; \mathbf{x}) = \frac{1}{2} \sum_i \pi_i s_i \bar{f}_i^2 + \frac{1}{2} \sum_{jk} \pi_{jk} t_{jk} f_{jk}^2. \quad (41)$$

In this expression, which looks almost exactly like the one in the previous theorem, the summation on the right is the quadratic approximation of the risk premium that would be assessed by a decision maker with state-dependent expected utility preferences whose risk neutral probabilities are $\{\pi_{jk}\}$ and whose state-dependent Arrow-Pratt risk aversion measures are $\{t_{jk}\}$. The summation on the left is the quadratic approximation of an ambiguity premium assessed by a decision maker who is also averse to the uncertainty about her credal state, as represented by a concave second-order utility function u .

An important difference between this model and the previous one is that it has many more free parameters. A constructive application of it would require an explicit enumeration of credal states and an assignment of both first-order and second-order probabilities rather than just the assessment of a second-order utility function to represent greater aversion to risks that are more ambiguous. Even Ellsberg's two-urn problem offers a rich set of possibilities. The decision maker might believe that the ambiguous urn contains either all red balls or all black balls, modeled by a binary distribution. Or she might believe that all proportions are equally likely, modeled by a uniform distribution. She might think there is some chance that the two urns are identical after all. Her second-order beliefs could be a mixture over these scenarios, and others can be imagined.

j and k are symmetric in the approximation formula (41) because no distinction has yet been made between subjectively ambiguous and unambiguous events. The following approach is suggestive. Suppose that \mathcal{B} -events are unambiguous in the sense that they are stochastically independent of the credal state, while \mathcal{A} -events are potentially ambiguous in the sense that they are dependent on the credal state but independent of the \mathcal{B} -events. Then $p_{jk|i}$ in equation (36) has the factorization $p_{jk|i} = p_{j|i} q_k$. This is a model-specific definition of ambiguity that refers to unobservable states of mind, rather than a behavioral definition, but it describes a scenario in which the second-order uncertainty

model should be expected to produce ambiguity-averse behavior. In particular, it might be expected that a decision maker with these beliefs and aversion to credal uncertainty would be more averse to betting on \mathcal{A} -events than \mathcal{B} -events in the sense of assigning higher ambiguity premia to $A:B$ -bets than to $B:A$ -bets, analogous to the property of the two-source utility function that was established in Corollary 2. However, the picture is not quite that simple. The decision maker's preferences among small bets in the vicinity of act \mathbf{x} are determined as usual by her local risk neutral probability distribution, whose value in state ijk now satisfies

$$\pi_{ijk}(\mathbf{x}) \propto (\mu_i u'_i)(p_{j|i} q_k) v'_{jk}, \quad (42)$$

where u'_i and v'_{jk} both depend \mathbf{x} . The corresponding conditional risk neutral probability of first-order state jk in credal state i satisfies

$$\pi_{jk|i}(\mathbf{x}) \propto p_{j|i} q_k v'_{jk}, \quad (43)$$

and the unconditional risk neutral probability of first-order state jk satisfies

$$\begin{aligned} \pi_{jk}(\mathbf{x}) &\propto \sum_i (\mu_i u'_i) (q_k p_{j|i}) v'_{jk} \\ &= q_k v'_{jk} \sum_i \mu_i u'_i p_{j|i} \\ &\propto q_k v'_{jk} \widehat{p}_j, \end{aligned} \quad (44)$$

where $\widehat{p}_j \propto \sum_i \mu_i u'_i p_{j|i}$ is in some sense an average risk neutral probability of A_j over the credal states.

An issue that emerges here is that background risk is important in this model. If the decision-maker's first-order marginal utility of money v'_{jk} is not independent between j and k , which might be the case if she has significant prior stakes that are correlated between the two partitions of events, then $\pi_{jk}(\mathbf{x})$ is not independent between j and k . So, even if the decision maker privately believes the events are independent, the fact may not be revealed by her preferences for bets. This is also true under the two-source utility model and the subjective expected utility model. However, under the second-order uncertainty model, *unlike* the two-source model, the decision maker's local aversion to ambiguity can sometimes get obscured by this effect even for bets that are vanishingly small. It is possible to construct patterns of correlated background risk against which some $B:A$ -bets are actually more sensitive to ambiguity than the corresponding $A:B$ -bets, and hence the $B:A$ -bets have larger ambiguity premia, despite the fact that the decision maker is averse to credal uncertainty and her probabilities for the \mathcal{A} -events vary with her credal state while her probabilities for the \mathcal{B} -events do not.

The following example provides an illustration. Consider a 2-urn problem in which urn 1 contains equal numbers of red and black balls and urn 2 contains balls that may be light blue, dark blue, light green, and dark green, in proportions that are uncertain. Let $\mathcal{B} = \{B_1, B_2\} := \{red, black\}$ and $\mathcal{A} = \{A_1, A_2, A_3, A_4\} := \{light\ blue, dark\ blue, light\ green, dark\ green\}$.

Suppose the decision maker's preferences are represented by a second-order uncertainty model with the following two first-order probability distributions, each of which is independent between \mathcal{A} and \mathcal{B} :

		B_1	B_2			B_1	B_2
		0.5	0.5			0.5	0.5
A_1	.05	.025	.025	A_1	.50	.250	.250
A_2	.40	.200	.200	A_2	.05	.025	.025
A_3	.50	.250	.250	A_3	.05	.025	.025
A_4	.05	.025	.025	A_4	.40	.200	.200
		p_1				p_2	

(Both marginal and joint probabilities are shown.) Let them have equal second-order probabilities ($\mu = (\frac{1}{2}, \frac{1}{2})$) and let the first- and second-order utility functions be $v(x) = 1 - \exp(-2x)$ and $u(v) = 1 - \exp(-10v)$. In other words, the decision maker is risk averse with a first-order Arrow-Pratt risk aversion measure equal to 2 for all x , she is ambiguity averse with a second-order Arrow-Pratt measure equal to 10 for all v ,¹⁸ and she has credal uncertainty about the probabilities of the \mathcal{A} -events but not the \mathcal{B} -events. The former have the probability distributions (0.05, 0.4, 0.5, 0.05) and (0.5, 0.05, 0.05, 0.4) in credal states 1 and 2, while the latter have the probability distribution (0.5, 0.5) in both of the credal states.

To complete the picture, we need to specify the status quo act \mathbf{x} , which is the local background risk. Suppose that it is the following:

	B_1	B_2
A_1	.615	0
A_2	0	1
A_3	0	.615
A_4	1	0
	\mathbf{x}	

The payoffs here are correlated between \mathcal{A} and \mathcal{B} : the decision maker's wealth is high in states $\{A_1B_1, A_2B_2, A_3B_2, A_4B_1\}$ and low in the other four states. What is most significant about this pattern is that it makes the values of bets on B_1 or B_2 dependent on the credal state. In credal state 1 there is a very high first-order probability (0.9) that either A_2 or A_3 will occur, and in both of those events the decision maker will be much wealthier if B_2 occurs than if B_1 occurs. In credal state 2 it is the other way around. Insofar as marginal utilities go down when wealth goes up, this means that the first-order expected utility of a small bet on B_1 or a small bet on B_2 depends strongly on the credal state, despite the fact that the decision maker's first-order probabilities for those events are the same in both credal states.

¹⁸In this parameterization, the first-order Arrow-Pratt measure has units of 1/money, and the second-order measure has units of 1/(first-order-utiles). The unit of money is not a dollar or euro. Rather, it is the unit in terms of which the decision maker's first-order risk tolerance (the reciprocal of the first-order Arrow-Pratt measure) is expressed.

The local risk neutral probability distribution determined by the first- and second-order probabilities, the first- and second-order utility functions, and the background risk is as follows:

	B_1	B_2
A_1	.066	.225
A_2	.184	.025
A_3	.225	.066
A_4	.025	.184
	$\pi(\mathbf{x})$	

There are rather significant variations in the risk neutral probabilities at the level of individual states. For example, the value in state A_3B_1 is nearly 10 times as large as the value in state A_4B_1 , partly due to the fact that A_3B_1 has a larger average first-order probability between the two credal states but mostly due to the fact that the marginal utility of money is much higher in state A_3B_1 , because it is one of the lowest-payoff states.

Now let $A = \{A_1 \cup A_2\}$ and let $B = B_1$, i.e., A is the event "blue" (light or dark) and B is the event "red", and let's proceed to evaluate the risk premia that the decision maker assigns to $A:B$ -bets and $B:A$ -bets. A has a first-order probability of 0.45 ($=0.05+0.4$) in credal state 1 and 0.55 ($=0.5+0.05$) in credal state 2, while the first-order probability of B is the same (0.5) in both states, so we might expect that the decision maker would be more averse to betting on blue than red, but that remains to be seen. Although there is considerable variability in the risk neutral probabilities of the individual states, when we collect terms we find (not coincidentally) that the risk neutral probabilities of the events AB , $A\bar{B}$, $\bar{A}B$, and $\bar{A}\bar{B}$ are all equal to $\frac{1}{4}$. The $A:B$ -bet and $B:A$ -bet of size Δ therefore have the following payoff tables:

	B	\bar{B}		B	\bar{B}
A	+4 Δ	+4 Δ		+4 Δ	-4 Δ
\bar{A}	-4 Δ	-4 Δ		+4 Δ	-4 Δ
	$\mathbf{f}_{A:B}(\Delta)$			$\mathbf{f}_{B:A}(\Delta)$	

These are the standard Ellsberg bets in the 2-urn problem, re-centered so that both losses and gains are possible. So, what we have here is a situation in which the standard Ellsberg bets just happen to be standard $A:B$ -bets and $B:A$ -bets for this specification of A and B and $U(\mathbf{x})$, but there is more going on under the hood, and what is going on under the hood makes a difference. The values of bets on B or \bar{B} have a strong dependence on the credal state, for the reasons mentioned above, while the values of bets on A or \bar{A} do not depend as strongly on the credal state as might otherwise have been expected, because the first-order probability of A is not very different between the credal states and because the average marginal utility of money is the same in A and \bar{A} in both credal states. The punch line is that $\mathbf{f}_{B:A}(\Delta)$ actually has a much larger ambiguity premium than $\mathbf{f}_{A:B}(\Delta)$ for every small Δ , positive or negative,

despite the fact that A is the event whose probability is ambiguous, so betting on blue is preferred to betting on red. For example, if we choose $\Delta = \pm 0.04$, we obtain the following results from the approximation formula of Theorem 4:

	$\mathbf{f}_{A:B}(\Delta)$	$\mathbf{f}_{B:A}(\Delta)$
Risk Premium	.102	.102
+ Ambiguity premium	.014	.162
= Overall uncertainty premium	.116	.264

The ambiguity premium is almost negligible in comparison to the risk premium for $\mathbf{f}_{A:B}(\Delta)$, while it is much larger than the risk premium for $\mathbf{f}_{B:A}(\Delta)$. The same pattern holds as Δ approaches zero from either direction (it merely becomes a more accurate approximation), and it is confirmed by exact solutions.

To guarantee that ambiguity aversion is revealed by local preferences for Ellsberg-type bets under the second-order uncertainty model, it is necessary to impose some constraints on background risk, for example, to require the dependence of the risk neutral probabilities on k to be the same for any fixed i and j (as is also assumed here for the true probabilities), which will force $\bar{f}_i = 0$ when $\mathbf{f} = \mathbf{f}_{B:A}(\Delta)$, thus zeroing out the ambiguity premium exactly as in the source-dependent model. This can be done as follows:

Corollary 3: Under the second-order uncertainty model, the quadratic approximation of the local ambiguity premium of a $B:A$ -bet is zero at \mathbf{x} , and hence $\hat{b}(\mathbf{f}_{A:B}(\Delta); \mathbf{x}) \geq \hat{b}(\mathbf{f}_{B:A}(\Delta); \mathbf{x})$, if the following conditions hold:

- (i) Given the credal state, the first-order probability distributions are independent between \mathcal{A} and \mathcal{B} , and only the probabilities of \mathcal{A} -events depend on the credal state, i.e., $p_{jk|i}$ has the factorization $p_{jk|i} = p_{j|i}q_k$.
- (ii) The first-order marginal utilities at \mathbf{x} are multiplicatively separable between j and k within A and also within \bar{A} , i.e., they can be factored as $v'_{jk}(\mathbf{x}) = y_j z_{k|A}$ for all $j \in A$ and $v'_{jk}(\mathbf{x}) = y_j z_{k|\bar{A}}$ for all $j \in \bar{A}$, for some sets of positive numbers $\{y_j\}$, $\{z_{k|A}\}$, and $\{z_{k|\bar{A}}\}$.¹⁹

Condition (ii) is automatically satisfied for all \mathbf{x} if \mathcal{A} has only two elements, as in the two-urn experiment with only two colors of balls in the ambiguous urn (which is why a 4-color urn was needed in the example). The condition is also satisfied if the first-order payoffs are additively separable between j and k and the first-order utility function is an exponential function (an important special case), and it is satisfied if the local first-order marginal utilities depend only on j or only on k , which is plausible if the sources of risk are conceptually distinct and the decision maker has significant financial exposure to only one of them.

¹⁹Here "for all $j \in A$ " is shorthand for "for all j s.t. $A_j \in A$."

6 Conclusions

State-preference theory provides a convenient starting point for modeling ambiguity-averse preferences insofar as it doesn't give an especially privileged role to the independence axiom in the first place. Its formal structure is that of consumer theory, in which the additivity property of utility that follows from the independence axiom is quite often behaviorally unrealistic. The intuition it brings to the modeling of problems such as Ellsberg's urn is that consequences received in mutually exclusive events may be complementary goods under some conditions, and the decision maker's appetite for some risks may be satiated more quickly than others, which are familiar sentiments when choosing among other types of consumption bundles. It is well-suited to modeling decisions that take place in financial markets, because it gives a very privileged role to money, and the rationality principles that do or do not apply to the individual are largely the same as those that do or do not apply to groups of investors or the market as a whole. The natural parameters of belief in the state-preference framework are the decision maker's risk neutral probabilities, which are her relative marginal betting rates on events. These are the correct probabilities to use in computing local risk premia, even when utility is additive, because they adjust for background risk as well as intrinsic state-dependence of utility. The fact that they do not uniquely separate beliefs from tastes turns out not to be much of a problem despite running counter to conventions. The matrix of derivatives of the local risk neutral probabilities—which bears an almost-inverse relationship to the Slutsky matrix—describes the curvature of the decision maker's indifference curves and determines her local degree of aversion to risk and ambiguity. Models of smooth ambiguity-averse preferences are readily constructed in this framework, and two such models—the source-dependent risk aversion model and the second-order uncertainty (KMM) model—allow a decomposition of the overall uncertainty premium into distinct components for risk and ambiguity. Local aversion to ambiguity can be detected by measurements that are a state-dependent generalization of Ellsberg's two-urn experiment, although ambiguity aversion that is comparatively weak may sometimes be overshadowed by background risk.

References

- [1] S. AFRIAT, *Demand Theory and the Slutsky Matrix* (1980), Princeton University Press
- [2] F. ANSCOMBE, R. AUMANN, A definition of subjective probability, *Ann. Math. Stat.* **34** (1963), 199-205
- [3] K. ARROW, The role of securities in the optimal allocation of risk-bearing, *Quarterly J. Econ.* **31** (1964), 91-96
- [4] K. ARROW, *Aspects of the Theory of Risk Bearing* (1965), Yrjö Jahnssonin, Helsinki

- [5] R. AUMANN, Utility theory without the completeness axiom, *Econometrica* **30** (1962) 445-462
- [6] T. BEWLEY, Knightian Decision Theory Part I, Cowles Foundation Discussion Paper No. 807. (1986) Reprinted in *Decisions in Economics and Finance* **25** (2002), 79-110
- [7] S. CHEW, J. SAGI, Small worlds: modeling attitudes toward sources of uncertainty, *J. Econ. Theory* **139** (2008), 1-24
- [8] B. DE FINETTI, La Prévision: Ses Lois Logiques, Ses Sources Subjectives. *Ann. Inst. Henri Poincaré* **7** (1937), 1-68. Translation reprinted in H.E. Kyburg and H.E. Smokler, Eds., *Studies in Subjective Probability*, 2nd ed. (1980), Robert Krieger, New York, 53-118
- [9] B. DE FINETTI, Sulla Preferibilità, *Giornale degli Economisti e Annali di Economia* **11** (1952), 685-709
- [10] B. DE FINETTI, *Theory of Probability, Vol. 1.* (1974), Wiley, New York
- [11] G. DEBREU, *Theory of Value: An Axiomatic Analysis of Economic Equilibrium* (1959), Yale University Press, New Haven
- [12] P. DOOLEY, Slutsky's equation is Pareto's solution, *History of Political Economy* **15** (1983), 513-517
- [13] G. DUNCAN, A matrix measure of multivariate local risk aversion, *Econometrica* **45** (1977), 895-903
- [14] H. ERGIN, F. GUL, A subjective theory of compound lotteries, *J. Econ. Theory* **144** (2009), 899-929
- [15] L. EPSTEIN, M. SCHNEIDER, Recursive multiple-priors, *J. Econ. Theory* **113** (2003) 1-31
- [16] L. EPSTEIN, A paradox for the 'smooth ambiguity' model of preference, *Econometrica*, forthcoming (2010)
- [17] J. ETNER, M. JELEVA, J. TALLON, Decision theory under ambiguity, *J. Econ. Surveys*, forthcoming (2010)
- [18] T. GALAABAATAR, E. KARNI, Objective and subjective expected utility with incomplete preferences, working paper (2010)
- [19] I. GILBOA, D. SCHMEIDLER, Maxmin expected utility with non-unique prior, *J. Math. Econ.* **18** (1989), 141-153
- [20] E. KARNI, On multivariate risk aversion, *Econometrica* **47** (1979), 1391-1401
- [21] E. KARNI, *Decision Making Under Uncertainty: The Case of State-Dependent Preferences* (1985), Harvard University Press

- [22] E. KARNI, A definition of subjective probability with state-dependent preferences, *Econometrica* **61** (1993), 187-198
- [23] E. KARNI, D. SCHMEIDLER, K. VIND, On state-dependent preferences and subjective probabilities, *Econometrica* **51** (1983) 1021-1031
- [24] P. KLIBANOFF, M. MARINACCI, S. MUKERJI, A smooth model of decision making under ambiguity, *Econometrica* **73** (2005), 1849-1892
- [25] P. KLIBANOFF, M. MARINACCI, S. MUKERJI, On the smooth ambiguity model: a reply, working paper (2010)
- [26] F. KNIGHT, *Uncertainty and profit* (1921), Houghton Mifflin, Boston
- [27] B. KOOPMAN, Axioms and algebra of intuitive probability, *Ann. Math.* **41** (1940), 269-292
- [28] I. LEVI, *The Enterprise of Knowledge* (1980), MIT Press, Cambridge
- [29] R. NAU, Uncertainty aversion with second-order probabilities and utilities, *Electronic Proceedings of the Second International Symposium on Imprecise Probabilities and their Applications*, <http://www.sipta.org/isipta01/proceedings/063.html> (2001a)
- [30] R. NAU, de Finetti was right: Probability does not exist, *Theory and Decision* **51** (2001b), 89-124
- [31] R. NAU, A generalization of Pratt-Arrow measure to non-expected-utility preferences and inseparable probability and utility, *Management Sci.* **49** (2005), 1089-1104
- [32] R. NAU, The shape of incomplete preferences. *Ann. Statist.* **34** (2006a), 2430-2448
- [33] R. NAU, Uncertainty aversion with second-order utilities and probabilities, *Management Sci.* **52** (2006b), 136-145
- [34] E. OK, P. ORTOLEVA, G. RIELLA, Incomplete preferences under uncertainty: indecisiveness in beliefs vs. tastes, working paper (2008)
- [35] J. PRATT, Risk aversion in the small and in the large, *Econometrica* **32** (1964), 122-136
- [36] L. RIGOTTI, C. SHANNON, Uncertainty and risk in financial markets, *Econometrica* **73** (2005), 203-243
- [37] M. SCHERVISH, T. SEIDENFELD, J. KADANE, State-dependent utilities, *J. Amer. Statist. Assoc.* **85** (1990), 840-847
- [38] T. SEIDENFELD, M. SCHERVISH, J. KADANE, Decisions without ordering, in W. Sieg, ed., *Acting and Reflecting* (1990.) 143-170. Reidel, Dordrecht

- [39] T. SEIDENFELD, M. SCHERVISH, J. KADANE, A representation of partially ordered preferences, *Ann. Statist.* **23** (1995), 2168-2217
- [40] K. SEO, Ambiguity and second-order belief, *Econometrica* **77** (2009) 1575-1605
- [41] C. SMITH, Consistency in statistical inference and decision. *J. Roy. Stat. Soc. B* **23** (1961), 1-25
- [42] N. STERN, A note on commodity taxation: the choice of variable and the Slutsky, Hessian, and Antonelli matrices (SHAM), *Rev. Econ. Stud.* **53** (1986), 293-299
- [43] P. WAKKER, *Additive Representations of Preferences: a New Foundation for Decision Analysis* (1989), Kluwer, Dordrecht
- [44] P. WALLEY, *Statistical Reasoning With Imprecise Probabilities* (1989), Chapman and Hall, London
- [45] M. YAARI, Some remarks on measures of risk aversion and their uses, *J. Econ. Theory* **1** (1969), 315-329

Appendix: Proofs of theorems

Theorem 1: For a strictly risk averse decision maker, both $D\boldsymbol{\pi}(\mathbf{x})$ and $S(\mathbf{x})$ have rank $n - 1$, while both bordered matrices have rank $n + 1$, hence they are invertible. To see that the bordered $D\boldsymbol{\pi}$ matrix has rank $n + 1$, note that $\boldsymbol{\pi}(\mathbf{x})$ is linearly independent from the columns of $D\boldsymbol{\pi}(\mathbf{x})$, because its elements do not sum to zero, so their concatenation $[D\boldsymbol{\pi}(\mathbf{x}) \quad \boldsymbol{\pi}(\mathbf{x})]$ has rank n . In turn, $[\boldsymbol{\pi}(\mathbf{x})^T \quad 0]$ is linearly independent from the rows of $[D\boldsymbol{\pi}(\mathbf{x}) \quad \boldsymbol{\pi}(\mathbf{x})]$, because otherwise there would exist a non-zero n -vector \mathbf{f} satisfying $\mathbf{f} [D\boldsymbol{\pi}(\mathbf{x}) \quad \boldsymbol{\pi}(\mathbf{x})] = [\boldsymbol{\pi}(\mathbf{x})^T \quad 0]$, which would imply $\mathbf{f} \cdot \boldsymbol{\pi}(\mathbf{x}) = 0$ and $\mathbf{f} \cdot D\boldsymbol{\pi}(\mathbf{x}) = \boldsymbol{\pi}(\mathbf{x})^T$, which together would also imply $\mathbf{f} \cdot D\boldsymbol{\pi}(\mathbf{x})\mathbf{f} = 0$, which is impossible because $\mathbf{f} \cdot D\boldsymbol{\pi}(\mathbf{x})\mathbf{f} < 0$ for every \mathbf{f} that satisfies $\mathbf{f} \cdot \boldsymbol{\pi}(\mathbf{x}) = 0$ if the decision maker is strictly risk averse. Similarly for the bordered Slutsky matrix, $\mathbf{w}(\mathbf{x})$ is linearly independent from the columns of $S(\mathbf{x})$, because the latter all have a zero inner product with $\boldsymbol{\pi}(\mathbf{x})$ while the former does not, and $[\mathbf{1}^T \quad 0]$ is linearly independent from the rows of $[S(\mathbf{x}) \quad \mathbf{w}(\mathbf{x})]$ because the latter all have a zero inner product with $[\boldsymbol{\pi}(\mathbf{x})^T \quad 0]$ while the former does not.

Next, if $\boldsymbol{\pi}(\mathbf{x})$ is strictly positive, any vector \mathbf{y} can be expressed as

$$\mathbf{y} = \begin{bmatrix} \alpha\boldsymbol{\pi}(\mathbf{x}) + \gamma d\boldsymbol{\pi} \\ \beta \end{bmatrix} \quad (45)$$

for some constants α , β , and γ , where $d\boldsymbol{\pi}$ is a vector satisfying $\mathbf{1} \cdot d\boldsymbol{\pi} = 0$ and $0 < \pi_i(\mathbf{x}) + d\pi_i < 1$ for all i , which can be interpreted as a feasible change in

relative prices. Let $d\mathbf{x} := S(\mathbf{x})d\boldsymbol{\pi}$ be the compensated change in consumption induced by $d\boldsymbol{\pi}$, and note that $S(\mathbf{x})\boldsymbol{\pi}(\mathbf{x}) = \mathbf{0}$ and $\boldsymbol{\pi}(\mathbf{x}) \cdot S(\mathbf{x}) = \mathbf{0}$ (because a change in prices proportional to the current price vector has no effect on consumption, and $S(\mathbf{x})$ is symmetric), and $\boldsymbol{\pi}(\mathbf{x}) \cdot d\mathbf{x} = 0$ ($d\mathbf{x}$ is self-financing at relative prices $\boldsymbol{\pi}(\mathbf{x})$). It follows that

$$\begin{bmatrix} S(\mathbf{x}) & \mathbf{w}(\mathbf{x}) \\ \mathbf{1}^T & 0 \end{bmatrix} \begin{bmatrix} \alpha\boldsymbol{\pi}(\mathbf{x}) + \gamma d\boldsymbol{\pi} \\ \beta \end{bmatrix} = \begin{bmatrix} \gamma d\mathbf{x} + \beta\mathbf{w}(\mathbf{x}) \\ \alpha \end{bmatrix}. \quad (46)$$

Multiplication by the bordered $D\boldsymbol{\pi}$ matrix yields

$$\begin{bmatrix} D\boldsymbol{\pi}(\mathbf{x}) & \boldsymbol{\pi}(\mathbf{x}) \\ \boldsymbol{\pi}(\mathbf{x})^T & 0 \end{bmatrix} \begin{bmatrix} \gamma d\mathbf{x} + \beta\mathbf{w}(\mathbf{x}) \\ \alpha \end{bmatrix} = \begin{bmatrix} \alpha\boldsymbol{\pi}(\mathbf{x}) + \gamma d\boldsymbol{\pi} \\ \beta \end{bmatrix} \quad (47)$$

because $D\boldsymbol{\pi}(\mathbf{x})d\mathbf{x} = d\boldsymbol{\pi}$ (first-order condition for $d\mathbf{x}$ to yield a change of $d\boldsymbol{\pi}$ in relative marginal rates of substitution), and $D\boldsymbol{\pi}(\mathbf{x}) \cdot \mathbf{w}(\mathbf{x}) = \mathbf{0}$ (marginal rates of substitution do not change along the wealth expansion path), and $\boldsymbol{\pi}(\mathbf{x}) \cdot d\mathbf{x} = 0$ (self-financing property of $d\mathbf{x}$). Hence,

$$\begin{bmatrix} D\boldsymbol{\pi}(\mathbf{x}) & \boldsymbol{\pi}(\mathbf{x}) \\ \boldsymbol{\pi}(\mathbf{x})^T & 0 \end{bmatrix} \begin{bmatrix} S(\mathbf{x}) & \mathbf{w}(\mathbf{x}) \\ \mathbf{1}^T & 0 \end{bmatrix} \mathbf{y} = \mathbf{y} \quad \text{for all } \mathbf{y}, \quad (48)$$

which implies

$$\begin{bmatrix} D\boldsymbol{\pi}(\mathbf{x}) & \boldsymbol{\pi}(\mathbf{x}) \\ \boldsymbol{\pi}(\mathbf{x})^T & 0 \end{bmatrix}^{-1} = \begin{bmatrix} S(\mathbf{x}) & \mathbf{w}(\mathbf{x}) \\ \mathbf{1}^T & 0 \end{bmatrix}. \quad \blacksquare \quad (49)$$

Corollary 1: $\Delta\boldsymbol{\pi}$ and $\Delta\mathbf{x}$ must satisfy $D\boldsymbol{\pi}(\mathbf{x})\Delta\mathbf{x} \approx \Delta\boldsymbol{\pi}$ (the change in \mathbf{x} must yield a change in relative marginal utilities that matches the change in relative prices, to a first-order approximation) and $(\boldsymbol{\pi}(\mathbf{x}) + \Delta\boldsymbol{\pi})\Delta\mathbf{x} = 0$ ($\Delta\mathbf{x}$ must be self-financing at the new relative prices). This can be written as the system of equations:

$$\begin{bmatrix} D\boldsymbol{\pi}(\mathbf{x}) & \boldsymbol{\pi}(\mathbf{x}) + \Delta\boldsymbol{\pi} \\ (\boldsymbol{\pi}(\mathbf{x}) + \Delta\boldsymbol{\pi})^T & 0 \end{bmatrix} \begin{bmatrix} \Delta\mathbf{x} \\ 0 \end{bmatrix} \approx \begin{bmatrix} \Delta\boldsymbol{\pi} \\ 0 \end{bmatrix}. \quad (50)$$

The bordered $D\boldsymbol{\pi}$ matrix is non-singular, by the same argument as before, and multiplication by its inverse yields the result. \blacksquare

Theorem 2: The first equality (22) is Proposition 2(a) in Nau (2003). A proof of a more general result is given there, but the short version is as follows. By assumption, $\mathbf{x} + \mathbf{f} + b(\mathbf{f}, \mathbf{x})\mathbf{1}$ lies on the same indifference curve as \mathbf{x} , and in principle the decision maker could move from \mathbf{x} to $\mathbf{x} + \mathbf{f} + b(\mathbf{f}, \mathbf{x})\mathbf{1}$ along the curve by receiving small increments of \mathbf{f} and being compensated for them at their marginal prices. (Here $b(\mathbf{f}, \mathbf{x})\mathbf{1}$ is the constant vector $\mathbf{1}$ scaled by $b(\mathbf{f}, \mathbf{x})$. The marginal prices are negative and increasingly so, and the equivalent compensation has the opposite sign, yielding a positive number.) To a first-order approximation, the price vector changes linearly with the amount purchased, and

in particular it changes from $\boldsymbol{\pi}(\mathbf{x})$ to approximately $\boldsymbol{\pi}(\mathbf{x}) + D\boldsymbol{\pi}(\mathbf{x})(\mathbf{f} + b(\mathbf{f}, \mathbf{x})\mathbf{1})$ from start to finish. The average value of the price vector along this linear trajectory is the midpoint, $\boldsymbol{\pi}(\mathbf{x}) + \frac{1}{2}D\boldsymbol{\pi}(\mathbf{x})(\mathbf{f} + b(\mathbf{f}, \mathbf{x})\mathbf{1})$. So the total compensation she must receive to accept \mathbf{f} is approximately $-\mathbf{f} \cdot (\boldsymbol{\pi}(\mathbf{x}) + \frac{1}{2}D\boldsymbol{\pi}(\mathbf{x})(\mathbf{f} + b(\mathbf{f}, \mathbf{x})\mathbf{1}))$, which reduces to $-\frac{1}{2}\mathbf{f} \cdot D\boldsymbol{\pi}(\mathbf{x})(\mathbf{f} + b(\mathbf{f}, \mathbf{x})\mathbf{1})$ because $\mathbf{f} \cdot \boldsymbol{\pi}(\mathbf{x}) = 0$ for a neutral asset. The exact amount of compensation she must receive is $b(\mathbf{f}, \mathbf{x})$ by definition, which yields the approximation:

$$b(\mathbf{f}, \mathbf{x}) \approx -\frac{1}{2}\mathbf{f} \cdot D\boldsymbol{\pi}(\mathbf{x})(\mathbf{f} + b(\mathbf{f}, \mathbf{x})\mathbf{1}).$$

The second term in parentheses on the RHS can be ignored in the limit as \mathbf{f} becomes small because $b(\mathbf{f}, \mathbf{x}) = O(\|\mathbf{f}\|^2)$, which yields $b(\mathbf{f}, \mathbf{x}) \approx -\frac{1}{2}\mathbf{f} \cdot D\boldsymbol{\pi}(\mathbf{x})\mathbf{f}$ as claimed. The second equality (23) follows by applying (7) to rewrite $D\boldsymbol{\pi}(\mathbf{x})$ in terms of $DU(\mathbf{x})$ and $D^2U(\mathbf{x})$ and noting that the second term in (7) must yield $\mathbf{0}$ when it is premultiplied by \mathbf{f} because each of its columns is proportional to $\boldsymbol{\pi}(\mathbf{x})$, and $\mathbf{f} \cdot \boldsymbol{\pi}(\mathbf{x}) = 0$ if \mathbf{f} is a neutral asset. The third equality (24) follows from the fact that

$$\begin{aligned} \left[\frac{D^2U(\mathbf{x})}{\mathbf{1} \cdot DU(\mathbf{x})} \right]_{ij} &= \frac{\partial^2 U(\mathbf{x}) / \partial x_i \partial x_j}{\mathbf{1} \cdot DU(\mathbf{x})} \\ &= \frac{\partial U(\mathbf{x}) / \partial x_i}{\mathbf{1} \cdot DU(\mathbf{x})} \frac{\partial^2 U(\mathbf{x}) / \partial x_j \partial x_k}{\partial U(\mathbf{x}) / \partial x_i} = \pi_i(\mathbf{x}) r_{ij}(\mathbf{x}). \quad \blacksquare \end{aligned}$$

Corollary 2: By construction f_{jk}^2 is the same for both bets, so they have the same risk premium, and $\bar{f}_j = 0$ for all j in the $B:A$ bet—hence its ambiguity premium is zero—while $|\bar{f}_j| > 0$ for all j in the $A:B$ bet. If u_j is strictly concave, then $s_j > 0$, so the ambiguity premium is strictly positive for the $A:B$ bet if this is true for one or more j . \blacksquare

Theorem 3: This is Theorem 1(d) in Nau (2006b), but the full details of the proof are not given there, so here they are. Suppose that U is the two-source utility function

$$U(\mathbf{x}) = \sum_j u_j \left(\sum_k v_{jk}(x_{jk}) \right). \quad (51)$$

Its gradient, whose normalization yields the local risk neutral distribution, is

$$[DU(\mathbf{x})]_{jk} = u'_j v'_{jk}, \quad (52)$$

where u'_j is shorthand for $u'_j(\sum_k v_{jk}(x_{jk}))$ and v'_{jk} is shorthand for $v'_{jk}(x_{jk})$, both of which are positive. Its Hessian $D^2U(\mathbf{x})$ is the sum of a diagonal matrix and a block-diagonal matrix constructed as follows. Let the rows of $D^2U(\mathbf{x})$ be indexed by jk and let the columns be indexed by mn , where row jk corresponds to event $A_j B_k$ and column mn corresponds to event $A_m B_n$. The element in row jk and column mn is

$$[D^2U(\mathbf{x})]_{jk,mn} = u'_j v''_{jk} \mathbf{1}_{jk=mn} + u'_j v'_{jk} v'_{jn} \mathbf{1}_{j=m}, \quad (53)$$

where u_j'' and v_{jk}' are shorthands for $u_j''(\sum_k v_{jk}(x_{jk}))$ and $v''(x_{jk})$ respectively, both of which are negative. ($1_{jk=mn}$ is the indicator function for terms in which $jk = mn$, and similarly for $1_{j=m}$.) The corresponding term in the risk aversion matrix \mathbf{R} is by definition the quotient of (53) and (52):

$$r_{jk,mn} = \frac{u_j' v_{jk}'' 1_{jk=mn} + u_j'' v_{jk}' v_{jn}' 1_{j=m}}{u_j' v_{jk}'} = \frac{v_{jk}''}{v_{jk}'} 1_{jk=mn} + \frac{u_j''}{u_j'} v_{jn}' 1_{j=m}$$

By Theorem 2, the local quadratic approximation of the overall uncertainty premium is:

$$\widehat{b}(\mathbf{f}; \mathbf{x}) = \frac{1}{2} \mathbf{f} \cdot \mathbf{\Pi}(\mathbf{x}) \mathbf{R}(\mathbf{x}) \mathbf{f} \quad (54)$$

$$= \frac{1}{2} \sum_{jk} \sum_{mn} \pi_{jk} r_{jk,mn} f_{jk} f_{mn} \quad (55)$$

$$= \frac{1}{2} \sum_{jk} \pi_{jk} \frac{v_{jk}''}{v_{jk}'} f_{jk}^2 + \frac{1}{2} \sum_{jkn} \pi_{jk} \frac{u_j''}{u_j'} v_{jn}' f_{jk} f_{jn} \quad (56)$$

$$= \frac{1}{2} \sum_{jk} \pi_{jk} t_{jk} f_{jk}^2 + \frac{1}{2} \sum_{jk} \pi_{jk} \frac{u_j''}{u_j'} f_{jk} \sum_n v_{jn}' f_{jn} \quad (57)$$

The first term on the RHS is the risk premium. To unpack the second term, let $\alpha = (\sum_{jk} u_j' v_{jk}')^{-1}$, in terms of which $\pi_{jn} = \alpha u_j' v_{jn}'$ and $\pi_j = \sum_n \pi_{jn} = \alpha u_j' \sum_n v_{jn}' = \alpha u_j' v_{j\bullet}'$, whence $v_{jn}' = v_{j\bullet}' (\pi_{jn} / \pi_j)$. Then:

$$\frac{1}{2} \sum_{jk} \pi_{jk} \frac{u_j''}{u_j'} f_{jk} \sum_n v_{jn}' f_{jn} = \frac{1}{2} \sum_j \pi_j \frac{u_j''}{u_j'} \sum_k \frac{\pi_{jk}}{\pi_j} f_{jk} \sum_n v_{jn}' f_{jn} \quad (58)$$

$$= \frac{1}{2} \sum_j \pi_j \frac{u_j''}{u_j'} \bar{f}_j \sum_n v_{jn}' f_{jn} \quad (59)$$

$$= \frac{1}{2} \sum_j \pi_j \frac{u_j''}{u_j'} \bar{f}_j \left(v_{j\bullet}' \sum_n \frac{\pi_{jn}}{\pi_j} f_{jn} \right) \quad (60)$$

$$= \frac{1}{2} \sum_j \pi_j \frac{u_j''}{u_j'} v_{j\bullet}' \bar{f}_j^2 = \frac{1}{2} \sum_i \pi_i s_i \bar{f}_i^2 \quad (61)$$

which is the ambiguity premium. \blacksquare

Theorem 4: The second-order utility model can be viewed as a special case of the source-dependent utility model in which the first source of risk is the first-order payoff and the second source of risk is the credal state. With a change of notation (i taking the place formerly occupied by j , and jk taking the place formerly occupied by k alone), Theorem 3 yields

$$\widehat{b}(\mathbf{f}; \mathbf{x}) = \frac{1}{2} \sum_i \pi_i s_i \bar{f}_i^2 + \frac{1}{2} \sum_{ijk} \pi_{ijk} t_{jk} f_{jk}^2. \quad (62)$$

and i can then be integrated out of the second term because the first-order payoffs do not depend on it:

$$\sum_{ijk} \pi_{ijk} t_{jk} f_{jk}^2 = \sum_{jk} t_{jk} f_{jk}^2 \sum_i \pi_{ijk} = \sum_{jk} \pi_{jk} t_{jk} f_{jk}^2. \quad \blacksquare \quad (63)$$

Corollary 3: First consider the simpler situation in which $v'_{jk} = y_j z_k$ (i.e., z_k is the same within A and \bar{A}). Then the joint risk neutral probabilities satisfy

$$\pi_{ijk}(\mathbf{x}) \propto (\mu_i u'_i)(p_{j|i} q_k) y_j z_k \quad (64)$$

$$= (\mu_i u'_i)(p_{j|i} y_j)(q_k z_k) \quad (65)$$

which implies that $\pi_{jk|i} \propto q_k z_k$ for any fixed i and j , and also $\pi_{jk} \propto q_k z_k$ for any fixed j , and therefore also $\pi_{jk} \propto \pi_{jk|i}$ for any fixed i and j . If $\mathbf{f} = \mathbf{f}_{B:A}(\Delta)$, then by definition $\sum_{j \in A} \sum_k \pi_{jk} f_{jk} = 0$ and also $\sum_{j \in \bar{A}} \sum_k \pi_{jk} f_{jk} = 0$. Because $\pi_{jk} \propto \pi_{jk|i}$ for fixed i and j , it follows that $\sum_{j \in A} \sum_k \pi_{jk|i} f_{jk} = 0$ and $\sum_{j \in \bar{A}} \sum_k \pi_{jk|i} f_{jk} = 0$ for every i , which together imply $\sum_j \sum_k \pi_{jk|i} f_{jk} = 0$ for every i , which means $\bar{f}_i = 0$ for every i . Hence the ambiguity premium, which is a weighted sum of $\{\bar{f}_i^2\}$, must be zero. The proof is completed by observing that the same chain of implication holds if z_k is allowed to depend on whether $j \in A$ or $j \in \bar{A}$, because the cases $j \in A$ and $j \in \bar{A}$ can be treated separately in proving that $\sum_{j \in A} \sum_k \pi_{jk|i} f_{jk} = 0$ and $\sum_{j \in \bar{A}} \sum_k \pi_{jk|i} f_{jk} = 0$. Note that this argument does not apply to $\mathbf{f} = \mathbf{f}_{A:B}(\Delta)$ in general, because π_{jk} need not be proportional to $\pi_{jk|i}$ for fixed i and k . \blacksquare