

BELIEFS AS STATE PRICES: REVISITED

Simon Grant

Australian National University, Canberra, AUSTRALIA

simon.grant@anu.edu.au

Abstract

In a finite-state, multi-good economy, when do local state prices reveal beliefs rather than state-dependent tastes? Under the sure-thing principle, suppose small Arrow claims quote one relative state-price book at every sure bundle, on purchase and sale sides, and in every commodity currency. If one state has finite positive one-sided marginal values, this condition forces the additive state utilities to be positive affine transformations of a common Bernoulli utility. State Independence is therefore implied and the normalized book is the unique subjective probability vector. No restriction on risk attitudes is imposed, and finite-slope kinks are allowed.

JEL Classification: D80, D81

Key words: subjective expected utility, state prices, subjective probability, State Independence, kinks.

1 Introduction

When does the price of a state-contingent claim deserve to be called a belief? In the familiar Arrow–Debreu formulation of subjective expected utility (SEU) (Arrow, 1964; Debreu, 1959; Savage, 1954), an act is a vector of state-contingent consumption bundles and, at a sure bundle c , the marginal rate of substitution between good i delivered in states s and t is

$$\frac{p_s \partial_i u(c)}{p_t \partial_i u(c)} = \frac{p_s}{p_t}.$$

The marginal utility of the physical good cancels, so relative prices of state claims reveal subjective odds. This is an appealing economic reading of subjective probability: beliefs, like tastes in consumer theory, are inferred from tradeoffs. The state-preference tradition going back at least to Hirshleifer (1965) and Yaari (1969), and especially the local-belief analysis of Nau (2011), makes this interpretation transparent.

There is, however, a circularity hidden in the one-line cancellation. It presumes that the same Bernoulli utility u evaluates a physical bundle in every state. That is precisely the state-independence structure needed to distinguish beliefs from state-dependent tastes. If a unit of food, health care, or wealth has a different marginal utility in different states for reasons other than likelihood, an observed inter-state marginal rate of substitution need not identify a probability ratio at all. The standard price formula therefore tells us how to read beliefs once SEU has been obtained; by itself it does not tell us when that reading is justified.

This paper reverses the argument. Suppose first that the sure-thing principle gives an additive representation across states. Now impose one economically simple requirement: at every sure bundle, on either side of the market, and whatever physical good denominates the claim, the relative price of state s in terms of state t never reverses. If one state has finite positive one-sided marginal values, this single-book condition implies that there are numbers $q_s > 0$ and one common Bernoulli utility u such that every additive state utility has the form

$$u_s(c) = q_s u(c) + a_s.$$

Normalizing q gives a unique probability vector p , and preferences are represented by

$$\sum_s p_s u(x_s).$$

State Independence, Savage’s P3, is a consequence rather than an independent assumption.

The result imposes no attitude toward risk: concavity, convexity, local linearity, and changing curvature of u are all permitted. It also permits finite-slope kinks. Same-side comparisons still recover the same odds because the right-hand marginal utility cancels on the buy side and the left-hand marginal utility on the sell side. Thus familiar piecewise-linear utilities are not excluded by differentiability.

The closest antecedents clarify what is new. Wakker (1984, 1989) showed that Cardinal Coordinate Independence (CCI) characterizes finite-state SEU without separately assuming State Independence or differentiability. On the common domain, CCI is an alternative route to the same representation, but through cross-coordinate strength-of-preference comparisons rather than local Arrow-claim prices. Hens (1992) combines the sure-thing principle and State Independence with constant inter-state marginal rates of substitution along certainty; here the price condition itself supplies State Independence. Werner’s (2002) differentiable small-risk compensation condition likewise forces proportional state utilities and extends to multiple goods, but is indexed by candidate probabilities. The present condition instead recovers probabilities from the observed book and permits finite-slope kinks.

These distinctions matter beyond axiomatization. SEU combines beliefs p and tastes u . The theorem shows that coherent local state prices not only identify the former but force the separation from tastes that makes those prices interpretable as probabilities. Section 2 formulates the single-book condition; Section 3 proves the characterization and redundancy of State Independence; Section 4 treats kinks and risk attitudes; Section 5 revisits the architecture of SEU; and Section 6 relates the result to existing foundations.

2 State-Contingent Consumption and One Price Book

Let $S = \{1, \dots, n\}$ be a finite set of states, with $n \geq 3$, and let the consequence space be $C = \mathbb{R}_+^k$. An act is a vector

$$x = (x_1, \dots, x_n) \in X := C^n,$$

where x_s is the bundle consumed if state s occurs. We identify $c \in C$ with the constant act $(c, \dots, c)$. Preferences $\succsim$ on X are complete, transitive, continuous, and strongly monotone: $x \geq y$ and $x \neq y$ imply $x \succ y$.

For an event $B \subseteq S$, let $x_B y$ be the act that agrees with x on B and with y on $S \setminus B$. For a singleton $\{s\}$ and $c \in C$, write $c_s x$ for the act that yields c in state s and agrees with x elsewhere. The first condition is Savage's (1954) P2.

P2 (No Common-Consequence Reversals) *For all $x, y \in X$ and every $B \subseteq S$,*

$$x \succsim y_B x \implies x_B y \succsim y.$$

The condition says that a preference between two acts cannot depend on what they have in common off the event on which they differ. With the maintained continuity and monotonicity assumptions and at least three essential state coordinates, standard additive-representation results give

$$V(x) = \sum_{s=1}^n u_s(x_s), \tag{1}$$

where each $u_s : C \rightarrow \mathbb{R}$ is continuous and strictly increasing in every coordinate; the component functions are jointly cardinal, so another additive representation has components $au_s + b_s$ with the same $a > 0$ for every state (Debreu, 1960, Theorem 3); see also Wakker (1989, Theorem III.6.6). At this stage there is no reason for u_s and u_t to rank two multi-good bundles in the same way. The additive representation separates states; it has not yet separated beliefs from tastes.

The information used to perform that separation will come from small trades at certainty. Fix an interior sure bundle $c \in \mathbb{R}_{++}^k$, a state s , a good i , and a small amount $a > 0$. Let e_i denote the i th unit vector and let $x_{c,s,i}^+(a)$ (respectively, $x_{c,s,i}^-(a)$) be the act that delivers $c + ae_i$ (respectively, $c - ae_i$) in state s and c in every other state. That is, $x_{c,s,i}^+(a)$ is a small purchase of an Arrow claim on state s , denominated in good i , while $x_{c,s,i}^-(a)$ is a small sale of the same claim.

A rate $\rho > 0$ sets one unit of the claim on s against ρ units of the otherwise identical claim on t . Each comparison sorts a rate into one of three classes: it favors s , it favors t , or it does neither.

Definition 1 (Same-side state-claim comparisons) *Fix distinct states s, t , $c \in \mathbb{R}_{++}^k$, a good i , and $\rho > 0$.*

(i) On the buy side, ρ strictly favors s over t at (c, i) if there are $\rho' > \rho$ and $\bar{a} > 0$ such that

$$x_{c,s,i}^+(a) \succ x_{c,t,i}^+(\rho'a) \quad \text{for every } a \in (0, \bar{a});$$

and ρ strictly favors t over s at (c, i) if there are $\rho' \in (0, \rho)$ and $\bar{a} > 0$ such that

$$x_{c,t,i}^+(\rho'a) \succ x_{c,s,i}^+(a) \quad \text{for every } a \in (0, \bar{a}).$$

(ii) On the sell side, the same comparisons are made with sales in place of purchases and the preferences reversed: ρ strictly favors s over t at (c, i) if there are $\rho' > \rho$ and $\bar{a} > 0$ such that

$$x_{c,t,i}^-(\rho'a) \succ x_{c,s,i}^-(a) \quad \text{for every } a \in (0, \bar{a});$$

and ρ strictly favors t over s at (c, i) if there are $\rho' \in (0, \rho)$ and $\bar{a} > 0$ such that

$$x_{c,s,i}^-(a) \succ x_{c,t,i}^-(\rho'a) \quad \text{for every } a \in (0, \bar{a}).$$

In words, ρ favors s when one small claim on s still beats ρ' small claims on t at some rate ρ' above ρ , and favors t when ρ' small claims on t already beat one small claim on s at some rate ρ' below ρ . The gap between ρ and ρ' is what leaves the true marginal rate unclassified: under regularity, rates below the ratio of marginal utilities favor s and rates above it favor t (Lemma 1). On the sell side “beats” means that giving the claim up is the smaller sacrifice.

The economic restriction can be stated before its formal version. If one state commands a higher price than another when small claims are traded in one good at one sure bundle, that relative ranking may not reverse when the consumer moves to another sure bundle, switches from buying to selling, or denominates the claims in another good. Beliefs should not change merely because an Arrow claim pays bread rather than shelter, nor because the consumer is richer, nor because the trade is approached from the other side of a kink.

One State-Price Book *For every pair of distinct states s, t , if a rate $\rho > 0$ strictly favors s over t in one same-side comparison, it does not strictly favor t over s in any other same-side comparison, whatever the interior sure bundle, good, or side of the market.*

The condition is stated entirely in terms of preference. To interpret it as a marginal-price restriction we need local regularity, but not smoothness, and we need it for one state only. Under an additive representation (1), call a state t *price regular* if, for every $c \in \mathbb{R}_{++}^k$ and every good i , the coordinate section of u_t has finite strictly positive left and right derivatives

$$D_i^- u_t(c), \quad D_i^+ u_t(c).$$

The property is invariant to the joint positive affine freedom of the additive representation, so that for preferences satisfying P2 it is a property of the preference relation itself, in the same way that risk aversion is a property of a preference through the Bernoulli utility that the preference determines up to scale. The two derivatives need not coincide or vary continuously. Nothing is assumed about the other states: their utilities need only be continuous and strictly increasing, and the theorem below shows that they inherit price regularity from the one state that has it.

The requirement that the book be global is substantive when $k \geq 2$. Observing state prices only on one certainty ray cannot discipline how the u_s behave in directions transverse to that ray; one can alter state-specific tastes off the ray without changing any quote observed on it. The role of the global condition is to use the entire certainty set, and a spanning collection of commodity currencies, to close that gap.¹

3 From Local State Prices to Subjective Expected Utility

The proof has two steps. First, the behavioral condition pins down, for each state, the rates at which small claims on that state exchange against claims on the one price regular state; because that state has one-sided marginal utilities, the comparison brackets the marginal utility of every other state without assuming it exists. Second, because the bracket is the same in every physical currency and at every certainty position, it integrates into a common Bernoulli utility.

Throughout, fix a price regular state and call it t^* . For any other state s , an interior bundle c , and a good i , write $\Delta_s(a) := u_s(c + ae_i) - u_s(c)$ for the utility gain

¹For example, with two goods a perturbation proportional to a bounded increasing transformation of $c_1 - c_2$ can be chosen to vanish on the diagonal $c_1 = c_2$ while changing state-specific rankings away from it. No collection of quotes confined to that diagonal can detect the perturbation.

in state s from acquiring a units of good i , and let

$$\underline{D}_i^+ u_s(c) := \liminf_{a \downarrow 0} \frac{\Delta_s(a)}{a}, \quad \overline{D}_i^+ u_s(c) := \limsup_{a \downarrow 0} \frac{\Delta_s(a)}{a}$$

be the lower and upper right Dini derivates of the coordinate section of u_s . They exist in $[0, \infty]$ because u_s is increasing in good i , and they coincide with the ordinary right derivative whenever it exists. The lower and upper left Dini derivates $\underline{D}_i^- u_s(c)$ and $\overline{D}_i^- u_s(c)$ are defined in the same way from $u_s(c) - u_s(c - ae_i)$.

Lemma 1 (Local comparisons bracket the book) *Assume P2 and let t^* be price regular. Fix a state $s \neq t^*$, an interior bundle c , and a good i . On the buy side, a rate ρ strictly favors s over t^* if and only if $\rho < \underline{D}_i^+ u_s(c) / D_i^+ u_{t^*}(c)$, and strictly favors t^* over s if and only if $\rho > \overline{D}_i^+ u_s(c) / D_i^+ u_{t^*}(c)$. The sell side is identical with left derivates and $D_i^- u_{t^*}(c)$. Consequently, under One State-Price Book there are numbers $0 \leq L_s \leq U_s \leq \infty$ such that, for every finite $q \in [L_s, U_s]$, every interior c , every good i , and $\sigma \in \{+, -\}$,*

$$\underline{D}_i^\sigma u_s(c) \leq q D_i^\sigma u_{t^*}(c) \leq \overline{D}_i^\sigma u_s(c). \quad (2)$$

Proof. Write $\beta := D_i^+ u_{t^*}(c)$, a finite positive number, and $\Delta_{t^*}(b) := u_{t^*}(c + be_i) - u_{t^*}(c)$, so that $\Delta_{t^*}(b)/b \rightarrow \beta$ as $b \downarrow 0$. By additivity the two purchases in Definition 1 differ in value by $\Delta_s(a) - \Delta_{t^*}(\rho'a)$. If $\underline{D}_i^+ u_s(c) > \beta\rho$, choose ρ' with $\rho < \rho' < \underline{D}_i^+ u_s(c)/\beta$; for all small a the quotient $\Delta_s(a)/a$ then exceeds $\beta\rho'$ by a fixed margin while $\Delta_{t^*}(\rho'a)/a$ converges to $\beta\rho'$, so the s -claim is strictly preferred and ρ favors s . Conversely, if ρ favors s at some rate $\rho' > \rho$, then $\Delta_s(a)/a > \Delta_{t^*}(\rho'a)/a$ for all small a , and taking the lower limit gives $\underline{D}_i^+ u_s(c) \geq \beta\rho' > \beta\rho$. The statement for t^* is symmetric with the upper derivate, and the sell side replaces gains from purchases by losses from sales.

For the bracket, call a triple $\kappa = (c, i, \sigma)$ a context and let $\underline{r}_\kappa$ and $\bar{r}_\kappa$ denote the two thresholds just found, the lower and upper Dini derivates of u_s divided by the corresponding one-sided derivative of u_{t^*} . Rates below $\underline{r}_\kappa$ favor s in context κ and rates above $\bar{r}_{\kappa'}$ favor t^* in context κ' , so One State-Price Book forbids $\underline{r}_\kappa > \bar{r}_{\kappa'}$ for any two contexts. Hence $L_s := \sup_\kappa \underline{r}_\kappa \leq \inf_\kappa \bar{r}_\kappa =: U_s$, and any finite $q \in [L_s, U_s]$ satisfies $\underline{r}_\kappa \leq q \leq \bar{r}_\kappa$ in every context, which is (2). ■

When u_s also has one-sided derivatives, the bracket collapses and the switch point is the ratio $D_i^\sigma u_s(c) / D_i^\sigma u_{t^*}(c)$ of same-side marginal utilities, exactly as the one-

line cancellation in the introduction suggests. The lemma shows that nothing about u_s needs to be assumed for the condition to have bite: the regular state supplies the yardstick, and the no-reversal requirement forces every other state's marginal valuations to fit between two constant multiples of it.

The next step turns the bracket into an equation. No mixed derivatives, differentiability across a kink, or path independence theorem is needed. Coordinate directions are enough because any two points in the interior of $\mathbb{R}_+^k$ are connected by an axis-parallel path, and the mean value theorem is replaced by an elementary fact about one-sided derivatives.

Lemma 2 (Local proportionality integrates globally) *Let f be a continuous real function on an open interval whose upper right Dini derivate is nonnegative at every point. Then f is nondecreasing. Consequently, if a continuous function $h : C \rightarrow \mathbb{R}$ has, along every coordinate line through every interior bundle, lower right Dini derivate at most zero and upper right Dini derivate at least zero, then h is constant on C .*

Proof. Suppose $f(b) < f(a)$ for some $a < b$ in the interval. Choose $\epsilon > 0$ so small that $g(\tau) := f(\tau) + \epsilon(\tau - a)$ still satisfies $g(b) < g(a)$; the upper right derivate of g is then at least ϵ at every point. Pick v strictly between $g(b)$ and $g(a)$ and let z be the last point of $[a, b]$ at which $g \geq v$. By continuity $g(z) \geq v$, and $z < b$ because $g(b) < v$. Every point to the right of z has $g < v \leq g(z)$, so every right difference quotient of g at z is negative and the upper right derivate of g at z is at most zero, a contradiction. For the second statement, apply the first to h and to $-h$ along each coordinate line through an interior bundle, using that the upper right derivate of $-h$ is the negative of the lower right derivate of h : the section is nondecreasing and nonincreasing, hence constant. Any two interior points can be connected by finitely many axis-parallel segments through interior points, so h is constant on $\mathbb{R}_{++}^k$, and continuity extends the same constant to the boundary of $C = \mathbb{R}_+^k$. ■

Theorem 1 (Separating Beliefs from Tastes) *Let preferences satisfy the maintained assumptions and P2, and let some state be price regular. The following are equivalent.*

- (A) *Preferences satisfy One State-Price Book.*
- (B) *There exist strictly positive probabilities $p_1, \dots, p_n$ and a continuous, strongly monotone Bernoulli utility $u : C \rightarrow \mathbb{R}$, with finite strictly positive one-sided*

coordinate derivatives on $\mathbb{R}_{++}^k$, such that

$$x \succsim y \iff \sum_{s=1}^n p_s u(x_s) \geq \sum_{s=1}^n p_s u(y_s). \quad (3)$$

The probability vector is unique and u is unique up to a positive affine transformation. Moreover, in every context the book quotes p_s/p_t for a claim on s against a claim on t .

Proof. Suppose (A). P2 gives the additive representation (1); let t^* be price regular and fix $s \neq t^*$.

First, $U_s < \infty$. Otherwise, since $U_s = \inf_{\kappa} \bar{r}_{\kappa}$, we would have $\bar{r}_{\kappa} = \infty$ in every context κ . In particular, $\bar{D}_i^+ u_s(c) = \infty$ for every interior c and every good i . Hence, for every $M > 0$, each coordinate section $z \mapsto u_s(c + ze_i) - Mz$ has nonnegative upper right Dini derivate throughout its interior domain and is therefore nondecreasing by Lemma 2. Thus, for every sufficiently small $a > 0$,

$$u_s(c + ae_i) - u_s(c) \geq Ma.$$

Since this holds for every $M > 0$, it contradicts the finite-valuedness of u_s . Hence $U_s < \infty$.

Now take any $q \in [L_s, U_s]$ and put $h := u_s - q u_{t^*}$, a continuous function on C . Because u_{t^*} has a right derivative β along good i at c , the right Dini derivatives of h there are those of u_s shifted by $q\beta$, so by (2) the lower one is at most zero and the upper one at least zero. Lemma 2 makes h constant on C . If $L_s < U_s$, the same argument applied to two distinct values $q < q'$ in $[L_s, U_s]$ would make the difference $(q' - q)u_{t^*}$ constant, contradicting strict monotonicity. Hence $L_s = U_s =: q_s$. If $q_s = 0$, then (2) says the lower and upper right derivatives of u_s itself straddle zero everywhere, and Lemma 2 makes u_s constant along coordinate lines, again a contradiction. Thus $q_s \in (0, \infty)$ and there is a constant a_s with

$$u_s(c) = q_s u_{t^*}(c) + a_s \quad \text{for every } c \in C. \quad (4)$$

Put $q_{t^*} = 1$ and $a_{t^*} = 0$. Then

$$V(x) = \sum_s q_s u_{t^*}(x_s) + \sum_s a_s.$$

Set $u = u_{t^*}$ and $p_s = q_s / \sum_r q_r$ to obtain (3); u has the stated one-sided derivatives because t^* is price regular, and by (4) so does every other state utility.

Conversely, suppose (B). An additive representation has components $u_s = p_s u$, so every state is price regular and, for every c, i, σ ,

$$\frac{D_i^\sigma u_s(c)}{D_i^\sigma u_t(c)} = \frac{p_s}{p_t}.$$

By Lemma 1, with any state as the regular one, a rate strictly favors s over t in a context exactly when it is below p_s/p_t and t over s exactly when it is above; no rate does both, which is what One State-Price Book requires. Uniqueness follows from the joint cardinal uniqueness of additive representations with essential coordinates: two SEU representations must induce proportional state weights, and normalization makes the probability vectors identical; the common Bernoulli indices then differ by a positive affine transformation. ■

The restriction $n \geq 3$ enters only in using P2 to obtain an additive representation. The state-price and integration arguments themselves require only an additive representation. Thus, with two states Theorem 1 continues to hold if additive separability is assumed directly.

The economic proof is shorter than the formal one. P2 gives one state utility for each coordinate. The single-book condition says that every local marginal valuation in state s is the same state-specific multiple of the corresponding valuation in the reference state, regardless of where the comparison is made, from which side it is approached, or which good supplies the units. Those local multiples therefore integrate to positive affine transformations of one common utility. The multiples are the state prices, and after normalization they are the probabilities.

Corollary 2 (State Independence is implied) *Under the maintained assumptions and P2, if some state is price regular then One State-Price Book implies Savage's State Independence condition: for all consequences c, c' , states s , and acts x ,*

$$c \succsim c' \iff c_s x \succsim c'_s x.$$

Every state is then price regular, and in every context the same-side switch point is well defined.

Proof. By Theorem 1, preferences have representation (3). Since $p_s > 0$, the right-hand comparison is equivalent to $u(c) \geq u(c')$, which is also the ranking of the

two constant acts. The remaining claims follow from $u_s = p_s u$ and the regularity of u . ■

Two parts of One State-Price Book do different work. Invariance across goods prevents the apparent state weights from absorbing state-specific relative tastes for physical commodities. Invariance across positions prevents them from absorbing non-linear state-specific transformations of a common index. Together they turn a collection of local exchange rates into one global decomposition of beliefs from tastes. Buy–sell invariance is not what admits a kink; the one-sided derivatives do that. Indeed the proof of Theorem 1 uses buy-side contexts only, together with the right-hand derivatives of the reference state: position and currency invariance on one side of the market already deliver the representation, and equality of the two sides’ books then follows from it wherever the Bernoulli utility has slopes on the other side. The condition is nonetheless stated on both sides, because the two-sided book is what separates a common kink in tastes from a set of priors in Section 5.

4 Kinks, Risk Attitudes, and the Price Interpretation

Nothing in Theorem 1 says that the individual is risk averse. Nor does it say that she is risk neutral “in the small” in the usual Arrow–Pratt sense. For a one-good consequence, concavity of u delivers risk aversion, convexity delivers risk seeking, and linearity delivers risk neutrality. The theorem leaves all three possibilities open, as well as Bernoulli utilities whose curvature changes with wealth. With several goods the language of risk attitudes requires more care, but the same point remains: the price condition restricts the separation of states from tastes, not the shape of tastes themselves.

This distinction is easy to see at a kink. Let $b \gg 0$ and define $m(c) = b \cdot c$. For $\theta > 0$, define

$$g(z) = \begin{cases} z, & z \leq 1, \\ 1 + \theta(z - 1), & z \geq 1, \end{cases}$$

and set $u(c) = g(m(c))$, after rescaling b so that the kink occurs at the sure bundle under consideration. If $\theta < 1$ the kink is downward; if $\theta > 1$ it is upward. At a bundle with $m(c) = 1$, adding a small amount of good i has marginal utility θb_i ,

while giving up a small amount has marginal utility b_i . Under SEU the corresponding state- s marginal values are $p_s\theta b_i$ and p_sb_i . Their levels differ across sides, possibly by a first-order amount, but

$$\frac{p_s\theta b_i}{p_t\theta b_i} = \frac{p_sb_i}{p_tb_i} = \frac{p_s}{p_t}.$$

The odds survive because same-side comparisons keep the taste term common to numerator and denominator.

This is more than a technical relaxation of differentiability. Piecewise-linear and kinked Bernoulli utilities are useful precisely because they let economists test what a proposition about risk requires. Smooth expected utility generates second-order compensation for sufficiently small mean-zero risks (Pratt, 1964); the distinction between first- and second-order local risk attitudes is central in Segal and Spivak (1990). A kink falls outside that smooth argument. Excluding it by assumption would make the state-price characterization least informative at exactly the examples used to probe the boundary of standard risk results.

There is also a revealing contrast with supporting-price definitions of beliefs. Riquetti et al. (2008) identify subjective beliefs with normalized supporting prices in a broad class of convex preferences. At a concave taste kink, the supporting set of an SEU preference can contain many normals because each state's marginal utility can be selected from the common subgradient interval. Same-side quotes ask a different local question. Comparing right derivatives with right derivatives, or left with left, cancels the common slope and returns the unique ratio p_s/p_t . A set of supporting prices and a unique same-side state-price book can therefore coexist even under ordinary SEU.

The pricing of small state claims is neutral in one respect only. Once a same-side comparison holds the physical marginal value of the currency fixed, the relative valuation of the two states does not drift with the position, the currency, or the side of the market. It says nothing about how that marginal value itself changes with consumption, which is where attitudes toward risk reside; the Bernoulli utility may be as curved or as kinked as the maintained one-sided regularity permits. This is a different neutrality from the vanishing first-order compensation for small risks that Werner (2002) uses to characterize differentiable expected utility, which is a statement about the second-order behavior of the utility of wealth and fails at any kink.

5 The Architecture of SEU: A Tale of Three Stories

SEU is often presented as one formula. It is more useful to see the formula as the compact expression of a three-story construction. On the ground floor is Savage’s consequentialist description of uncertain choice: an act specifies a consequence for each state of the world. The middle floor is probabilistic sophistication (Machina and Schmeidler, 1992): a single probability measure converts each act into a distribution over consequences, so the identity of a state matters only through the probability it contributes to that distribution. At the top is the expected-utility rule: those induced lotteries are evaluated by the expectation of one Bernoulli utility index. In an SEU representation

$$\sum_s p_s u(x_s),$$

the state-contingent consequences x_s come from the ground floor, the probability vector p from the middle floor, and the Bernoulli utility u from the top.

The state-price theorem illuminates the seam between them. When a small amount of the same physical good is moved across states at certainty, the taste term cancels. What remains is the relative state price. Requiring one such book everywhere does more than recover p : it forces the local taste terms to fit together into a single u . Beliefs are therefore not whatever coefficients happen to appear in an additive representation. They are the state weights that remain invariant when the physical currency, the sure consumption bundle, and the side of the local trade are changed.

This decomposition helps locate familiar departures from SEU. The Allais pattern attacks the top. A decision maker can retain a single probability measure and still reject the expected-utility rule for the lotteries that measure induces. Rank-dependent utility is the canonical example: probability remains meaningful, but it enters through rank-dependent decision weights rather than linearly as in expected utility (Quiggin, 1982). In state-price language, the difficulty is not necessarily the existence of beliefs; it is the rule by which probabilities and utilities are combined.

Ellsberg strikes lower. The familiar urn choices cannot be reconciled with one probability measure over the colors while maintaining the natural dominance ordering of bets. Multiple-prior models respond by replacing a single p with a set of priors, as in Gilboa and Schmeidler (1989). A local state-price analysis then need not deliver

one invariant book. Buy-side and sell-side valuations can be governed by different supporting priors, so relative odds may split. The theorem should not be read as a universal diagnostic for ambiguity—some non-singleton prior sets can generate coincident local ratios—but a split book identifies a failure of the single-probability structure that SEU requires.

State-dependent utility creates a different problem. The ground floor description of acts can remain intact and preferences can even be additively separable, yet the same physical bundle may have different ordinal significance in different states. Then an inter-state marginal rate of substitution mixes two things one would like to keep separate: how likely a state is and how valuable the commodity is in that state. This is the identification problem emphasized in the state-dependent-utility literature (Karni, 1985; Schervish et al., 1990); Schervish et al. make it vivid with a decision maker who bets coherently in dollars and coherently in yen and reveals a different probability in each currency. One State-Price Book is deliberately demanding at precisely this point. If the apparent probability of a state changes when the Arrow claim is paid in a different commodity, the coefficient is not a belief in the state-independent SEU sense.

Finally, a kink in tastes need not damage any floor. A fixed piecewise-linear Bernoulli utility is still one function u shared across states. It may make the absolute marginal value of a gain differ from that of a loss, but same-side state odds remain unchanged. Reference dependence becomes a departure only when the reference itself moves with the choice problem in a way that changes how otherwise identical consequences are evaluated. Nondifferentiability and non-SEU behavior are therefore separate issues. The one-book characterization is useful partly because it keeps them separate.

The question “which floor gave way?” is not a taxonomy for its own sake. It disciplines the interpretation of alternatives. Allais-type models may preserve beliefs while changing risk evaluation; Ellsberg-type models question whether one probability suffices; state dependence questions whether beliefs can be separated from tastes at all. The state-price result gives that distinction a local economic reading: ask what survives when we look at the exchange rates of state-contingent consumption.

6 Antecedents and Relation to Existing Characterizations

The broad result that finite-state SEU can be derived without separately postulating Savage's P3 (State Independence) is not new. The strongest antecedent is Wakker's Cardinal Coordinate Independence (CCI): Wakker (1984) established the result for monetary consequences, and Wakker (1989, Theorem IV.2.7) extended it to a connected topological consequence space. For a continuous weak order on C^n with sufficiently rich C , CCI requires cardinal consistency across state coordinates and is equivalent to a representation $\sum_s p_s u(x_s)$, without the one-sided derivative regularity required here.

On their common domain, CCI and P2 plus One State-Price Book are alternative routes to the same representation. CCI asks whether preference differences across coordinates are cardinally consistent; the present condition asks whether local exchange rates between Arrow claims form one book across positions, sides, and commodity currencies. The latter pays for its price interpretation with one-sided regularity, but the probabilities emerge as marginal state prices rather than coefficients recovered from a global strength-of-preference test.

Hens (1992) is closer in economic language. The sure-thing principle and State Independence first give additively separable state utilities with common ordinal content; smoothness and constant inter-state marginal rates of substitution along certainty then make the state transformations affine. Here a sufficiently rich set of local state prices supplies both the state-independent ordinal content and the cardinal proportionality, so State Independence need not be assumed separately.

The closest local-to-global proof idea is in Werner (2002). For a given probability vector π , Werner defines a small-risk compensation condition: compensation for every π -mean-zero risk is of smaller order than the risk itself. Under differentiability and the sure-thing principle, the condition implies $v'_s(x) = (\pi_s/\pi_1)v'_1(x)$ and hence expected utility; his paper explicitly extends the argument to multiple goods by measuring compensation in each individual good, obtaining proportional state gradients. This is the same integration mechanism in a different behavioral language. There are three differences relevant here. Werner's condition is indexed by the candidate probabilities, whereas the state-price book reveals the p_s endogenously; it is expressed through compensation for mean-zero risks rather than direct comparisons of Arrow claims; and his general expected-utility theorem is differentiable, whereas same-side comparisons

retain finite positive one-sided slopes at kinks. The published version of his paper (Werner, 2005) presents the companion result that the sure-thing principle and risk aversion with respect to a given probability vector yield concave expected utility. The comparison also makes plain that the present condition restricts the separation of states from tastes and not the curvature of tastes.

The state-price interpretation itself has older roots in the contingent-commodity framework of Arrow (1964) and Debreu (1959). Hirshleifer (1965) developed it as the state-preference approach to investment, Yaari (1969) used it to study risk aversion and portfolio choice, and Nau (2011) makes the local interpretation explicit: revealed beliefs are encoded in risk-neutral probabilities, the relative marginal rates of substitution between states, and local risk attitudes in their derivatives. Rigotti et al. (2008) develop a supporting-price notion of beliefs that remains meaningful under ambiguity. The present result revisits this line from a different direction. It asks when a collection of local state prices can be integrated into the state-independent taste/belief decomposition of SEU itself.

There is also a logically distinct literature on the redundancy of P3. Hartmann (2020) proves that, in Savage’s system, P1, P2, P4, and P7 imply P3. A related observation appears in Gul and Pesendorfer (2014), who recover from preferences the set of ideal events: those events for which the sure-thing principle holds both for the event and for its complement. On this preference-defined event structure, with monetary prizes ordered by amount, monotonicity in prizes together with their comparative-probability axiom delivers P3, as they note in proving their representation.² In both arguments the work is done by dominance in the sure ranking of consequences, supplied by P7 in Hartmann’s case and by monotonicity in money in Gul and Pesendorfer’s, which P4 then transfers from one pair of consequences to another. With consequences that are bundles of several goods, dominance in the sure ranking of bundles is most of State Independence and is not implied by monotonicity in the goods, so that route is closed here. The question in this paper is whether the economically observable price structure can itself do the work usually assigned to State Independence.

Finally, several modern foundations separate utility from beliefs without marginal-price arguments. Gul (1992) gives a finite-state Savage theorem using an independence-type condition, and Chateauneuf et al. (2026) use betting consistency to identify cardinal utility independently of beliefs and, with the sure-thing principle, obtain SEU under continuity. Their separation runs in the opposite

²See Appendix B, Lemma B1 and the proof of Lemma B3.

direction from the one here: betting consistency pins down cardinal utility first, and probabilities then emerge as its coefficients, whereas coherent state prices pin down beliefs first and a common utility follows. These approaches reach subjective expected utility with no regularity beyond continuity, whereas the present theorem requires one state utility to have finite positive one-sided slopes. The reason to use the one-book route is not that finite-state SEU otherwise lacks a foundation. It is that state prices are an object economists already interpret, trade, and estimate, and their coherence provides a direct way to see where the separation of beliefs from tastes comes from.

7 Conclusion

The usual state-price formula for subjective probability is persuasive because the marginal utility of the commodity cancels. But that cancellation is normally written after state-independent expected utility has already been assumed or derived. The result here runs the logic in the other direction. Given additive separability and finite positive one-sided marginal values in one state, if local Arrow claims quote one relative state-price book at every sure bundle, on both sides of the market, and in every commodity currency, the local book integrates into the common Bernoulli utility and probability vector of SEU.

This leaves attitudes toward risk where they belong: in the shape of the Bernoulli utility. The book can be coherent whether that utility is concave, convex, changes curvature, or has finite-slope kinks. What coherence rules out is state dependence that masquerades as belief. Seen this way, subjective probabilities are not merely coefficients in a representation. They are the relative prices of states that survive every permissible change in the physical terms of the local trade.

References

- Arrow, K. J. (1964). The role of securities in the optimal allocation of risk-bearing. *Review of Economic Studies* 31(2), 91–96. Translation of “Le rôle des valeurs boursières pour la répartition la meilleure des risques,” *Économétrie*, Colloques Internationaux du CNRS, Paris, 1953, pp. 41–47.

- Chateauneuf, A. J., F. Maccheroni, and H. Zank (2026). A separation of utility and beliefs through betting consistency. *Management Science* 72(3), 1987–2005.
- Debreu, G. (1959). *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. New York: John Wiley & Sons.
- Debreu, G. (1960). Topological methods in cardinal utility theory. In K. J. Arrow, S. Karlin, and P. Suppes (Eds.), *Mathematical Methods in the Social Sciences, 1959*, pp. 16–26. Stanford University Press.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141–153.
- Gul, F. (1992). Savage’s theorem with a finite number of states. *Journal of Economic Theory* 57(1), 99–110.
- Gul, F. and W. Pesendorfer (2014). Expected uncertain utility theory. *Econometrica* 82(1), 1–39.
- Hartmann, L. (2020). Savage’s P3 is redundant. *Econometrica* 88(1), 203–205.
- Hens, T. (1992). A note on Savage’s theorem with a finite number of states. *Journal of Risk and Uncertainty* 5(1), 63–71.
- Hirshleifer, J. (1965). Investment decision under uncertainty: Choice-theoretic approaches. *Quarterly Journal of Economics* 79(4), 509–536.
- Karni, E. (1985). *Decision Making Under Uncertainty: The Case of State-Dependent Preferences*. Cambridge, MA: Harvard University Press.
- Machina, M. J. and D. Schmeidler (1992). A more robust definition of subjective probability. *Econometrica* 60(4), 745–780.
- Nau, R. F. (2011). Risk, ambiguity, and state-preference theory. *Economic Theory* 48(2–3), 437–467.
- Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica* 32(1/2), 122–136.
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization* 3(4), 323–343.

- Rigotti, L., C. Shannon, and T. Strzalecki (2008). Subjective beliefs and ex ante trade. *Econometrica* 76(5), 1167–1190.
- Savage, L. J. (1954). *The Foundations of Statistics*. New York: John Wiley & Sons.
- Schervish, M. J., T. Seidenfeld, and J. B. Kadane (1990). State-dependent utilities. *Journal of the American Statistical Association* 85(411), 840–847.
- Segal, U. and A. Spivak (1990). First order versus second order risk aversion. *Journal of Economic Theory* 51(1), 111–125.
- Wakker, P. P. (1984). Cardinal coordinate independence for expected utility. *Journal of Mathematical Psychology* 28(1), 110–117.
- Wakker, P. P. (1989). *Additive Representations of Preferences: A New Foundation of Decision Analysis*. Dordrecht: Kluwer Academic Publishers.
- Werner, J. (2002, October). Risk aversion as a foundation of expected utility. Working paper, University of Minnesota.
- Werner, J. (2005). A simple axiomatization of risk-averse expected utility. *Economics Letters* 88(1), 73–77.
- Yaari, M. E. (1969). Some remarks on measures of risk aversion and on their uses. *Journal of Economic Theory* 1(3), 315–329.