What's in a u?*

Larbi Alaoui† Antonio Penta‡

April 13, 2026

Abstract

We revisit the long-lasting debate about the meaning of the utility function used in Expected Utility (EU). Contrary to the common view that EU inherently links risk aversion to diminishing marginal utility, we demonstrate that these two concepts need not coincide. Marginal utility of money is an input into risk attitude, but it is not its sole determinant. An independent 'pure risk' attitude is also a contributing factor. We discuss some theoretical implications of this result for the following topics: non-neutral risk attitudes for profit maximizing firms; risk aversion over time lotteries in the presence of discounting, and convex time budget decisions; the equity premium puzzle. We also discuss matters of identification: for firms; via *proxies*; via MLE methods under parametric restrictions; in intertemporal choice problems; and cross-context elicitation in multi-dimensional settings, and its relationship with the methods and results from Psychology.

Keywords: utility function – risk aversion – marginal utility
JEL Codes: C72; C91; C92; D80; D91.

*We are especially indebted to Agustin Troccoli Moretti, for many excellent comments and for valuable inputs. We also thank Jose Apesteguia, Han Bleichrodt, Tommaso Denti, David Dillenberger, Jan Eeckhout, Kirby Nielsen, Marcus Pivato, Antonio Rosato, Charlie Sprenger, Michela Tincani, Elias Tsakas, and Leeat Yariv for very helpful comments, and participants in several conferences and seminars.

The BSE acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S), Larbi Alaoui from the Spanish Ministry of Science and Innovation under project PGC2018-098949- B-I00, and Antonio Penta from the ERC Consolidator Grant #101089139 (ERC, Pers&Pref, n. 101089139). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

†Universitat Pompeu Fabra and BSE. E-mail: larbi.alaoui@upf.edu

‡ICREA, Universitat Pompeu Fabra and BSE. E-mail: antonio.penta@upf.edu

1 Introduction

The meaning of the (Bernoulli) utility function used in the standard Expected Utility representation has long been debated, and this debate has typically been viewed as one of preferred interpretation of the object. On one side, a classical and natural interpretation is associated with Bernoulli ([1738] 1954): the marginal utility captures a notion of the marginal value, i.e., of pleasure or satisfaction (in a real sense, 'utility') of the additional unit. This interpretation is commonly held by economists and psychologists, and is in line with the compelling intuition that the reason we prefer 1 billion dollars for sure to a 50-50 chance of getting 2 billion dollars or nothing is that the value of what we can buy with that first billion is significantly higher than what we can buy with the second. In other words, risk aversion in Expected Utility (EU) is driven by the diminishing marginal utility (value). On the other side, the orthodox decision theorists' viewpoint (e.g., Friedman and Savage, 1952) cautions against this interpretation. It takes the utility function as a representation of risk attitude, and nothing more.[1]

Regardless of one's view of these interpretations, it is generally accepted that the notions of marginal value (utility) and risk attitude are inexorably tied together in the Expected Utility representation. In fact, this view is so well-established that it has motivated the development of now widely-used alternatives to EU which aim to separate them. This is summarized in Yaari (1987)'s well-known quote:

> "In expected utility theory, the agent's attitude towards risk and the agent's attitude towards wealth are forever bonded together." [The reason this is problematic is that] "at the level of fundamental principles, risk aversion and diminishing marginal utility of wealth, which are synonymous under expected utility theory, are horses of different colors."

In this paper we focus on understanding how this debate can be resolved, and what can formally be said about the meaning of the utility function as an abstract concept. Our first result may appear striking: it is *not* the case that risk aversion and diminishing marginal value of wealth are pinned together in EU. While marginal value of wealth is certainly a factor in risk aversion, it is not its sole determinant.[2]

[1] This debate goes back to the 1950s. For a thorough historical perspective, see Moscati (2018, 2023). A nice account of the various notions of utility can be found in Bleichrodt (2025), Ch.2, and for an experimental focus see Camerer (2013).

[2] Following Yaari (1987), a large and important literature has departed from EU to model the

These theoretical insights emerge from a *thought experiment*, within which we can define and separate 'pure risk' attitude from the marginal utility of money under certainty.[3] In particular, the utility function u can be viewed as the composition of two functions, $u = g \circ f$, where the curvature of f represents the marginal value of wealth, while g is a CARA transformation, whose parameter represents the 'pure risk' attitude (Section 2). Thus, two individuals could have identical values of wealth, but different risk attitudes (in the standard sense of which gambles they accept), because they differ in how much they like or dislike risk, perhaps due to different personalities. By the same token, two individuals could have the same 'dislike' of risk, but value money in different ways, which in turn implies that they will have different risk attitudes over money. Hence, despite the well-known and common view that EU forces risk aversion and diminishing marginal utility of wealth to be pegged to one another, this is in fact not the case. Even within EU, horses of different colors are horses of different colors.

But it is not just that the commonly held interpretation is inaccurate and should be revised purely for the sake of interpretation. As we show in Section 3, once expected utility is viewed through the perspective of our approach, it also allows us to provide a solution to classical open questions in the literature, such as the willingness to gamble or buy insurance, even at high wealth levels (e.g., Friedman and Savage, 1948), as well as to accommodate, within an expected utility framework, non-neutral risk attitudes for profit-maximizing firms (another famous objective of the classical literature – see footnote 17 on Yaari (1987)). Furthermore, in a dynamic setting, our approach provides a foundation to a specific version of Kihlstrom and Mirman (1974)'s preferences, which enables us to accommodate, *within EU*, the behavioral patterns from the experiments of Andreoni and Sprenger (2012a,b), as well as to reconcile exponential discounting in the certainty space with risk aversion over time lotteries (DeJarnette et al., 2020). Finally, we show how the separation between diminishing marginal utility of money and 'pure risk' attitude has important implications for the analysis of saving and investment decisions, which may provide novel insights (from within expected utility) about Mehra and Prescott (1985)'s Equity Premium Puzzle.

distinction between certain and risk preferences (see, e.g., Abdellaoui et al. (2007); Bleichrodt (2025) and references therein), mostly by separating the different roles of the probability weighting function from that of the utility function. Within this non-EU literature, the closest work in spirit is probably Schmidt and Zank (2022). But this entire literature is very different from the focus of this paper, which shows that such a separation is possible even *within* EU.

[3]Issues of *actual* identification are postponed to Section 4, where we discuss a few methods and their connection with Tsakas (2025)'s solution to the problem of belief-identification with state-dependent utility.

The conceptual significance of the results of Sections 2 and 3 stands independent of the issue of separately identifying the two components of the utility function in practice. Nonetheless, identifying these functions is itself a relevant exercise, both for understanding preferences and for enabling sharper predictions. We turn to this exercise in Section 4, and discuss three methods of identification with the corresponding datasets. We first consider profit-maximizing firms, where the identification exercise follows naturally from a concrete implementation of our thought experiment. Then, we focus on intertemporal settings, and show how the two components of the utility function can be separately identified, under the assumption that preferences under certainty take the standard discounted utility form, using standard datasets, as those in Andreoni and Sprenger (2012a,b). We then consider individuals more generally, and provide the formal conditions on the observable dataset under which identification can be conducted via a *proxy* (i.e., a commodity that satisfies suitably defined properties, similar to Tsakas, 2025). We also suggest an example of a proxy that can be used in practice. Lastly, we discuss identification under specific parametric assumptions, of the kind that is most frequent in empirical and experimental work (cf. Gill and Prowse (2012)).

Finally, in Section 5 we show that the separation between 'pure risk' and marginal utility sheds light on important questions that arise in the context of risk when multiple commodities are involved.[4] In particular, we explain why risk attitude elicited in one domain may not correlate with those elicited in other domains, and show how our approach can solve this problem and lead to more accurate predictions *across* domains (Einav et al., 2012). We also show how eliciting risk attitudes over multiple domains may enable a better identification of the 'pure risk' parameter, leading to a stable ordering of individuals' risk attitudes that is portable *across* domains. In doing so, we also contribute to the discussion concerning the distinction between the standard choice-based economics methodology to elicit risk attitudes and the method often used in psychology, which involves questions that refer to a wide range of domains and forms of behavior (e.g., Frey et al., 2017). Our discussion formally connects the two approaches, and shows how they can be combined more effectively to understand and predict behavior.

The concluding section discusses how this work could be further developed, both theoretically and in combination with neuroeconomics research (see, for instance, Glimcher and Rustichini (2004); Camerer (2008, 2007); Rustichini (2009)

[4]For a careful analysis on risk attitude in multiple domains, see Ke and Zhang (2024), who introduce the *hierarchical expected utility* representation. We view this important contribution as complementary to that of our paper, where we remain strictly within standard expected utility.

for classic references, and Glimcher and Tymula (2023) for recent work), and various directions for future research.

2 Theoretical framework

In the standard approach to risk, within the von-Neumann-Morgenstern (vNM) framework, it is customary to consider preferences $\succsim^*$ over simple lotteries over the real line, $p, q \in \Delta(\mathbb{R})$ that satisfy the standard axioms (weak order, the Archimidean property and the independence axiom).[5] Letting $m \in \mathbb{R}$ denote the quantity of money, vNM's representation theorem ensures that there exists a utility function $u^* : \mathbb{R} \to \mathbb{R}$ such that $p \succsim^* q$ if and only if $\sum p(m)u^*(m) \geq \sum q(m)u^*(m)$, and u^* is unique up to positive affine transformations.

2.1 A Conceptual Yardstick

In the setting above, suppose that two onlookers (scientists) observed data on the agent's choices and concluded, for instance, that he is risk-averse, so that his utility is concave. But suppose that they disagree on the meaning of this finding: The first observer argues that the curvature of u^*, and hence the attitude towards risk, is merely a consequence of the agent's decreasing 'marginal utility for money'. The second onlooker instead argues that this concave function represents the agent's dislike of risk alone, and has nothing to do with his valuation of m. In particular, he argues that nothing can be concluded concerning what each extra unit of money yields, in utility terms. How then could these onlookers settle this argument?[6]

While they disagree on their interpretations of the model, however, both onlookers engage in this debate. In particular, they do not discard the notion of marginal utility of money as meaningless; they only differ in its connection with the vNM's u. But the notion of 'marginal utility of money' is inherently one of comparison: a comparison of money and some concept of its value, however the

[5]The vNM framework for decision under risk of course applies to general outcome spaces. This restriction on the outcome space, however, is common in the analysis of risk, and it is natural here. For classical papers on EU and risk in multidimensional settings, see Kihlstrom and Mirman (1974) and Karni (1979).

[6] These two archetypes illustrate two opposing views of a classical debate within the theory of individual decision making. The first onlooker, for instance, embodies Bernoulli's original theory (Bernoulli, [1738] 1954), which views risk-aversion as stemming from the diminishing marginal utility of riskless money. (Bernoulli did not use the term 'utility', but *emolumentum*, cf. Moscati (2023)). Friedman, in contrast, argued that the u^* in vNM's representation should best be referred to as 'choice-generating function', precisely to separate it from the notion of utility under certainty (Moscati, 2018). See also Friedman and Savage (1952). For a more recent iteration of this debate, see, for instance Rabin (2000), footnote 2.

latter is determined for the individual. That is, there must be at least an abstract *yardstick* of comparison, first for this debate to be well-defined, and second, as we will show, to settle it.[7]

Here we shall consider, as a pure thought experiment, which properties must such an idealized *yardstick* satisfy, and what this will imply for the interpretation of the utility function u^*. For the moment, we set aside questions of measurement and identification, engaging only in this conceptual thought experiment, with the aim of settling the debate at hand. This exercise only requires accepting the possibility that such a yardstick exists, and does not rely on its measurement. We will return to the issue of measurement and identification, which are of course important in their own right, in Section 4.

2.1.1 A Yardstick under Certainty

If such a yardstick existed, call it y, it would effectively be a second commodity, besides money, which could be traded off against m to identify the marginal rate of substitution (MRS) in the certainty space, and hence irrespectively from the agent's attitude towards risk. This MRS would effectively pin down the utility of the agent for each unit of m, in terms of the chosen unit of measure, the yardstick.

Having elicited such an MRS, for any $m \in \mathbb{R}$, we could identify the quantity of the y-commodity that would make the agent indifferent between receiving m and that specific quantity of y. Let $f : \mathbb{R} \to \mathbb{R}$ denote such a function, with $f(m)$ being the yardstick-equivalent of m elicited from these preferences.

Formally, in this setting, the agent's *preferences* are defined over a richer commodity space, which includes both m and y. That is, let the space of outcomes be $Z = \mathbb{R}^2$, with typical element $z = (m, y) \in Z$, and $\succsim^c$ denote a weak order, strictly increasing in both components. For simplicity, we also maintain throughout that $\succsim^c$ is continuous. Then, for a 'utility of money' $f(m)$ expressed *in terms of the yardstick*, and for y to effectively have this role, it must be the case that, up to increasing transformations, the representation of $(\succsim^c, Z)$ takes the form $f(m) + y$. That is: $(m, y) \succsim^c (m', y')$ if and only if $f(m) + y \geq f(m') + y'$. If not, commodity y would not be a *yardstick*, with respect to which the utility of m is expressed. While we view it as transparent to provide this property in terms of functions, it can clearly be expressed in terms of preferences directly, as mere quasilinearity:

[7] This exercise can of course be dismissed if one believes that the notion of marginal utility does not exist. It goes without saying that, this view notwithstanding, such a notion is ubiquitous in economics textbooks and papers. Beyond this, however, as we will show in Sections 3.1 and 3.2, even the 'agnostic' view of the EU representation has in fact taken a stance on the existence of marginal utility, and implicitly derived risk attitudes from preferences under certainty.

Definition 1 (Yardstick under Certainty)**.** *If the y-commodity is an ideal* yardstick under certainty*, then the preferences over the certain space $Z = \mathbb{R}^2$, $(\succsim^c, Z)$, are such that $(m, y) \succsim^c (m', y')$ if and only if $(m, y+t) \succsim^c (m', y'+t)$, where $t \in \mathbb{R}$.*

Remark 1. *If y is a* yardstick under certainty*, then there exists a strictly increasing $f : \mathbb{R} \to \mathbb{R}$ such that $(m, y) \succ^c (m', y')$ if and only if $f(m) + y > f(m') + y'$. The shape of such f, however, is unconstrained (e.g., it need not be concave).*

Note that even under this thought experiment, we cannot yet make any statements concerning the utility function u^* of the EU-representation in the baseline setting, because f need not correspond to this function. It only captures the value of m, in units of the yardstick y, absent risk considerations.

This is a property that we find appropriate for an *ideal* yardstick; nonetheless, in Section 2.2.4 we discuss how this property can be weakened to mere linear separability, without significantly changing our results.

2.1.2 From certainty to risk

Now, let us extend the domain of preferences to a risk setting, i.e. to account for simple lotteries $p, q \in \Delta(Z)$, maintaining all the standard vNM axioms, and let $(\succsim, \Delta(Z))$ denote the corresponding preference system. Again, we let δ_z denote the degenerate lottery over outcome $z = (m, y) \in Z$, and write $(p^m, y) \in \Delta(Z)$ to denote a lottery that is degenerate over the y-component and that randomizes over m according to $p^m \in \Delta(\mathbb{R})$.

First, to ensure that $(\succsim, \Delta(Z))$ naturally embeds the preferences over the certain domain, $(\succsim^c, Z)$, we maintain that for any (m, y) and (m', y'), it holds that $\delta_{(m,y)} \succsim \delta_{(m',y')}$ if and only if $(m, y) \succsim^c (m', y')$. Second, for commodity y to play the role of an ideal yardstick, *and nothing more*, there should be no interaction between the level of the yardstick and the risk preferences (i.e., preferences over lotteries over m.) This means that the preference system must satisfy the following *yardstick neutrality* condition: for any $y, y' \in \mathbb{R}$ and for any $p^m, q^m \in \Delta(\mathbb{R})$, $(p^m, y) \succsim (q^m, y)$ if and only if $(p^m, y') \succsim (q^m, y')$. In summary, besides the standard vNM axioms, preferences $(\succsim^c, \Delta(Z))$ must satisfy the following properties:

1. **A Yardstick under Certainty:** there exists a $f : \mathbb{R} \to \mathbb{R}$ that is strictly increasing and such that $\delta_{(m,y)} \succ \delta_{(m',y')}$ if and only if $f(m)+y > f(m')+y'$.

2. **Yardstick Neutrality:** for any $y, y' \in \mathbb{R}$ and for any $p^m, q^m \in \Delta(\mathbb{R})$, $(p^m, y) \succsim (q^m, y)$ if and only if $(p^m, y') \succsim (q^m, y')$.

The first property follows directly from Definition 1 and the fact that $\succsim$ is an extension of $\succsim^c$. As for Yardstick Neutrality, it ensures that the addition of the ideal y-commodity, per se, does not alter the set of preferences over monetary lotteries that could be expressed in the baseline (one-dimensional) setting $(\succsim^*, \Delta(\mathbb{R}_+))$. This is key to ultimately be able to draw a meaningful connection between the Bernoulli utility in the EU representation of $(\succsim, \Delta(Z))$, and the baseline u^* obtained in the standard one-dimensional setting. In fact, if Yardstick Neutrality did not hold, then it would mean that the baseline vNM preferences (and, hence, their representation) would be affected by an omitted variable problem, and their meaning would not be clear. Yardstick Neutrality therefore is both a natural and a necessary property to hold, within the standard vNM framework.

Proposition 1. *Under the maintained assumptions, $u : Z \to \mathbb{R}$ provides an EU representation of $(\succsim, \Delta(Z))$ if and only if there exists $g : \mathbb{R} \to \mathbb{R}$ such that $u(m, y) = g(f(m)+y)$ for all $(m, y) \in Z$, and which can only be one of the following functional forms (up to positive affine transformations): either (i) $g(x) = x$, or (ii) $g(x) = \frac{1-e^{-\alpha x}}{\alpha}$, with $\alpha \neq 0$.*

The logic of this result is simple: from the assumptions on the certain preferences, it follows that any strictly increasing transformation of $f(m)+y$ represents $(\succsim^c, Z)$. Extending these preferences to $\Delta(Z)$, with the vNM axioms, ensures that there exists a $g : \mathbb{R} \to \mathbb{R}$ (unique now up to positive *affine* transformations) such that $p \succsim q$ if and only if $\sum p(m, y)g(f(m) + y) \geq \sum q(m, y)g(f(m) + y)$. Then, yardstick neutrality further ensures that preferences over monetary lotteries are invariant to y, and hence $g(f(\cdot))$ is an affine transformation of $g(f(\cdot)+y)$ for each y. It follows that $g : \mathbb{R} \to \mathbb{R}$ exhibits constant absolute risk aversion (CARA).

Yardstick neutrality also allows us to define preferences over monetary lotteries alone, $(\succsim^m, \Delta(\mathbb{R}))$: for any y, say that $p^m \succsim_y q^m$ if and only if $(p^m, y) \succsim (q^m, y)$, and let $\succsim^m$ coincide with the $\succsim_y$-ordering for $y = 0$. (Under yardstick neutrality, $\succsim_y = \succsim_0$ for all y.) With this, we can relate the u obtained in Proposition 1, as part of the representation of the preferences $(\succsim, \Delta(Z))$, with the u^* obtained from the representation of the standard one-dimensional preferences, $(\succsim^*, \Delta(\mathbb{R}))$, when $\succsim^m = \succsim^*$. Formally:

Proposition 2. *$u^* : \mathbb{R} \to \mathbb{R}$ is a (strictly increasing) Bernoulli utility associated with the EU-representation of some preference system $(\succsim^*, \Delta(\mathbb{R}))$ if and only if there exists $(\succsim, \Delta(Z))$ that satisfies the maintained axioms, with $\succsim^m = \succsim^*$, and a utility $u : \mathbb{R}^2 \to \mathbb{R}$ that represents it in the sense of Proposition 1, such that* *(i)* $u^*(m) = u(m, 0) = g(f(m))$ *for all* $m \in \mathbb{R}$ *and (ii)* $u^*(\cdot)$ *is a positive affine transformation of* $u(\cdot, y)$ *for all* y*.*

The first part of Proposition 2 says that, under the maintained assumptions, (including, in particular, yardstick neutrality), the set of Bernoulli utility functions generated in the two-dimensional setting is exactly the same as that generated in the standard (one-dimensional) setting for preferences over risk: The addition of the yardstick neither restricts nor enlarges the set of possible Bernoulli utilities. The second part states the implication of yardstick neutrality discussed above, now in the space of the representations: the vNM preferences over monetary lotteries are invariant to y, and hence if u represents $(\succsim, \Delta(Z))$ in the sense of Prop. 1, then $u(m, 0)$ (and, thus, $u^*(m)$) is an affine transformation of $u(m, y)$ for any y.

In light of this result, in the following we will simply write $u(m)$, to refer interchangeably to either $u^*(m)$ or $u(m, 0)$, and refer to it as the Bernoulli utility function, which by Proposition 1 is understood to be of the form $u = g \circ f$, where g is a CARA transformation. Then, we write $g(x) = \frac{1-e^{-\alpha x}}{\alpha}$ if $\alpha \neq 0$, and $g(x) = e^{\alpha}x$ otherwise, so that $\alpha \in \mathbb{R}$ denotes the (constant) coefficient of absolute risk aversion of g, with $\alpha = 0$ for risk-neutrality.

Furthermore, for the case where f and u are twice differentiable (given the CARA property of g, f is differentiable if and only if u is), let $\alpha_u(m) = -\frac{u''(m)}{u'(m)}$ and $\alpha_f(m) = -\frac{f''(m)}{f'(m)}$ denote the Arrow-Pratt coefficients of absolute risk-aversion at m, for u and f, respectively.[8]

Corollary 1. *Under the maintained assumptions, if* u *is a Bernoulli utility function of the EU-representation, and if it is differentiable, then, for each* $m \in \mathbb{R}$*,*

$$\alpha_u(m) = \alpha \cdot f'(m) + \alpha_f(m). \tag{1}$$

Our exercise and results may appear reminiscent of Epstein and Zin (1989), in that, while the functional form we obtain is of course clearly distinct, it uses an enriched space. But note that our exercise is entirely within EU, rather than being within a recursive utility setting (Kreps and Porteus, 1978). Moreover, while Epstein and Zin (1989) focus on the separation between risk attitude and intertemporal preferences, we study instead the connection between risk attitude

[8]From a formal viewpoint, the argument above involves an 'extended domain', which also includes the yardstick. In multidimensional settings, other notions of risk aversion may also be relevant. But here we use this extended domain to make a conceptual point on the baseline one-dimensional setting, which remains our main object of interest. In that domain, risk aversion retains its standard meaning, and that is why we focus on the standard Arrow-Pratt indices.

and the utility over the good (money). For this reason, our enrichment is of a different nature from the Epstein-Zin exercise.

2.2 Discussion and Extensions

We discuss next the interpretation and some implications of our results, before turning to a discussion of some relaxations of the maintained assumptions.

2.2.1 On the Interpretation of f and g

By construction, the f function represents the value of money, m, in units of the yardstick y, as elicited in the *certainty domain*, i.e. purely on the basis of preferences $(\succsim^c, Z)$. The function g instead, is pinned down by the agent's preferences over lotteries, and hence it is purely about risk. The u in the EU-representation, is the composition of these two functions: $u = g \circ f$. As usual, the decision maker's *risk attitude over money* is expressed by the curvature of the u function, which thus depends on the curvatures of two functions: 'value of money', f, and 'pure risk function' g.[9] Corollary 1 formalizes this idea, providing the decomposition of the Arrow-Pratt index of risk aversion of u in terms of the indices for g and f, respectively. The former is constant in m, due to the result that g is CARA, and we will refer to it as the *'pure risk' parameter.*

With this, *risk aversion* coincides with *diminishing marginal utility for money* if and only if $\alpha = 0$. In this case, the agent is neutral about risk, per se, and his overall risk-attitude (the curvature of u) is entirely due to his value for money (the curvature of f).[10] At the opposite extreme, even if f is linear (as it would, for instance, in the important special case of a profit-maximizing firm, which we discuss in the next section), the agent could still exhibit risk aversion (or be risk seeking), depending on the value of α.

[9]In this sense, our results provide a foundation to the introspection arguments in Bell and Raiffa (1982) and Dyer and Sarin (1982), who also argue for envisioning u as a composition of two functions, to separate 'strength of preferences' from 'intrinsic risk attitude.' An analogous composition is obtained by Aumann (2021), albeit with a different interpretation and through a distinct notion of risk-aversion. An example in Bell and Raiffa (1982) makes the case for the exponential form for the outer function, for which Smidts (1997) finds empirical support in a field experiment. We find it encouraging that this early literature supports our results. Recent developments of those early insights, withinsocial choice theory, can be found in Dietrich (2025), whereas Gerasimou (2021, 2022) revisits the issue of preference intensity comparisons.

[10]This case corresponds to Bernoulli ([1738] 1954)'s original theory (also the view of our first onlooker), according to which risk-aversion arises *because* of decreasing marginal utility. The second onlooker, in contrast, is agnostic and does not care to explain *why* risk-aversion arises: u is just a tool to make predictions, as in Friedman's 'choice-generating function' (see footnote 6).

More broadly, if g is not linear, then there is a wedge between the curvature of u and that of f. Take, for instance, an agent with an everywhere decreasing marginal utility for money, who therefore exhibits a globally concave f. This agent may still be risk-loving overall, if his 'pure risk' attitude is such that g is sufficiently convex. Alternatively, take another agent, with exactly the same preferences under certainty (and, hence, the same f), but who is averse to 'pure' risk. Then, his u would be more concave than what dictated by the concavity of f alone, since it would reflect both the curvature due to the decreasing marginal utility of money, and that of g, which only reflects the agent's dislike for risk, net of his preferences over money.

In summary, *diminishing marginal utility for money* (i.e., concavity of f) is a *reason for* risk-aversion, but it need not suffice for it, nor need it be its sole determinant. The attitude towards *pure risk*, as represented by the 'pure risk' parameter α, also contributes. Thus, despite the well-known and common view that EU forces risk aversion and diminishing marginal utility of wealth to be pegged to one another, this is in fact not the case. Even within EU, horses of different colors are horses of different colors.

2.2.2 On a richer notion of EU, but still EU

As discussed, our results suggest that the widespread view, which identifies the curvature of the utility function under certainty with overall risk aversion, is overly restrictive: the two notions coexist, within the u, but they remain conceptually distinct, even within EU. As we will explain in the following sections, such a richer notion of EU enables us to reconcile this framework with several important open questions. Nonetheless, we have deliberately developed our argument so as to remain within EU. It follows that our result cannot explain violations of that framework, such as Allais, Ellsberg, or Rabin's paradoxes.[11] Of course, this does not diminish the importance of those results, nor of the enormous literature on violations of independence or of other components of EU. We only aim to reassess what we understand about this central framework in economics, and explore some important implications of this newly found richness.

[11]The distinction that we identify between f and g may lead to novel ways of departing from EU and to new insights into such paradoxes. Some of these are explored in ongoing work, but they remain beyond the scope of this paper, which maintains the EU axioms throughout.

2.2.3 From u to f, and the St. Petersburg Paradox revisited

While questions of identification are left to Section 4, here we make the following observations concerning partial inferences that can be made. Since the function g is of CARA form, one of the following three must hold: i) it is bounded above if $\alpha > 0$, ii) it is bounded below if $\alpha < 0$, or iii) it is unbounded if $\alpha = 0$. This means that if u is assumed to have a specific parametric form that is unbounded on one side or another, than we can make partial inferences on the α parameter. Specifically: if $\alpha > 0$, then u would have to be bounded above, and if $\alpha < 0$, then u would have to be bounded below. Hence, only a linear g can be consistent with unboundedness both above and below.

Let us now return to the famous St. Petersburg Paradox, which initially led Bernoulli to the Expected Utility formulation in the first place.[12] The paradox is the following: suppose a fair coin is flipped, and if it lands heads then the game ends and the agent receives 2 ducats (or euros), otherwise it is flipped again and if it lands heads than the agent receives 4 ducats, and so forth. It has infinite expected value, but most individuals would not be willing to pay an infinite amount for it. In proposing an EU formulation, Bernoulli further proposed using the functional form $u(x) = log(x)$. However, since it is unbounded, it has itself been criticized for allowing the emergence of a variation of the same paradox, in which the agent would still be willing to pay an infinite amount. The proposed solution was precisely to resort to utility functions that are bounded above (CRRA functions with a risk attitude higher than that of the log function, as commonly used in practice, satisfy this property).

From the viewpoint of our approach, the reason why the (adjusted) St. Petersburg Paradox would still be an issue with an unbounded u is precisely that it forces g to be linear, and hence it does *not* allow for any aversion with respect to 'pure risk', and it forces f to be log. The proposed solution of a bounded above u, such as a CRRA function with a high enough risk attitude, allows for $\alpha \geq 0$, and hence allows for aversion to pure risk. It also allows for f to be bounded above, which the log u does not.

[12]A version of the paradox was initially put forward by Nicolas Bernoulli, with the version that became 'definitive' being proposed by Gabriel Cramer. The EU resolution was proposed by Daniel Bernoulli, who favored the log utility representation. (Cramer's favored square root representation suffers from the same unboundedness issue as the log function). See Moscati (2023) for a detailed discussion.

2.2.4 On the Relaxation of Quasilinearity

The first property, quasilinearity, is the one we find appropriate to demonstrate the point cleanly, since one can clearly see how f maps to units of an *ideal* yardstick, y. However, suppose y were regarded as an extra (actual) commodity, or a yardstick with a non-linear scale. Then, if quasilinearity were weakened to linear separability, with $(\succsim^c, Z)$ being (ordinally) represented by $f(m) + h(y)$ for some strictly increasing h, then the exercise would still go through as is. Intuitively, the exercise relies on the unit of measurement (the yardstick) as being conceptually separated from what it measures (the utility of money). It is simpler to understand this separation using a linear scale than a non-linear one, hence by using the form $f(m) + y$ rather than $f(m) + h(y)$. This is because with a non-linear scale the MRS would be $f'(m)/h'(y)$ rather than $f'(m)$. Moreover, one could simply transform y into another unit to recover quasilinearity. We therefore maintain quasilinearity of the yardstick for the rest of the paper, noting only that linearity of that good is not in itself crucial to our analysis.

2.2.5 On Uniqueness

It is worth noting that, while vNM's u is of course unique up to positive affine transformations, and so is the g in the representation in Propositions 1 and 2, the f is instead *unique*, and hence it is *cardinal* in the classical, Pre-Samuelsonian sense of the word (as opposed to the notion that has become standard since von Neumann and Morgenstern (1954); cf. Moscati, 2018). This is natural, since f represents the marginal utility of money in terms of the yardstick: as we will further clarify in Section 4, the choice of the yardstick pins down the units in which the MRS is expressed, and given such a choice, the marginal utility of money is cardinal in the classical sense. In contrast, the 'pure risk' attitude is cardinal in the sense of being unique up to positive affine transformations.[13]

[13] In his critical account of the history of expected utility theory, Moscati (2018, 2023) distinguishes the classical notion of cardinality from von Neumann and Morgenstern's, which he refers to as 'scale-invariance'. It is interesting that the uniqueness(-es) in our result reflect, in a formal sense, the two main historical perspectives: the latter notion applies to the g function, which captures the 'pure risk' attitude, whereas the f function, which in a sense formalizes Bernoulli's rationale for risk-aversion, is cardinal in the sense of the classics. See also Bleichrodt (2025), Ch.2, for a similar account of the various cardinality notions.

2.2.6 Multiple commodities

The logic above clearly extends to multi-dimensional settings, i.e. where preferences in the certainty space are defined over multiple commodities.[14] Specifically, let $X \subseteq \mathbb{R}^n$ denote the commodity space with typical element $x = (x_1, ..., x_n)$, and $(\succsim^c, X)$ denote the preferences in the certainty space. These preferences can be identified in the standard way, with a utility representation unique up to increasing transformations. Then, for the *ideal yardstick* y (which, as above, should satisfy its key defining property: here $(x, y) \succsim^c (x', y')$ if and only if $(x, y+t) \succsim^c (x', y'+t)$), the MRS between the utility index and such y would pin down the function $f : X \to \mathbb{R}$, i.e. the utility under certainty. Then, under yardstick neutrality (which is formulated exactly as above, provided that Z is redefined as $Z = X \times \mathbb{R}$, one obtains results analogous to Propositions 1 and 2, so that the overall $u : X \to \mathbb{R}$ in the vNM representation can be written as $u(x) = g(f(x))$, where $g : \mathbb{R} \to \mathbb{R}$ is a CARA transformation with parameter $\alpha \in \mathbb{R}$.

3 So What Else?

The results above shed light on several important open questions, of both conceptual and practical significance. At a minimum, they enable us to reconcile seemingly contrasting views that are based on insights over different domain of preferences. For instance, it is typically argued that the marginal utility of money should become constant at sufficiently high levels of wealth. Yet, even billionaires buy insurance, or others perhaps gamble instead, in ways that would be inconsistent with a Bernoulli utility function that is approximately linear (e.g., Friedman and Savage, 1948). As discussed, these two views can be reconciled within our framework: depending on the α parameter, the agent could still exhibit risk aversion (or risk seeking-ness), even if the value of wealth were perfectly linear. Hence, insights based on observations made on the certainty space, need not coincide with, nor translate to, unique implications in the risky domain. There are two conceptually distinct forces at play.

In this section we discuss a few further implications of these general observations, specifically for three important problems: first, we discuss the possibility of accommodating, within the expected utility framework, non-neutral risk-attitudes for profit-maximizing firms; second, in the context of intertemporal choice, the

[14] In order to keep the focus on the distinction between ‘pure risk’ and certainty preferences, we do not dwell here upon questions of risk attitudes in multi-dimensional settings (and the possible departures from EU that may ensue), which have been recently studied in Ke and Zhang (2024).

results above support a specific version of Kihlstrom and Mirman (1974)'s preferences, which enable us both to reconcile exponential discounting under certainty with risk aversion over time lotteries (DeJarnette et al., 2020), and to accommodate, *within EU*, the behavioral patterns from the experiments of Andreoni and Sprenger (2012a,b); finally, we discuss some implications for saving and investment decisions, and for the *equity premium puzzle* (Mehra and Prescott, 1985).

3.1 Firms and Risk

Consider first the case where the agent is a profit-maximizing firm. The thought exercise of the previous section becomes simple and very concrete in this scenario: the unit of account (say, US dollars for a US firm) *is* the yardstick. Hence, if m denotes its money holdings in USD, the MRS between m and y under certainty in this case is the identity function, $f(m) = m$. It is both standard and uncontroversial that these are the (neoclassical) firm's preferences under certainty.[15]

It has also been customary to jump from this uncontroversial observation about the firm's preferences under certainty, to the conclusion that a profit-maximizing firm under expected utility must necessarily be risk-neutral.[16] While this confusion seems natural from the viewpoint of our first onlooker (according to which risk aversion stems from a decreasing marginal utility of money), it strikes as peculiar from the viewpoint of our second onlooker – according to which the u in the vNM representation has no connection with the underlying preferences under certainty (cf. footnotes 6 and 10) – and which embodies a widely held decision-theoretic view at least since Friedman and Savage (1952). For the case of firms, the common practice clearly departs from the agnostic view of a mere index to represent risk attitude: for firms, a stronger position has been taken that the latter can be derived from the preferences under certainty. In fact, this view is so deep-rooted that it has led economists to accept a model which is clearly at odds with the evidence that several firms do buy insurance.[17]

[15] The finance literature has sometimes departed from this standard approach, for instance by taking f to be a quadratic function, as a way to provide foundations to mean-variance analysis, which is frequently used in practice (see Eeckhout and Veldkamp (2023) for a recent example). These preferences are typically not considered as standard, since they violate for instance monotonicity.

[16] A distinct argument for why firms should act as expected profit maximizers is that they have access to a perfectly competitive (i.e., actuarially fair) insurance market, where they can fully insure against *all* risk. This argument is obviously limited from a positive perspective (insurance markets are neither complete nor perfectly competitive). From a theoretical perspective, one can qualify the points we are making as being relevant whenever there exists some residual risk for which no actuarially fair insurance is available.

[17] Yaari (1987), for instance, took it as an extra motivation for venturing outside of expected

To be clear, this $g(\cdot)$ transformation should be applied to the standard profit function (say, for a monopolist, $\pi(q) = P(q) \cdot q - C(q)$, where q denotes the quantity, and $P(\cdot)$ and $C(\cdot)$ the inverse demand and cost functions, respectively). In the face of uncertainty, the firm here would maximize $\mathbb{E}[g(\pi)]$, not $\mathbb{E}[\pi]$. Since $g(\cdot)$ is an increasing transformation, this change in the objective function has no bearing on choices under certainty, but if $\alpha \neq 0$, it does affect choices under uncertainty.

As an example, consider a monopolist facing a stochastic demand, where the price associated with output q is $P(q)+\epsilon$ or $P(q)-\epsilon$, with equal probability. In this case, the optimal choice for a risk-neutral firm would be the same, independent of the magnitude of $\epsilon \geq 0$, and hence it would coincide with the optimal quantity under certainty, q^*. That a 'real world' monopolist would be indifferent over any magnitude of $\epsilon \geq 0$, and that it would not react to it, seems contrary to common sense (cf. footnote 17). But, as explained, a profit-maximizing firm need not be risk-neutral. For any 'pure risk' parameter $\alpha \neq 0$, the optimal choice in this setting would respond to the magnitude of the demand shocks.[18]

Overall, this discussion makes a few points. First and foremost is the observation that a profit-maximizing firm need not be risk-neutral, even within expected utility. The second is that in this context the yardstick (dollars) is natural and easily observed. Hence, the pure risk parameter, α, is directly identified here from standard choice data under risk. (We will return to this point in Section 4.) Lastly, if it is true that a profit-maximizing firm need not be risk neutral, it is also true that its risk attitude can only take the CARA form. Hence, while for individuals a CARA utility function is often viewed as unrealistic (if perhaps convenient in applications), a strong (and falsifiable) implication of our model is that here it is the *only* appropriate functional form for profit-maximizing firms, under EU.[19]

utility, with his 'dual theory' of choice under risk: "Under the dual theory, maximization of a linear function of profits can be entertained simultaneously with risk aversion. How often has the desire to retain profit maximization led to contrived arguments about firms' risk neutrality?" But, as discussed above, a profit-maximizing firm must also be risk-neutral only if the g function is linear. In general, even a profit-maximizing firm may have a pure risk parameter $\alpha \neq 0$, in which case its optimization problem in the risky space is captured by $g(m) = \frac{1-e^{-\alpha m}}{\alpha}$.

[18] For example, with non-decreasing marginal costs, the optimal choice $q^*(\alpha, \epsilon)$ would be such that $q^*(0, \epsilon) = q^*$ under risk neutrality, but for any $\alpha > 0$ (resp., $\alpha < 0$) $q^*(\alpha, \epsilon)$ would be decreasing (resp., increasing) in ϵ, and such that $q^*(\alpha, 0) = q^*$ for any α. Moreover, for any $\epsilon > 0$, $q^*(\alpha, \epsilon)$ is strictly decreasing in α, and $lim_{\alpha \to \infty} q^*(\alpha, \epsilon) = q^-$ and $lim_{\alpha \to -\infty} q^*(\alpha, \epsilon) = q^+$, where q^- and q^+ denote the optimal choices conditional on the realization of $-\epsilon$ and $+\epsilon$, respectively.

[19] Since $u = g \circ f$, g being CARA in general does not imply that the overall utility is CARA, unless when f is linear, as for the case of firms. But if one is willing to assume that f is not linear (as in the literature in footnote 15), then u may take many forms, even if g is CARA. These include the CRRA form, which is often convenient in applications.

3.2 Intertemporal Choice Problems

Another setting where the results in Section 2 have immediate implications is given by intertemporal choice problems. For instance, let the multiple commodities in Section 2.2.6 consist of consumption levels in different periods. That is, the commodity space is $X = \mathbb{R}^{\mathbb{N}}_{+}$, with typical element $x = (c_1, c_2, ...)$, where c_t denotes consumption in period t. In this setting, it is standard to assume the following specification for preferences under certainty: for any $x = (c_1, c_2, ...)$, $f(x) = \sum_{t=1}^{\infty} \beta^{t-1} v(c_t)$, where $v(\cdot)$ is the within-period utility function, and $\beta \in (0, 1)$ the discount factor. Then, taking this function to represent the agent's utility (in *yardstick* terms, and thus imposing a cardinal assumption) from consumption streams under certainty, our general representation $u = g \circ f$ in this case takes the following form: for any $x = (c_1, c_2, ...) \in X$,

$$u(x) = g\left(\sum_{t=1}^{\infty} \beta^{t-1} v(c_t)\right), \tag{2}$$

where $g(\cdot)$ is a CARA transformation. The standard expected discounted utility (EDU) model, where risk-attitudes are entirely driven by the curvature of v, once again obtains for the special case where g is linear. We consider two applications within these settings: Andreoni and Sprenger (2012a,b)'s choices under convex budget sets, and risk attitudes over time lotteries (e.g., (DeJarnette et al., 2020)).

3.2.1 Time Lotteries

As a number of recent influential papers have pointed out (e.g., DeJarnette et al. (2020); see also Strzalecki (2024), and references therein), the standard model has somewhat disappointing implications when it comes to risk-attitude over *time-lotteries* (i.e., lotteries that pay a fixed prize at a random time). Intuitively, due to the convexity of the exponential discounting with respect to t, the standard model implies that the agent must be risk-seeking over time lotteries, which clashes both with our introspection and with a substantial body of experimental evidence.

If the 'pure risk' parameter of the $g(\cdot)$ transformation is sufficiently high, however, it is easy to see that one can retain the exponential discounting in the certainty space (for instance, to maintain both the tractability and cogency of the standard EDU model in terms of saving decisions, consumption smoothing, etc.), while at the same time accommodating risk-aversion over time lotteries. The functional form in eq. (2) is known as the Kihlstrom and Mirman (1974)'s representation, with the added qualification that the g transformation must be

CARA.[20] Interestingly, with this added property the model is immune to one of the main criticisms of the general KM model, which is that period-t attitudes toward future risk depend on bygone consumption levels (Epstein, 1992).

This CARA property imposes more structure on the preferences, for instance by entailing a specific relationship between the coefficient of risk-aversions over *money-lotteries* (i.e., lotteries with random prize paid at a fixed time), across different time periods. Specifically, letting $\alpha_v(m_t) = -\frac{v''(m_t)}{v'(m_t)}$, and letting $\alpha_t(m_t)$ denote the coefficient of absolute risk-aversion over money-lotteries that pay out at time t, evaluated at period 1, when the period-t money holding is m_t, we have:

$$\alpha_t(m_t) = \alpha \cdot \beta^{t-1} v'(m_t) + \alpha_v(m_t). \tag{3}$$

Thus, an interesting implication of this model, which could be immediately obtained from this equation, is that if both g and v are concave, then the agent becomes progressively less risk-averse as the horizon of the money-lotteries is postponed (the opposite is true if v is concave and g convex).

3.2.2 The Andreoni-Sprenger Convex Budgets Choice Tasks

In two influential papers, Andreoni and Sprenger (2012a,b) consider agents facing a set of choices of the following nature:

Case 1: A budget to allocate between consumption at time t and at time $t + k$ (respectively, $c_t^{(1)}$ and $c_{t+k}^{(1)}$), at a given interest rate r, when the payment is certain (i.e., with probabilities $p_t^{(1)} = p_{t+k}^{(1)} = 1$).

Case 2: A budget to allocate between consumption at time t and at $t+k$ (respectively, $c_t^{(2)}$ and $c_{t+k}^{(2)}$), at a given interest rate r, when the payment in each period is made with (independent) probability equal to $p_t^{(2)} = p_{t+k}^{(2)} = 0.5$)

As discussed in Andreoni and Sprenger (2012b), while it is intuitive to expect different behavior in Cases 1 and 2 above, the standard EDU model cannot explain

[20] The KM-representation was recently given a novel foundation by Dillenberger et al. (forth.), precisely as the answer to the problem of accommodating risk aversion over time-lotteries. In this sense, our results applied to this setting provide an alternative ‘foundation’ to their arguments, as well as the extra restriction to the g. See also Apesteguia et al. (2019) for a discussion of this functional form, and Proposition 5 in particular, for a family of models separating risk and time in different ways. Finally, as discussed above (p.10), while the connection with Epstein and Zin (1989) is tempting, the two models are distinct, as our analysis is completely within EU, rather than a generalization of recursive utility. In fact, as shown by DeJarnette et al. (2020), risk-aversion over time lotteries (which, as discussed, can be accommodated within our model) is incompatible with the Epstein-Zin preferences.

a difference in choices of budget allocations between them: under EDU, the choices must be the same whenever the ratio of probabilities is the same (as is the case for Cases 1 and 2 above, since $p_t^{(1)}/p_{t+k}^{(1)} = p_t^{(2)}/p_{t+k}^{(2)} = 1$). In fact, Andreoni and Sprenger (2012b) further discuss that well-known alternative models, cannot explain these patterns either. They are, however, fully in line with the model here, because the time preferences in Case 1 do not pin down the preferences in Case 2, which depend not only on the v and the β, but also on the risk parameter α.

Andreoni and Sprenger (2012b) also present other cases, namely:

Case 3: A budget to allocate between consumption at time t and at time $t+k$ (respectively, $c_t^{(3)}$ and $c_{t+k}^{(3)}$), at a given interest rate r, when the payment is certain in period t and occurs with probability 0.8 in period $t+k$ (i.e., $p_t^{(3)} = 1$ and $p_{t+k}^{(3)} = 0.8$).

Case 4: The same as Case 3, but with halved probabilities (i.e., $p_t^{(4)} = 0.5$ and $p_{t+k}^{(4)} = 0.4$, with corresponding consumptions $c_t^{(4)}$ and $c_{t+k}^{(4)}$)

Case 5: A budget to allocate between consumption at time t and time $t+k$ (respectively, $c_t^{(5)}$ and $c_{t+k}^{(5)}$), at a given interest rate r, when the payments occur with probability 0.8 in period t and 1 in period $t+k$ ($p_t^{(5)} = 0.8$ and $p_{t+k}^{(5)} = 1$, with corresponding consumptions $c_t^{(5)}$ and $c_{t+k}^{(5)}$).

Case 6: The same as Case 5, but with halved probabilities (i.e., $p_t^{(6)} = 0.4$ and $p_{t+k}^{(6)} = 0.5$, with corresponding consumptions $c_t^{(6)}$ and $c_{t+k}^{(6)}$)

Here as well, Cases 3 and 4 should lead to the same choices in the standard model, since they have the same ratio of probabilities (similarly for Cases 5 and 6). In Andreoni and Sprenger (2012a)'s data, however, this is not observed: in Case 3, subjects place higher weight on the earlier, certain consumption compared to Case 4, and in Case 6 they place higher weight on the later, certain consumption compared to Case 6. These patterns, however, are fully in line with our framework, and can be jointly accommodated under the same set of parameters (the technical details can be provided upon request).

In summary, in intertemporal choice problems, our results can *jointly* accommodate, within expected utility, both risk-aversion over time lotteries (e.g., DeJarnette et al. (2020) and Andreoni and Sprenger (2012b)'s choice patters.

3.3 Investment and Savings

In this section we discuss some implications of the above findings for saving and investment choices, and we show that they also play a role in another classical economics question, namely the Equity Premium Puzzle (Mehra and Prescott, 1985). For simplicity, let us consider the intertemporal setting above, but with two periods only, 'today' and 'tomorrow', with corresponding consumption levels denoted by c_t and c_{t+1}. Thus, preferences in the certainty space are such that $f(x) = v(c_t) + \beta v(c_{t+1})$, and as standard within this branch of the literature we also assume that the within-period utility function, $v : \mathbb{R}_+ \to \mathbb{R}$, is CRRA with parameter γ. Then, the overall preferences in the risky domain are

$$\mathbb{E}[u(c_t, c_{t+1})] = \mathbb{E}\Big[g\Big(v(c_t) + \beta v(c_{t+1})\Big)\Big], \tag{4}$$

where v is CRRA and g is CARA:

$$v(c) = \frac{c^{1-\gamma}}{1-\gamma} \text{ with } \gamma > 0; \text{ and } g(z) = \frac{1 - e^{-\alpha z}}{\alpha} \text{ if } \alpha \neq 0, \text{ and } g(z) = z \text{ if } \alpha = 0.$$

The standard model again obtains with $\alpha = 0$. There, the γ parameter is the sole determinant of both the agents' savings under certainty, i.e. their 'consumption smoothing' motives, and their attitudes towards risk, and hence their investment decisions over risky asset. Standard estimations of the CRRA parameter, based on typical preferences over (within period) money lotteries, normally yield values in the range of $[1, 4]$. As is well-known, values of γ in this range entail substantially lower risk-premia for risky investments than are observed empirically. This is the celebrated Equity Premium Puzzle (Mehra and Prescott, 1985).

As we explain next, if $\alpha \neq 0$, the expected utility model may account for larger risk-premia, even holding constant the agents' risk aversion over (within period) money lotteries, The intuition is that the g function plays no role when considering risk-free saving decisions, since it cancels out from the Euler equation associated with the optimal saving problem under certainty. The agents' consumption smoothing motive therefore is solely driven by the curvature of the v function, i.e. by the γ parameter. Preferences over risky assets, in contrast, depend on *both* γ and α. Hence, in the determination of the 'equity premium', the latter parameter only matters for the returns of the risky assets; the returns of risk-free assets are only linked to γ.[21] This effect, however, has bite provided that

[21] By continuity, the same logic also applies to the case where one considers the risk-premium compared to a low-risk (though not completely riskless) asset. Intuitively, the role of the g

risk follows a *rare disasters* distribution, with a 'thin' but sufficiently long left tail, thereby connecting the ability of our approach to account for a larger share of the Equity Premium Puzzle, with the shape of the risk distribution.

We illustrate this point within a simple example (for a general analysis, see Alaoui, Penta, and Troccoli-Moretti, 2026). In particular, consider an agent choosing the optimal level of investment, first with a risk-free bond, and then with a risky asset. Formally, letting y denote current income, the agent solves:

$$\begin{aligned}\max_{s\in[0,y]} \mathbb{E}\Big[g\Big(v(c_t)+\beta v(c_{t+1})\Big)\Big]\\ \text{subj.to: } c_t = y-s\\ c_{t+1}=(1+R)s\end{aligned}$$

In the first setting, the agent invests in a *risk-free asset*, that pays a fixed return $R = R_f > 0$, which we assume to be at the so called *natural* level ($\beta = 1/(1+R_f)$), so that consumption is stable across periods). In the second setting, the agent invests in a *risky asset*, that pays a random interest rate $R = \tilde{R}$ where

$$\widetilde{R} = \begin{cases} \bar{R}+\delta & \text{with Prob. } \dfrac{\lambda\varepsilon}{1+\lambda}, \\ \bar{R} & \text{with Prob. } 1-\varepsilon, \\ \bar{R}-\lambda\delta & \text{with Prob. } \dfrac{\varepsilon}{1+\lambda}, \end{cases} \qquad \text{where } \varepsilon,\delta \geq 0 \text{ and } \lambda \geq 1.$$

First consider the case where $\lambda = 1$. Then, the risk distribution is symmetric around the mean, and for a fixed $\delta > 0$, increments in risk are parametrized by the value of ε. Also, for any $\lambda \geq 1$, increments in ε keep the mean $\mathbb{E}[\widetilde{R}] = \bar{R}$ unchanged, while shifting probability mass from the center to the tails. Given $\varepsilon > 0$ and $\delta > 0$, instead, increasing λ makes the bad state both *more severe* and *less likely*, again keeping the mean return constant. In this sense, while increasing each of these parameters induces a mean-preserving spread, λ in this example parametrizes progressively more extreme *rare-disaster risk*.

Letting $c_f = (c_t^f, c_{t+1}^f)$ and $c_r = (c_t^r, c_{t+1}^r)$ denote, respectively, the optimal consumption in the risk-free and risky asset problems, we let $R^*(\alpha,\gamma,\varepsilon,\delta,\lambda)$ denote the value of the mean return of the risky asset, $\bar{R}$, for which $\mathbb{E}[u(c^f)] = \mathbb{E}[u(c^r)]$. Within this example, one could thus define the 'equity premium' as:

$$EP(\alpha,\gamma,\varepsilon,\delta,\lambda) := R^*(\alpha,\gamma,\varepsilon,\delta,\lambda) - R_f. \tag{5}$$

function becomes smaller, and eventually vanishes, as we approach the risk-free benchmark.

Consistent with the insights from Kihlstrom and Mirman (1974), it is easy to verify that this 'equity premium' is increasing in α.[22] But this is not the relevant exercise for the Equity Premium Puzzle, since increasing α would also increase the risk attitude over (within period) money lotteries. To properly address the Equity Premium Puzzle, one must instead increase α without changing the overall risk-attitude over money lotteries.

Applying the formula for the AP-index of risk aversion in (8) for the current consumption, c_t, evaluated at the optimal level in the risk-free setting, which we normalize to one (i.e., $c_t^f \equiv 1$), we obtain that the overall coefficient of risk aversion in this case is $AP_t(c_t^f) = \alpha + \gamma$. Thus, to vary the curvatures of the g and v functions, while keeping the overall risk-attitude over (within period) money lotteries constant, it suffices to impose that the sum of the two parameters is constant: $\alpha + \gamma = k$ for some k (with $k \in (1,4)$, according to the standard estimates). For a fixed value of k (and holding ε and δ, constant), the function

$$\Lambda(\alpha;\lambda) := R^*(\alpha, k-\alpha, \varepsilon, \delta, \lambda) \tag{6}$$

therefore describes how the equity premium is affected by shifting some of the curvature from the g to the v function, keeping constant the overall risk-attitude over money lotteries, for a given specification of the 'rare disaster' parameter λ. The following result summarizes a few key implications of our model:

Proposition 3. *For any $k \in (1,4)$, and for any $\varepsilon > 0$ and $\delta > 0$, there exists a finite $\lambda^* \in [1,\infty)$ such that $\frac{\partial\Lambda}{\partial\alpha}(0;\lambda) > 0$ for all $\lambda > \lambda^*$.*

To understand the significance of this result, recall that the standard model corresponds to the case where $\alpha = 0$ and $\gamma = k$. Hence, the result that $\frac{\partial\Lambda}{\partial\alpha}(0;\lambda) > 0$ means that, holding constant the risk attitude over money lotteries, moving from the standard model to one where $\alpha > 0$ increases the equity premium entailed by the EU model: shifting some of the curvature from the v to the g function, *holding risk attitude constant*, increases the model's ability to account for a higher equity premium. This effect, however, holds provided that the environment features a *sufficiently severe* rare disaster (i.e., if λ is high enough). The intuition is the

[22] As discussed, taking f to be equal to the standard model of intertemporal preferences, our approach yields a special case of Kihlstrom and Mirman (1974)'s model, with the extra restriction that g is CARA. Kihlstrom and Mirman (1974) show that increasing the concavity of the outer function raises risk premia for fixed risky prospects. Although their analysis is distinct, as it consists of a consumption–savings problem with uncertain returns, the qualitative effect of increasing α on the equity premium is the same as in this simple example. For a more general analysis, see Alaoui, Penta, and Troccoli-Moretti (2026).

following: once g is nonlinear (i.e., $\alpha > 0$), pricing depends not only on the curvature through v, but also on the variation in lifetime utility, $v(c_t) + \beta v(c_{t+1})$. Hence, for the shift of curvature from the inner to the outer function to have a quantitative effect, it must be that lifetime utility varies enough across states. In this example, a larger λ makes the bad state become progressively worse, so low-lifetime-utility states become increasingly important for valuation. Thus, in that region, the distinction between inner and outer curvature becomes first-order for pricing. This example therefore suggests a formal connection between *rare-disaster* environments, and the possibility that the distinction between 'pure risk' attitude and marginal utility matters for the pricing of equity, and particularly to explain larger equity premia (for a general analysis, see Alaoui et al. (2026)).

Needless to say, a proper assessment of the extent to which these observations may provide an expected-utility explanation to the equity premium puzzle requires a careful empirical exercise, which is beyond the scope of this section. Nonetheless, we stress that this analysis is fully within the vNM framework, and hence the insights discussed above do not rely on any deviation from Expected Utility theory.

3.4 Discussion

The applications discussed in this section are only some of the theoretical implications of our main results, but there are several others. For instance, in relation to our earlier work on EU models of reference-dependent preferences (Alaoui and Penta, 2026), the results presented here imply that the reversals of risk attitude that occur in each of the representations of that paper (which include, among others, the standard S-shape utility function with *loss aversion* and *diminishing sensitivity*, as well as Genicot and Ray (2017)'s *aspiration model* and Diecidue and Van De Ven (2008)'s utility with a *discontinuity*), must all be coming from the f function, or in other words they must all be due to reference-dependence effects in the certainty domain.[23] That is, unless one posits that 'pure risk' preferences may feature reference dependence with respect to the yardstick itself, in which case one may want to consider a weakening of yardstick neutrality.

Also note that the separation between f and g need not apply to EU only. Aside from its historical role and the central position it still occupies within economics, we focus on EU because the conceptual points we raised are *most trans-*

[23] For an empirical analysis of such expected utility models of reference dependence, see Alaoui, Hervy, Kariv, and Penta (2025), where we perform individual estimation of such preferences in the context of portfolio choices of Arrow securities. Our results highlight a great deal of individual heterogeneity and show a striking predictive power of these models.

parent within it. That is because there is 'only u' within EU (as opposed to other components of the representations in other theories; see, e.g., footnote 2). But extending our approach beyond EU, for instance in combination with other known mechanisms that may contribute to addressing the equity premium puzzle, is clearly a promising direction for future research. Furthermore, as discussed in Alaoui, Penta, and Troccoli-Moretti (2025), the separation we identify between f and g may also suggest novel ways of departing from expected utility, to shed new light on classical 'paradoxes' (e.g., Rabin, 2000) as well as to explain new ones (cf. Alaoui, Penta, and Troccoli-Moretti, 2025).

These are only some examples of further theoretical implications of our results, and there are many others to explore. These developments are clearly beyond the scope of this paper, but we think they are a promising direction for future research.

4 Identification

Our theoretical exercise thus far has relied on the use of a *conceptual yardstick*, not necessarily an actual one. This thought experiment serves to make the point that the 'two horses of different colors' live together within expected utility. This theoretical insight is valid independent of matters of identification. But, from the viewpoint of identification, that thought experiment also serves as a benchmark to clarify what kind of data would be best suited to separately identify the value of wealth (function f) from pure risk aversion (parameter α): if we had the yardstick readily available and agreed upon, as is the case with actual physical yardsticks (or meter sticks) to measure distance, then it would be straightforward to identify the parameters. But in practice, it is perhaps unclear at this stage which commodity can attain the benchmark conditions for a yardstick, or even how it can be verified that such a commodity fulfills the desired properties. In this section we will discuss how this can be done, and how f and α can be identified from choice data.

4.1 Identification for firms

Consider the case of the profit-maximizing firm discussed in Section 3.1, which is both important in itself, and useful to introduce a few conceptual points. There, the yardstick y is the home currency (e.g., dollars), and we let m denote the quantity of money held in a possibly different currency (e.g., euros). Then, letting d denote the exchange rate, in the certain space the firm maximizes $dm+y$. Under risk, as discussed, it maximizes either $u(m,y) = \frac{1-e^{-\alpha(dm+y)}}{\alpha}$ with $\alpha \neq 0$, or $dm+y$

(in which case we will say that $\alpha = 0$ ($d = 1$ if m is money held in home currency). Of course, with enough data on y, all the relevant parameters can be identified in the usual manner. Instead, to illustrate what kind of identification is possible when data on the yardstick is not available, suppose now that we have data on m only, while y is held fixed.

If d is observable, then α can be identified in standard way, from choice data over lotteries over m. But note that the firm's overall risk attitude over m (the foreign currency) now depends on both d and α. The overall Arrow-Pratt coefficient for m, for instance, will be equal to $\alpha_u = \alpha d$.[24] This is intuitive: for instance, consider a gamble between 10 and 0 units of the foreign currency, m. In terms of the firm's own currency, this would amount to a gamble between \$10d and \$0. So, even though the function f here is linear in m, the pure risk coefficient is unique only up to the multiplicative constant $1/d$. Clearly, if d is observable, this does not create any practical issue. Otherwise, only the product αd can be jointly identified from lotteries over m; not α and d separately. At the same time, if our interest were in comparative statics across firms with the same exchange rate, then this product αd would suffice to have an ordering of their risk attitudes, independent of the (unknown) actual value of d.

In this discussion, we took the units of the yardstick to be dollars. Suppose now that we know the exchange rate of foreign to home currency, but we take the yardstick to be (dollar) *cents* instead of dollars, which would be just as plausible of a yardstick. This would be changing the y, and therefore all the coefficients would have to adjust: the exchange rate d would be adjusted by a factor of 100 for cents compared to dollars, and so the elicited Arrow-Pratt coefficient of absolute risk aversion would also adjust accordingly. While this does not have any implications in terms of the fundamentals, it illustrates that without taking a stance on the units of account, identification can only be made up to multiplicative constants.

For firms, the availability of an obvious yardstick also allows a clean identification for more complicated cases, where the firm may also have further objectives besides profits, such as environmental concerns, corporate social responsibility, etc. Letting x denote the variable the firm is concerned with, then the f function could be identified from the MRS between x and money (in own currency) under certainty. In this case, f would not necessarily be linear, but the α parameter could still be identified from lotteries over own currency, or from lotteries over x,

[24]The Arrow-Pratt indices are obviously invariant to positive affine transformations of the utility function, but they are not necessarily invariant to rescaling the units of account. For instance, let $\hat{m} = d \cdot m$ for some $d > 0$. If u is of the CARA family, then $\alpha_u(\hat{m}) = d \cdot \alpha_u(m)$.

adapting the formula in Corollary 1 to x. Namely, $\alpha_u(x) = \alpha f'(x) + \alpha_f(x)$.

4.2 Identification via Intertemporal Choice Problems

Take a standard intertemporal setting, as in Section 3.2, where for any $x = (c_1, c_2, ...) \in X$, the function $f(x) = \sum_{t=1}^{\infty} \beta^{t-1} v(c_t)$ for some 'flow consumption utility' $v : \mathbb{R}_+ \to \mathbb{R}$, so that the overall utility function is

$$u(x) = g\left(\sum_{t=1}^{\infty} \beta^{t-1} v(c_t)\right), \tag{7}$$

for some CARA transformation, $g : \mathbb{R} \to \mathbb{R}$, with parameter α. The f function in this case has a useful structure for identification. First, using standard techniques, we can identify the v function and β using consumption data in the certain space alone, given enough variance in the budget problem faced by the agent, which will lead to different tangency conditions (see below). Second, by the linear separability over time, the risk attitude for lotteries over c_t is the same for any fixed level of $(c_1, ..., c_{t-1}, c_{t+1}, ...)$. Hence, once v and β are identified, we can take lotteries over c_t (holding all other periods consumption fixed), and identify the pure risk parameter α by finding the certainty equivalent of a lottery.

We note as well that existing datasets can be used for the purposes of a similar identification exercise. In particular, the convex time budget decisions in Andreoni and Sprenger (2012a,b) discussed in Section 3.2.2 are particularly apt for this exercise. Using these datasets and standard parametric assumptions on v (specifically, that it is of the CRRA form), but with the inclusion of the g function in the estimates, one can elicit the parameters of risk and time preferences.

More precisely, the choices in the certain tradeoff between the present and the future (cf. Case 1 in Section 3.2.2) can be used to elicit the parameters of the v function and the discount factor β: given a budget m and interest rate r, optimality under the constraint $(1+r)\, c_t^{(1)} + c_{t+k}^{(1)} = m$ yields the tangency condition $\frac{v'(c_t^{(1)})}{\beta^k v'(c_{t+k}^{(1)})} = (1+r)$, exactly as in Andreoni and Sprenger (2012b)'s certain space. Choices made for different interest rates can thus be used to estimate the CRRA parameter of the v and the β. Once this is achieved, the individual's choice of $c_t^{(2)}$ and $c_{t+k}^{(2)}$ for the risky lotteries in Case 2, given the probabilities $p_t^{(2)}$ and $p_{t+k}^{(2)}$, then serve to elicit the risk parameter α of the g function.

4.3 Identification via Proxy

Turning to more general decision-makers (i.e. not necessarily profit maximizers), our first method consists of finding (or constructing) a reliable yardstick, or a proxy for it, to achieve identification to the extent possible. For this exercise, we return to the multiple goods case introduced in Section 2.2.6. In particular, let $x = (x_1, x_2)$, and assume that f is (strictly) increasing in each dimension. Suppose that we are primarily interested in good 1 (e.g., x_1 represents quantity of money, as in most of the paper), while good 2 is the 'candidate proxy' for the yardstick. Under the maintained assumptions on the yardstick, it must be that $u(x,y) = g(f(x)+y)$, where $g : \mathbb{R} \to \mathbb{R}$ is a CARA transformation with parameter $\alpha \in \mathbb{R}$, as above. Now, for good 2 to be an appropriate substitute for the yardstick, it would need to satisfy a quasilinearity property with respect to good 1, but also with respect to the yardstick itself:

Full Quasilinearity (FQ): Good 2 satisfies *full quasilinearity* if, for any x_1, x_2, x_1', $x_2', y, y', t \in \mathbb{R}$, $\delta_{x_1,x_2,y} \succsim \delta_{x_1',x_2',y'}$ if and only if $\delta_{x_1,x_2+t,y} \succsim \delta_{x_1',x_2'+t,y'}$.

If this property holds, under the maintained assumptions on the yardstick, then the $f(\cdot)$ function above must be such that $f(x) = f_1(x_1) + d_2x_2$ for some function f_1 and constant $d_2 > 0$. Thus, the overall utility function for (x, y) must take the form $u(x,y) = g(f_1(x_1)+d_2x_2+y)$. Hence, if one could find a commodity that satisfies FQ, then it would serve as a suitable proxy for the yardstick, and be effectively equivalent to it, up to a multiplicative constant $d_2 > 0$ (just like the exchange rate in the firm example above).

The issue with FQ, however, is that it cannot be directly tested with data that only involve commodities 1 and 2. But the following two implications of FQ are testable with data that only involve commodities 1 and 2:

Proxy Neutrality (PN): Good 2 satisfies *proxy neutrality* if, for any $x_2, x_2', y \in \mathbb{R}$, and for any $p^{x_1}, q^{x_1} \in \Delta(\mathbb{R})$, $(p^{x_1}, x_2, y) \succsim (q^{x_1}, x_2, y)$ if and only if $(p^{x_1}, x_2', y) \succsim (q^{x_1}, x_2', y)$.

Proxy Quasilinearity (PQ): Good 2 satisfies *partial (proxy) quasilinearity* if $\delta_{x_1,x_2,y} \succsim \delta_{x_1',x_2',y}$ if and only if $\delta_{x_1,x_2+t,y} \succsim \delta_{x_1',x_2'+t,y}$ for all $x_1, x_2, x_1', x_2', y, t \in \mathbb{R}$.

That PN is implied by FQ follows from the representation of the u function that holds under FQ (namely, $u(x,y) = g(f_1(x_1)+d_2x_2+y)$), since g being CARA implies that the risk attitudes on the x_1 component is not affected by the level of x_2 in this representation. Interestingly, while FQ only refers to the certainty space, PN is about preferences over lotteries over good 1. This connection from the certain to the uncertain space occurs thanks to the yardstick, which serves as a bridge between the two (FQ connects certain preferences over x to y, which in turn is connected to preferences over the uncertain space via yardstick neutrality).

PQ instead is clearly a weakening of FQ, as it characterizes the strongest implications of FQ that are testable in the certain space with data over x_1 and x_2 alone. PQ has a straightforward interpretation: under PQ, good 2 is only quasilinear with respect to good 1, not necessarily with respect to the yardstick. Since both PN and PQ are implied by FQ, both conditions also provide a way of falsifying FQ indirectly, from choice over lotteries over x_1 and x_2. A commodity that satisfies both PN and PQ therefore is *viable*, as choice data do not rule out that it might be a suitable proxy, as they do not falsify FQ.

Now, suppose that there are multiple goods beyond x_1, which could take the role of x_2 in the discussion above. Say that there is a set $\mathcal{K} = \{1, ..., K\}$ of candidate proxies, over which preferences are increasing, each of which is a *viable proxy* in the sense that it satisfies both PN and PQ. If commodity k also satisfies FQ, then we say that it is a *suitable proxy*.[25] As we show next, under the maintained assumptions on the yardstick, the following holds: if FQ holds for some $k^* \in \mathcal{K}$, then *any* candidate proxy $k \in \mathcal{K}$ satisfies FQ if and only if it satisfies PN and PQ. Putting everything together, we thus obtain the following result:

Proposition 4 (Detecting a Proxy)**.** *Under the maintained assumptions:*

1. *If a suitable proxy exists, any viable proxy is suitable.*
2. *If no proxy is viable, then no suitable proxy exists.*

Note that this means that, if a suitable proxy exists, then it can be detected by testing FQ from choice data over x, via PN and PQ. From now on, we will assume that an ideal proxy exists. Then, from the discussion above, if good 2 satisfies both PN and PQ, then the representation of $u(x, y)$ involves three distinct objects: the function $f_1 : \mathbb{R} \to \mathbb{R}$, the scalar $d_2 > 0$, and CARA parameter $\alpha \in \mathbb{R}$ of the g function. Note, however, that such a representation involves y as well, and hence it contains more information than what could be gathered based on data on (x_1, x_2) alone. Hence, we introduce the following definition:

Definition 2. *Let good 2 be a* suitable proxy*. Then we say that* (f_1, d_2, α) observationally represents $\succsim$ *if, for all* $p^x, q^x \in \Delta(\mathbb{R}^2)$ *and* $y \in \mathbb{R}$*,* $(p^x, y) \succsim (q^x, y)$

[25] To be clear, letting x_2^k denote the quantity of good $k \in \mathcal{K}$, we say that commodity k satisfies FQ (or PQ or PN) if the definition of FQ (resp., PQ or PN) given above holds using quantity x_2^k in the role of x_2.

iff $\mathbb{E}_{p^x} u(x,y) \geq \mathbb{E}_{q^x} u(x,y)$*, where the utility function* u *takes the form* $u(x,y) = g(f_1(x) + d_2 x_2 + y)$ *with CARA function* $g(\cdot)$ *having parameter* α*.*

In words, this notion of *observational representation* allows only for inference on preferences from the available data on x_1 and x_2, without allowing the unobserved y to be varied. As we show next, under the maintained assumptions, the three parameters of the utility representation are identifiable from standard choice data up to a single multiplicative constant (plus an additive one for f_1):

Proposition 5 (Identification via Proxy)**.** *Under the maintained assumptions, if commodity 2 is a* viable proxy*, and* (f_1, d_2, α) *and* (f_1', d_2', α') *both observationally represent* $\succsim$*, there exists* $a > 0$ *and* $b \in \mathbb{R}$ *s.t.* $f_1' = a f_1 + b$*,* $d_2' = a d_2$ *and* $\alpha' = \alpha / a$*.*

Note that since the identification of α and the curvature of f_1 are identifiable up to the same multiplicative constant, essentially all the relevant information is identified save the yardstick's units of account. Such an exercise is particularly useful for inferences on behavior and risk-attitude across domains. We will return to this point in Section 5 below.

We note that while the exercise above concerns finding a suitable proxy for y for the entire range, in practice we need much less. For identification of the relevant parameters, it suffices to have a proxy on *some* range for identification. In particular, suppose that we have identified a proxy that is 'locally' viable, in the sense that PQ and PN hold over some range of commodities, say for all $(x_1, x_2) \in [\underline{x}_1, \overline{x}_1] \times [\underline{x}_2, \overline{x}_2]$. This is all we need to then identify α, from which we can then identify f (in the sense of Proposition 5). Hence, if we have identified a natural proxy for which we believe PN and PQ hold at least 'locally', on some range, then we can fully conduct the entire identification exercise.

In the absence of a 'natural' commodity that can serve as a viable proxy, one could think of constructing an artificial commodity that would plausibly satisfy properties PN and PQ (at least on some range). Consider the following example:[26]

Example 1 (An Artificial Proxy)**.** Take a unit mass of anonymous third parties, whose marginal utilities are constant over some good, let's call it *Token*. Then, $x_2 \in \mathbb{R}_+$ could denote the fraction of such individuals that receive a fixed amount of tokens (say, one).[27] Under the assumption that the decision maker's preferences

[26] We thank Elias Tsakas for inspiring us this example.

[27] Note that no randomization is involved in this setting. In particular, for $x_2 \in (0,1)$, it is not the case that a given agent receives one token with probability x_2, and nothing otherwise. Rather, a fraction x_2 of the anonymous agents receive one token for sure. For $x_2 > 1$, the interpretation is that a fraction $(x_2 - \lfloor x_2 \rfloor)$ of agents get $\lfloor x_2 \rfloor + 1$ tokens, and the rest get $\lfloor x_2 \rfloor$.

are increasing in such a number, and that under anonymity PQ holds, then PN amounts to assuming that the value of x_2 is orthogonal to the decision maker's preferences over lotteries of x_1. Moreover, if we believe that this artificial proxy does not satisfy PN for the full range (for instance, if there are endpoint effect near 100% of the third party receiving the token), but rather for a limited one (say, if between 45% and 65% of the third party receives a token), then we can still fully conduct the identification exercise, as discussed above. □

Proposition 5 also implies that, in order to *uniquely* identify α, we would require at least one data point on y. This is for similar reasons to those discussed in Section 4.1: in this case, without such a data point, we have no way of distinguishing between (f_1, d_2, α) and $(af_1 + b, ad_2, \alpha/a)$ (for $a > 0$) because we have no way of knowing what the correct exchange rate between the proxy and the yardstick is. However, if we could have, say one indifference condition involving y, then we could establish the proper exchange rate (for example, if $(x_1, x_2, 0) \sim^c (x_1, 0, y)$ for some x_1, x_2 and y, then $d_2 = y/x_2$), and hence pin down $(f_1 + b, d_2, \alpha)$ uniquely – aside for the additive constant b – which can of course be normalized to 0).

Alternatively, and perhaps more relevantly for empirical work, suppose that we do not have such a data point, but instead are willing to say that there is a viable proxy that is common to different individuals, and that its relationship with the yardstick is identical across individuals (i.e., we are willing to assume a common normalization). Then, one would still be able to perform interpersonal comparisons of both the α parameters and of the f_1 functions: for any choice of d_2, the α parameter and f_1 function would be uniquely identified, and hence their ordering across agents would be unaffected by the particular choice of d_2.

If, furthermore, one is willing to take a stance on the 'exchange rate' between the proxy and the yardstick – what we would call a *pegged proxy* (formally, a viable proxy in which the d_2 parameter is known, or normalized to one) – then Proposition 5 implies that both the f and α are uniquely identified:

Corollary 2 (Identification via Peg)**.** *Under the maintained assumptions, if commodity 2 is a* pegged proxy*, then f_1 and α are uniquely identifiable.*

Lastly, we use this discussion to illustrate an additional point. Suppose that one were concerned whether, depending on which proxy we use, the function f

But this is only done to accommodate the technical requirement that x_2 lives in an unbounded space. For practical elicitation, the x_2 in this example could be taken from the unit interval, and the token could be a fixed amount of dollars (if $x_2 \in [0, 1]$, the requirement that the third parties have a constant marginal utility over the amount of the token is not needed, as PQ would be ensured by the anonymity of the receiving third parties).

could be concave with respect to one (say), but convex with respect to another, so that all the conclusions drawn would rely on the proxy used. Our discussion shows that this cannot be the case: if both proxies satisfy PN and PQ, they can only be positive affine transformations of one another, and hence the function can only differ by a positive affine transformation. The shape, then, will be the same irrespective of the proxy used.

4.4 Parametric assumptions

Turning to parametric assumptions, suppose that we assume f to be of the CRRA form, i.e. such that that $f(m) = \frac{m^{1-\gamma}}{1-\gamma}$. Then, using Corollary 1, we have:

$$\alpha_u(m) = \alpha m^{-\gamma} + \gamma m^{-1}.$$

Then α and γ can be identified with the following two-step procedure:

1. First, estimate $\alpha_u(m)$ for a set of m's, say $M = \{m_1, ..., m_K\} \subset \mathbb{R}$. Since $\alpha_u(m)$ is the local risk attitude parameter over money, at any given m, it can be estimated using standard methods (e.g., Baillon and L'Haridon, 2019).

2. Then, using the collection of $\{(m,\ \alpha_u(m))_{m \in M}\}$ as the independent and dependent variables, respectively, a maximum likelihood estimation (MLE) procedure can be used to estimate α and γ.

Alternatively, if f is CRRA, then Proposition 1 ensures that u is as follows:

$$u(m;\alpha,\gamma) = \begin{cases} \frac{m^{1-\gamma}}{1-\gamma} & \text{if } \alpha = 0 \\ \frac{1-e^{-\alpha\left(\frac{m^{1-\gamma}}{1-\gamma}\right)}}{\alpha} & \text{if } \alpha \neq 0 \end{cases}.$$

Hence, instead of the above two-step procedure, α and γ could also be estimated directly, using standard MLE methods. Note that the α parameter here is pinned down due to a parametric normalization of the CRRA function f. The γ parameter, however, is pinned down uniquely, irrespective of this normalization.

Another important parametric setting, particularly for macroeconomics applications, is the one in equation (2), with a CRRA v function with parameter γ. In such settings, γ and β can be identified or calibrated using savings decisions under certainty, adapting commonly used methods. In addition to that (or in conjunction, using an MLE procedure to estimate all parameters simultaneously), choice data over time lotteries can be employed to estimate the α parameter. Intuitively, let c^* be steady state consumption, and take $c^+ \geq c^*$ that can be either obtained in period t for sure, or in period $t-1$ with probability p or $t+1$ with probability $1-p$, and suppose that one identifies the *'time lottery equivalent'*, p_{TL}, that makes the agent indifferent between the two options. Then, given β and γ, the specific α that leads to this indifference can be identified.[28] The discussion in this last paragraph is closely related to the discussion in Section 4.2, but with a different identification method, since one varies p to obtain the indifference condition.

5 Cross-domain inference and elicitation

We now return to the case of multiple commodities (domains) $x = (x_1, x_2, .., x_n)$ introduced in Section 2.2. We discuss first how predictions of risk attitude from one domain to another could be conducted. Next we discuss how multiple domains can be used to elicit α, which in turn can also be used to make predictions to additional domains.

5.1 Cross-domain predictions

Suppose now that we are interested, as is common in economics and psychology, in using risk attitudes elicited in one domain to make predictions of risk attitudes in another. It is common to simply take the risk attitude in one domain and analyze the correlation and predictions with the risk attitude on other domains (see e.g. Frey et al. (2017); Mata et al. (2018); Einav et al. (2012) and Vieider et al., 2015). But note that our results suggest caution in interpreting the meaning of these correlations. To show this formally, suppose that preferences under certainty are linearly separable between different domains, and maintain the same assumptions over the yardstick (quasilinearity and neutrality) as we have done throughout. Then, in the certainty space, preferences are represented by

$$f(x) = \sum_{k=1}^{n} f_k(x_k),$$

for some collection $(f_k)_{k=1,\dots,n}$ of functions $f_k : \mathbb{R} \to \mathbb{R}$. Hence, for the overall preferences over risk, we have that $u(x) = g(\sum_k f_k(x))$, where g is CARA with parameter α. It follows that the Arrow-Pratt index for good k, at consumption level x_k, is

$$\alpha_{u,k}(x_k) = \alpha f_k' + \alpha_{f_k}(x_k), \tag{8}$$

[28] That is because $e^{-\alpha\beta^t v(c^*+c^+;\gamma)} = p_{TL} \cdot e^{-\alpha\beta^{t-1} v(c^*+c^+;\gamma)} + (1-p_{TL}) \cdot e^{-\alpha\beta^{t+1} v(c^*+c^+;\gamma)}$.

where $\alpha_{f_i}(x_k) = -\frac{f_k''(x_k)}{f_k'(x_k)}$. In words, for any k, the (overall) risk attitude over that good depends both on the curvature of f_k and on the common parameter α. It follows that, with no further information, eliciting an agent's risk attitude over a good, say k, cannot be used to make predictions about the agent's preferences over lotteries of some other good, say l, unless f_k and f_l happen to be identical.

This is in fact consistent with well-known findings in the psychology literature, which states that choice-based measures of risk attitudes commonly used in economics, which typically pertain to the money domain, are often not predictive of risk attitudes in other domains (see, e.g., Frey et al., 2017, Mata et al., 2018). In that literature, it is common to argue that this is a weakness of the economics approach. But, as we see here, it is rather to be expected, given the way that such analyses are often conducted.

Example 2 (A Formula 1 driver)**.** An F.1 driver has preferences over two domains: a financial one, and a racing one. For simplicity, say that x_1 represents quantity of money, and x_2 is a measure of his success as a racer (e.g., points in Championship, holding everything else constant). This particular driver is sufficiently wealthy that the marginal utility of money is essentially constant, and small. Nonetheless, his financial behavior suggests that $\alpha_{u,1} > 0$ (e.g., he buys insurance, diversifies his investments and portfolio holdings, etc.). Yet, on the race-track, he seems willing to take a great deal of risk, which would suggest that perhaps $\alpha_{u,2} < 0$.

Now, suppose we ask ourselves: Is this F.1 driver risk-averse or not? One could give different answers, and it is apparent that they would depend on which domain appears most salient in our mind. But surely enough, taking this driver to the laboratory to estimate his coefficient of risk aversion, $\alpha_{u,1} > 0$, would not be of much use to predict the amount of risk he would be willing to take on the racing track, or even to predict one's willingness to race in F.1 to begin with.

Our approach, however, gives a more nuanced view: If this F.1 driver is wealthy enough that we are willing to assume his utility for money, f_1, is essentially linear, we can conclude that the evidence from his investment behavior, which suggests an overall risk-aversion over money, must come from his 'pure risk' parameter $\alpha > 0$. With this, the fact that he is willing to take so much risk while racing suggests that his preferences over the 'racing greatness' domain must be such that f_2 is convex. Then, his risk-seeking behavior over the racing domain, in this case, must come from an *increasing* marginal utility of 'racing greatness', which in fact must be strong enough as to offset the fact that, when looking at 'pure risk' per se (i.e., the $\alpha > 0$ parameter), this F.1 driver is actually risk-averse.

This suggests that, in order to make predictions *across* domains, one should at least be able to elicit (or be willing to assume) something about α, f_1, and f_2. Mere risk attitudes over one dimension have no direct bearing on other dimensions. $\square$

To make predictions *across* domains, we thus propose the following method instead. Suppose that we have conducted one of the identification exercises discussed in the subsections above. Let us say that we have used the *Identification by Proxy* method (Section 4.3), where good n is the proxy. Suppose that we have used it with data on good 1 only. This has allowed us to elicit (f_1, d_n, α), up to transformation $(af_1 + b, ad_n, \alpha/a)$ for $a > 0$.

Now suppose that we are willing to make some assumption on the shape of f_2 (e.g., that it is linear), or that alternatively we have used the proxy n and data on the certain space 2 to also identify f_2 (up to transformation af_2, where a is the same as above). We now have estimates for (f_1, f_2, d_n, α), which is all that is required to make predictions on the risk attitude over good 2, as we can now use our estimates to obtain $\alpha_{u,2} = \alpha f_2' + \alpha_{f_2}(x_2)$.

This exercise also shows that the estimates being up to constant a in no way interferes with the predictions that we wish to make. This is because $\alpha_{u,2} = \alpha f_2' + \alpha_{f_2} = (\alpha/a)(af_2') + \alpha_{af_2}$. Intuitively, since the multiplicative constant that defines the non-uniqueness of the identification of the α parameter is the same across dimensions, it crosses out when going from one dimension to another. The estimate of α obtained from one domain is therefore portable to another, once combined with an estimate of the preferences under certainty in the latter.

5.2 Cross-domain elicitation

Within the psychology literature, the typical methodology to elicit risk attitudes is based on stated preferences, and frequently involves asking a number of questions about different domains (see, e.g., Frey et al., 2017). It is also argued in that literature that the measures of risk attitude elicited in this way are often more stable and more predictive than the ones elicited from standard economics methods, despite the latter being choice-based.

Putting aside for a moment the well-known issues of stated preferences, which we will return to at the end of this section, here we discuss how cross-domain elicitation may effectively help in identifying the 'pure risk' parameter α, and the sense in which this may lead to greater stability and predictability.

Specifically, maintain that there are n goods, and now suppose that there is a population of I individuals, and let $x^i = (x_1^i, ..., x_n^i)$ denote the bundle of agent

i, with x_k^i being i's quantity of good k. Suppose again that i's preferences in the certainty space linearly separable, i.e. $f^i(x) = \sum_{k=1}^n f_k^i(x_k^i)$, for some collection $(f_k^i)_{k=1,...,n}$ of functions $f_k^i : \mathbb{R} \to \mathbb{R}$. Although not necessary for the argument that follows, let us assume for simplicity that such preferences are linear in each good, i.e. that $f_k^i(x_k^i) = d_k^i \cdot x_k^i$ for all $i = 1, ..., I$ and $k = 1, ..., n$. We further assume that, for each k, the preference parameters in this population are drawn from some distribution with some (unknown) mean $d_k > 0$, with a noise that is independent and identically distributed across individuals and commodities. More precisely, for each i and k, we have $d_k^i = d_k + \epsilon_k^i$, where ϵ_k^i is i's idiosyncratic preference parameter for good k, drawn from a distribution with 0 mean that is i.i.d. across goods and agents. Under these assumptions, i's utility function is

$$u^i(x^i) = g^i\left(\sum_{k=1}^n (d_k + \epsilon_k^i)x_k^i\right),$$

where g^i is CARA with 'pure risk' parameter α^i.

Next we show that, by eliciting individuals' risk-attitudes over each dimension in this setting, we can obtain a ranking of the individuals' pure risk attitudes. To see this, note that applying the formula for the Arrow-Pratt (AP) index in eq. (8) to this setting, for each i and for each component k, we have:

$$\alpha_{u^i,k}(x_k^i) = \alpha^i(d_k + \epsilon_k^i).$$

Hence, if the number n of dimensions is large enough, for each individual i, the expectation of the average AP-index across dimensions is the following:

$$\bar{\alpha}_{u^i} := \frac{\sum_{k=1}^n \alpha_{u^i,k}(x_k)}{n} = \frac{\sum_{k=1}^n \alpha^i d_k}{n} + \frac{\sum_{k=1}^n \epsilon_k^i}{n} \simeq \frac{\alpha^i \sum_{k=1}^n d_k}{n}, \tag{9}$$

Since the term in the summation is the same for all agents, if the number of elicited dimensions is large enough, then the average AP-indices, $(\bar{\alpha}_{u^i})_{i=1,...,I}$, are fully ranked by the individuals' pure risk parameters, $(\alpha_i)_{i=1,...,I}$. Notice that this argument does not require knowing what the $(d_k)_{k=1,...,n}$'s are. The main identification assumption is that the ϵ_k^i's are i.i.d. across goods and agents.

Of course, this exercise does not map exactly to how risk indices are measured in Psychology, which uses a large number of questions on stated preferences in several domains, and creates an index based on the average of the responses. Rather, it serves to show why eliciting risk-attitudes over multiple dimensions can lead to approximate orderings of the pure risk parameter: Since, unlike the k-specific

f_k^i-terms, such pure risk parameters affect the risk attitudes across components, identifying a reliable ranking of the α^i across different individuals may yield better predictive power across domains. Moreover, this ranking could be correlated with risk-taking behavior in domains for which data has not been collected.

As discussed above, it is common in the psychology literature to ask subjects to state their willingness to take risks in various domains, rather than to use choice-based measures used in economics. While we believe that such methods are within the general spirit of the choice-based measures we assumed in the above analysis, they would not be as precise as carrying out the exercise discussed here. In fact, the exercise above can be seen as a *hybrid* of the two methodologies, which maintains the multidimensional logic from psychology and combines it with the choice-based economics methodology of eliciting Arrow-Pratt indices from choice over lotteries over each dimension separately.

We take the tradeoff between the single-domain choice-based measures and multiple-domain stated preferences to be as follows: commonly used choice-based measures are useful because they are precise and based on behavior, but it is difficult to make predictions from one domain to another for the reasons stated here and in the previous Section. It is also impractical to conduct them fully across domain, simply because it may be difficult to implement choice task over some of these relevant domains. As for multiple-domain stated preferences, these are useful for the reasons discussed in this subsection, but suffer from not being based on behavior. Hence, they may not represent true preferences, and it may often be unclear to the agent what the question is precisely asking for. Overall, to the extent that it can be carried out, it seems to us that a choice-based measure over multiple domains, or even a combination of the two approaches, may well be the most promising avenue, if done in a manner consistent with the logic we discussed in this section.

In closing this section, we note that Qualitative Self-Assessments (QSAs) – which involve a single, simple question on stated-preferences – have been increasingly used in recent years (e.g., Dohmen et al., 2011). When it comes to risk attitudes, the discussion above suggest caution in their use, because it is unclear whether they are eliciting the pure risk attitude α, or the composition with the f_k^i functions in any one domain, or some average composition across domains. In this sense, this observation on the meaning of QSA methods is complementary to the analysis contained in Camerer, Chapman, Ortoleva, Snowberg, and Yariv (2025).

6 Conclusion

The main idea in this paper has been to separate the vNM's utility function, u, into two components, one corresponding to preferences under certainty and the idea of a 'marginal utility' of money, and the other corresponding to 'pure risk' attitude (cf, Section 2). This has allowed us to derive several implications for economic theory, and to show that beyond interpretation, not accounting for these components of the u function has had important implications in the development of economic thought (cf. Section 3). For instance, the natural application of expected utility theory to profit-maximizing firms has typically assumed that firms must be risk neutral as a consequence. But this conclusion follows only if one abandons the typical 'agnostic' position of taking EU to be only a representation (which captures a common decision theoretic view), and hence by implicitly making stronger assumption on preferences. We have shown, in contrast, that profit-maximizing firms need not be risk neutral, even within expected utility.

Likewise, we have argued that the natural application of EU to intertemporal settings is not to take the expected discounted utility (EDU) form, but rather the form discussed in Section 3.2. Interestingly, this is precisely within the form of Kihlstrom and Mirman (1974), which the recent literature has revived to resolve EDU's inability to account for risk-aversion over time lotteries (cf. DeJarnette et al., 2020). As discussed in Section 3.3, the same formulation may also provide novel insights on the celebrated Equity Premium Puzzle (Mehra and Prescott, 1985), and it reconciles the broader EU framework (enriched through the lenses of our results), with the challenging behavioral findings in Andreoni and Sprenger (2012a,b) (cf. Sections 3.2.2 and 4.2). These are only some of the theoretical implications of our main results, but exploring further ones, both within EU and outside of it, is likely a promising direction for future research (cf. Section 3.4).

Our analysis has remained close to the classical economics approach, in the sense that the identification methods discussed in Sections 4 and 5 involve ways of identifying some commodity that could serve as a *proxy* for the yardstick (somewhat reminiscent of Tsakas, 2025), and to verify that this candidate proxy is appropriate based only on 'standard', fully *choice-based* datasets. But given the extensive research in economics and neuroeconomics that focuses on other, more direct ways to measure subjective value and satisfaction (e.g., Glimcher and Rustichini (2004); Camerer (2007, 2008); Rustichini (2009), and Glimcher and Tymula, 2023), we believe that a promising avenue for research could employ these techniques as well. For instance, suppose that we are willing to use auxiliary data on physical expressions of utility (such as firing rates, dopamine levels, etc.) and that we accept one such (non-choice based) measure as a common 'unit of account' for the utility index. Then, such a measure could directly serve as a yardstick, and the f function be identified directly as in our thought experiment from Section 2, from which g can be identified using standard choice data over lotteries.

Clearly, identifying which physical expression of utility is best suited to play the role of a yardstick is inherently a neuroeconomics question, and, as such, it is obviously beyond the scope of this paper. Nonetheless, the key properties for our conceptual yardstick still serve as theoretical guidelines for the properties that a physical measurement should have, in order to serve as a useful unit of measure of utility. Once such a 'physical yardstick' is identified, our exercise can easily be enriched to accommodate these domains, as can the practical identification methods discussed in this paper. This, we think, is a promising direction to further develop, through the lens of our approach.

Appendix

Proof of Proposition 1. Under the vNM axioms and monotonocity, preferences $\succsim$ have an EU representation, where the utility function $u(m, y)$ is strictly increasing in both m and y. By the property of yardstick under certainty, $\delta_{m,y} \succsim \delta_{m',y'}$ iff $f(m) + y \geq f(m') + y'$, which holds iff $u(m, y) \geq u(m', y') \Leftrightarrow g(f(m) + y) \geq g(f(m') + y')$ for strictly increasing $g : \mathbb{R} \to \mathbb{R}$ (note that if g is not strictly increasing, then there exists an (m, y) and (m', y') for which $u(m, y) > u(m', y')$ but $g(f(m)+y) \leq g(f(m')+y')$). Since u is unique up to positive affine transformation, it must then be that $u(m, y) = g(f(m) + y)$, or a strictly positive affine transformation of g. Next, by yardstick neutrality, for any p^m, q^m, y, y': $\sum p^m(m)g(f(m)+y) \geq \sum q^m(m)g(f(m)+y)$ iff $\sum p^m(m)g(f(m)+y') \geq \sum q^m(m)g(f(m)+y')$, meaning that function g must be CARA in y. Using the standard results on CARA, g must take the form $g(x) = \frac{1-e^{-\alpha x}}{\alpha}$ for $\alpha \neq 0$, or $g(x) = x$, or a positive affine transformation thereof. ■

Proof of Proposition 2. Let $u^* : \mathbb{R} \to \mathbb{R}$ denote a utility function in a EU representation of preference system $(\succsim^*, \Delta(\mathbb{R}))$. Now, construct the preference system $(\succsim, \Delta(Z))$ that is represented by $u(m, y) = g(f(m) + y)$ where $g(x) = x$ and $f(m) = u^*(m)$. Clearly, this utility function satisfies all the maintained axioms, and $u(m, 0) = g(f(m)) = u^*(m)$.

In the other direction, take any preference system $(\succsim, \Delta(Z))$ that satisfies the maintained assumptions, and let $\succsim^m$ be defined as in the text, i.e., $p^m \succsim^m q^m$ iff $(p^m, 0) \succsim (q^m, 0)$. Letting $u(m, y) = g(f(m) + y)$ represent $\succsim$, clearly $u^*(m) = u(m, 0)$ is strictly increasing and represents preferences $\succsim^* = \succsim^m$, and furthermore the preference system $(\succsim^*, \Delta(Z))$ satisfies all the vNM axioms. ■

Proof of Proposition 3. See Online Appendix

Proof of Proposition 4.

Part 1. Let $z = (x_2^1, ..., x_2^K) \in \mathbb{R}^K$ be a vector denoting quantities of viable proxies, so that $x = (x_1, z) \in \mathbb{R}^{K+1}$ denote the vector with the good of interest x_1 and the candidate proxies, where $\succsim^c$ are the certain preferences on $\mathbb{R}^{K+1}$. Assume that a suitable proxy exists, and w.l.o.g. let it be proxy K. Furthermore, assume that all candidate proxies are viable. Noting that here, preferences $\succsim^c$ are defined over outcomes $(x, y) \in \mathbb{R}^{K+1} \times \mathbb{R}$, we define $\succsim^{c,0}$ in the following manner: $x \succsim^{c,0} x'$ iff $(x, 0) \succsim^c (x', 0)$. Then $\succsim^{c,0}$ is clearly (strongly) monotonic and continuous. Moreover, PQ of each viable proxy in $\{1, ..., K\}$ implies that each of these goods satisfies quasilinearity in $\succsim^{c,0}$. Hence, by standard quasilinearity results, it must be that $\succsim^{c,0}$ is represented by $f_1(x_1) + \sum_{k \in \{1,...,K\}} d_2^k x_2^k$ for a strictly increasing function $f_1 : \mathbb{R} \to \mathbb{R}$ and constants $d_2^k > 0$ for all $k \in \{1, ..., K\}$ and that any other function that represents it must have the form $h(f_1(x_1) + \sum_{k \in \{1,...,K\}} d_2^k x_2^k)$, where h is strictly increasing. Returning to $\succsim^c$, it must then also be that $(x, 0) \succsim^c (x', 0)$ iff $f_1(x_1) + \sum_{k \in \{1,...,K\}} d_2^k x_2^k \geq f_1(x_1) + \sum_{k \in \{1,...,K\}} d_2^k x_2^k$, and since, by the maintained assumptions on the yardstick, $(x, 0) \succsim^c (x', 0)$ iff $(x, y) \succsim^c (x', y)$ for all y, it also follows that $(x, y) \succsim^c (x', y)$ iff $f_1(x_1) + \sum_{k \in \{1,...,K\}} d_2^k x_2^k \geq f_1(x_1) + \sum_{k \in \{1,...,K\}} d_2^k x_2^k$, and furthermore that a representation of $\succsim^c$ must have the form $h(f_1(x_1) + \sum_{k \in \{1,...,K\}} d_2^k x_2^k) + y$ for strictly increasing h.

Then, under the maintained assumptions on the yardstick (which imply quasilinearity of the yardstick), full quasilinearity of good K, (strong) monotonicity and continuity of $\succsim^c$, we obtain immediately from known results that $(x, y) \succsim^c (x', y')$ iff $f_{-K}(x_1, x_2^1 ..., x_2^{K-1}) + b_K x_2^K + y \geq f_{-K}(x'_1, x_2'^1 ...., x_2'^{K-1}) + b_K x_2'^K + y'$, where function $f_{-K} : \mathbb{R}^{K-1} \to \mathbb{R}$ is strictly increasing and $b_K > 0$.

Hence, combining this last result with the necessity of the form

$$h\left(f_1(x_1) + \sum_{k \in \{1,...,K\}} d_2^k x_2^k\right) + y,$$

it must be that h is a positive affine function, and hence that $\succsim^c$ is represented by the form $af_1(x_1) + a\sum_{k \in \{1,...,K\}} d_i x_2^k + y + \kappa$, where $a > 0$ and $\kappa \in \mathbb{R}$, which in turn means that every viable proxy satisfies FQ, and is thus a suitable proxy.

While this completes the proof that the existence of a suitable proxy implies that every viable proxy is suitable, we also show formally here why FQ suffices for PN. Considering now the full preferences $\succsim$, for yardstick neutrality to hold we must have the form $u(x,y) = g(af_1(x_1) + a\sum_{k\in\{1,...,K\}} d_k x_2^k + y)$ for CARA g (or a positive affine transformation of g; note that the constant e is thus included in this transformation). It thus follows that PN holds for each viable proxy $k \in \{1,...,K\}$.
Part 2. Suppose that no candidate proxy is viable, and hence every candidate proxy violates PQ or PN. Then, by the previous step, FQ cannot hold for any candidate proxy, and hence none can be ideal. ■

Proof of Proposition 5. Returning to the notation x_2 for the proxy, suppose (f_1, d_2, α) and (f_1', d_2', α') both observationally represent $\succsim$. Then in the certain domain, for all x_1, x_2 for which $(x_1, 0) \sim^c (0, x_2)$, it must be that $f_1(x_1) = f_1(0) + d_2 x_2$ and $f_1'(x_1) = f_1'(0) + d_2' x_2$, and hence that $x_2 = \frac{f_1(x_1) - f_1(0)}{d_2} = \frac{f_1'(x_1) - f_1'(0)}{d_2'}$. Rearranging, we obtain $f_1'(x) = \frac{d_2' f_1(x)}{d_2} + \left(f_1'(0) - \frac{d_2' f(0)}{d_2}\right) = af_1(x) + b$, where $a = \frac{d_2'}{d_2}$ and $b = f_1'(0) - \frac{d_2' f(0)}{d_2}$. Hence, it must be that $f_1'(x_1) = af_1(x_1) + b$ and $d_2' = ad_2$, where $a > 0$ and b is a constant. Now let α_2 denote the Arrow Pratt coefficient of $\succsim$ with respect to good 2, respectively. Note that it is must be the same for any (f_1, d_2, α) and (f_1', d_2', α') that both observationally represent $\succsim$. Hence, $\alpha_2 = \alpha d_2 = \alpha' d_2'$ must hold, and thus $\alpha d_2 = \alpha' a d_2$, so that $\alpha' = \frac{\alpha}{a}$. ■

References

Abdellaoui, M., C. Barrios, and P. P. Wakker (2007): "Reconciling introspective utility with revealed preference: Experimental arguments based on prospect theory," *Journal of Econometrics*, 138, 356–378.

Alaoui, L., C. Hervy, S. Kariv, and A. Penta (2025): "Expected Utility with Reference Dependence: Individual-Level Experimental Analysis," *mimeo.*

Alaoui, L. and A. Penta (2026): "Attitudes towards Success and Failure," *American Economic Journal: Microeconomics.*

Alaoui, L., A. Penta, and A. Troccoli-Moretti (2026): "What's in an Equity Premium," *Working Paper.*

Andreoni, J. and C. Sprenger (2012a): "Estimating time preferences from convex budgets," *American Economic Review*, 102, 3333–3356.

——— (2012b): "Risk preferences are not time preferences," *American Economic Review*, 102, 3357–3376.

Apesteguia, J., M. A. Ballester, and A. Gutiérrez (2019): "Random models for the joint treatment of risk and time preferences," *UPF working paper No. 1671.*

Aumann, Y. (2021): "An Axiomatic Definition of Risk Aversion with Applications to Saving Behavior," *mimeo.*

Bell, D. E. and H. Raiffa (1982): "Value and Intrinsic Risk Aversion,"," in *Decision Making: Descriptive, Normative, and Prescriptive Interactions*, ed. by A. Tverski and H. Raiffa, Cambridge University Press, 384–397.

Bernoulli, D. ([1738] 1954): "Exposition of a New Theory of the Measurement of Risk," *Econometrica*, 22, 23–36.

Bleichrodt, H. (2025): *Behavioural Economics*, Palgrave Macmillan Cham.

Camerer, C., J. Chapman, P. Ortoleva, E. Snowberg, and L. Yariv (2025): "Reassessing Qualitative Self-Assessments and Experimental Validation," *mimeo.*

Camerer, C. F. (2007): "Neuroeconomics: using neuroscience to make economic predictions," *The Economic Journal*, 117, C26–C42.

——— (2008): "The potential of neuroeconomics," *Economics & Philosophy*, 24, 369–379.

——— (2013): "Expected Utility Theory," *Utility Theories: Measurements and Applications*, 3, 207.

DeJarnette, P., D. Dillenberger, D. Gottlieb, and P. Ortoleva (2020): "Time Lotteries and Stochastic Impatience," *Econometrica*, 88, 619–656.

Diecidue, E. and J. Van De Ven (2008): "Aspiration level, probability of success and failure, and expected utility," *International Economic Review*, 49, 683–700.

Dietrich, F. (2025): "Welfare vs. Utility," *mimeo.*

Dillenberger, D., D. Gottlieb, and P. Ortoleva (forth.): "Stochastic Impatience and the Separation of Time and Risk Preferences," *Theoretical Economics.*

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. G. Wagner (2011): "Individual risk attitudes: Measurement, determinants, and behavioral consequences," *Journal of the European Economic Association*, 9, 522–550.

Dyer, J. S. and R. K. Sarin (1982): "Relative risk aversion," *Management Science*, 28, 875–886.

Eeckhout, J. and L. Veldkamp (2023): "Data and Markups: A Macro-Finance Perspective," *mimeo.*

Einav, L., A. Finkelstein, I. Pascu, and M. R. Cullen (2012): "How general are risk preferences? Choices under uncertainty in different domains," *American Economic Review*, 102, 2606–2638.

Epstein, L. G. (1992): "Behavior under risk: Recent developments in theory and applications," *Advances in economic theory*, 2, 1–63.

Epstein, L. G. and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica: Journal of the Econometric Society*, 937–969.

Frey, R., A. Pedroni, R. Mata, J. Rieskamp, and R. Hertwig (2017): "Risk preference shares the psychometric structure of major psychological traits," *Science advances*, 3, e1701381.

Friedman, M. and L. J. Savage (1948): "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy*, 56, 279–304.

——— (1952): "The Expected-Utility Hypothesis and the Measurability of Utility," *Journal of Political Economy*, 60, 463–474.

Genicot, G. and D. Ray (2017): "Aspirations and inequality," *Econometrica*, 85, 489–519.

Gerasimou, G. (2021): "Simple preference intensity comparisons," *Journal of Economic Theory*, 105199.

——— (2022): "Corrigendum to Gerasimou (2021) and comment on Banerjee (2022)," *Journal of Economic Theory*, 105542.

Gill, D. and V. Prowse (2012): "A structural analysis of disappointment aversion in a real effort competition," *American Economic Review*, 102, 469–503.

Glimcher, P. W. and A. Rustichini (2004): "Neuroeconomics: the consilience of brain and decision," *Science*, 306, 447–452.

Glimcher, P. W. and A. A. Tymula (2023): "Expected subjective value theory (ESVT): A representation of decision under risk and certainty," *Journal of Economic Behavior & Organization*, 207, 110–128.

Karni, E. (1979): "On Multivariate Risk Aversion," *Econometrica: Journal of the Econometric Society*, 1391–1401.

Ke, S. and M. Zhang (2024): "Decision Making under Multidimensional Risk," *working paper.*

Kihlstrom, R. E. and L. J. Mirman (1974): "Risk aversion with many commodities," *Journal of Economic Theory*, 361–388.

Kreps, D. M. and E. L. Porteus (1978): "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica: journal of the Econometric Society*, 185–200.

Mata, R., R. Frey, D. Richter, J. Schupp, and R. Hertwig (2018): "Risk preference: A view from psychology," *Journal of Economic Perspectives*, 32, 155–172.

Mehra, R. and E. Prescott (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145—-161.

Moscati, I. (2018): *Measuring Utility: From the Marginal Revolution to Behavioral Economics*, New York, NY: Oxford University Press.

——— (2023): *The History and Methodology of Expected Utility*, Cambridge, UK: Cambridge University Press.

Rabin, M. (2000): "Risk aversion and expected-utility theory: A calibration theorem," *Econometrica*, 68, 1281–1281.

Rustichini, A. (2009): "Neuroeconomics:: Formal Models of Decision Making and Cognitive Neuroscience," in *Neuroeconomics*, Elsevier, 33–46.

Schmidt, U. and H. Zank (2022): "Chance theory: A separation of riskless and risky utility," *Journal of Risk and Uncertainty*, 65, 1–32.

Smidts, A. (1997): "The relationship between risk attitude and strenght of preference: a test of intrinsic risk attitude," *Management Science*, 43, 357–370.

Strzalecki, T. (2024): *Decision Theory*, mimeo.

Tsakas, E. (2025): "Beliefs by proxy," *Review of Economic Studies.*

Vieider, F. M., M. Lefebvre, R. Bouchouicha, T. Chmura, R. Hakimov, M. Krawczyk, and P. Martinsson (2015): "Common components of risk and uncertainty attitudes across contexts and domains: Evidence from 30 countries," *Journal of the European Economic Association*, 13, 421–452.

von Neumann, J. and O. Morgenstern (1954): *Theory of Games and Economic Behavior*, Princeton, NJ: Princeton University Press.

Yaari, M. E. (1987): "The dual theory of choice under risk," *Econometrica: Journal of the Econometric Society*, 95–115.

What's in a u?: Online Appendix

Larbi Alaoui Antonio Penta

Proof of Proposition 3

For each $\lambda \geq 1$, define

$$Q(R,\lambda) := p_g(\lambda)(1+R+\delta)^{1-k} + p_m(1+R)^{1-k} + p_b(\lambda)(1+R-\lambda\delta)^{1-k}.$$

At $(\alpha,\gamma) = (0,k)$, homogeneity of v_k reduces the risky problem to

$$\max_{s\in(0,y)} \left\{ \frac{(y-s)^{1-k}}{1-k} + \beta\frac{s^{1-k}}{1-k}Q(R,\lambda) \right\}, \quad \text{so} \quad s_2^*(R,0,k,\varepsilon,\delta,\lambda) = \frac{y(\beta Q(R,\lambda))^{1/k}}{1+(\beta Q(R,\lambda))^{1/k}}.$$

The corresponding indirect utility is

$$U_2(R,0,k,\varepsilon,\delta,\lambda) = \frac{y^{1-k}}{1-k}\left(1+(\beta Q(R,\lambda))^{1/k}\right)^k,$$

while,

$$U_1(0,k) = \frac{y^{1-k}}{1-k}\left(1+(\beta(1+R_f)^{1-k})^{1/k}\right)^k.$$

Hence

$$F(R,0,k,\varepsilon,\delta,\lambda) := U_2(R,0,k,\varepsilon,\delta,\lambda) - U_1(0,k) = 0 \iff Q(R,\lambda) = (1+R_f)^{1-k}. \tag{1}$$

For fixed λ, $Q(\cdot,\lambda)$ is continuous and strictly decreasing on $(\lambda\delta-1,\infty)$, with $Q(R,\lambda) \to +\infty$ as $R \downarrow \lambda\delta - 1$ and $Q(R,\lambda) \to 0$ as $R \to \infty$. By the Intermediate Value Theorem, there is therefore a unique $R_0(\lambda)$ solving (1). Evaluating the formula for s_2^* at $R_0(\lambda)$ and using (1) yields

$$s_2^0(\lambda) := s_2^*(R_0(\lambda),0,k,\varepsilon,\delta,\lambda) = \frac{y(\beta(1+R_f)^{1-k})^{1/k}}{1+(\beta(1+R_f)^{1-k})^{1/k}} = s_1^0.$$

Thus risky saving at the EU corner is independent of λ and equal to certainty saving. In particular, under the normalisation $(c_t^f, c_{t+1}^f) = (1,1)$, first-period consumption remains equal to 1 along the benchmark path. Now define $x_\lambda :=$

$1+R_0(\lambda)-\lambda\delta>0$. Rewriting (1) at $R_0(\lambda)$ gives

$$\frac{\lambda\varepsilon}{1+\lambda}\Big((\lambda+1)\delta+x_\lambda\Big)^{1-k}+(1-\varepsilon)(\lambda\delta+x_\lambda)^{1-k}+\frac{\varepsilon}{1+\lambda}x_\lambda^{1-k}=(1+R_f)^{1-k}.$$

If $x_\lambda\nrightarrow 0$, some subsequence satisfies $x_{\lambda_n}\geq\eta>0$. Since $1-k<0$, each term on the left then converges to 0, a contradiction. Hence $x_\lambda\to 0$. Therefore, writing

$$c_b(\lambda):=s_1^0x_\lambda,\qquad c_m(\lambda):=s_1^0(\lambda\delta+x_\lambda),\qquad c_g(\lambda):=s_1^0((\lambda+1)\delta+x_\lambda),$$

we have

$$c_b(\lambda)\to 0,\qquad c_m(\lambda)\to\infty,\qquad c_g(\lambda)\to\infty. \tag{2}$$

By the Implicit Function Theorem,

$$\Lambda_\alpha(0,\varepsilon,\delta,\lambda)=\frac{F_\gamma-F_\alpha}{F_R},$$

with all derivatives evaluated at $(R,\alpha,\gamma)=(R_0(\lambda),0,k)$. By the envelope theorem,

$$F_R=\beta s_2^0(\lambda)\Big(p_g(\lambda)v_k'(c_g(\lambda))+p_mv_k'(c_m(\lambda))+p_b(\lambda)v_k'(c_b(\lambda))\Big)>0.$$

Moreover, if $U_t(\lambda):=v_k(1)+\beta v_k(\widetilde{c}_{t+1}(\lambda))$ and $U^f:=v_k(1)+\beta v_k(1)$, then $g_\alpha(z)=z-\frac{\alpha}{2}z^2+o(\alpha)$ implies

$$F_\alpha=-\frac{1}{2}\mathbb{E}_\lambda[U_t(\lambda)^2]+\frac{1}{2}(U^f)^2.$$

Since $F(R_0(\lambda),0,k,\varepsilon,\delta,\lambda)=0$, one has $\mathbb{E}_\lambda[U_t(\lambda)]=U^f$, hence

$$F_\alpha=-\frac{1}{2}\mathrm{Var}(U_t(\lambda)).$$

Finally,

$$F_\gamma=\beta\left(p_g(\lambda)\,\frac{\partial v_\gamma}{\partial\gamma}(c_g(\lambda))\bigg|_{\gamma=k}+p_m\,\frac{\partial v_\gamma}{\partial\gamma}(c_m(\lambda))\bigg|_{\gamma=k}+p_b(\lambda)\,\frac{\partial v_\gamma}{\partial\gamma}(c_b(\lambda))\bigg|_{\gamma=k}-\frac{\partial v_\gamma}{\partial\gamma}(1)\bigg|_{\gamma=k}\right),$$

because the first-period terms cancel. Therefore

$$\Lambda_\alpha(0,\varepsilon,\delta,\lambda)>0\iff F_\gamma\Big(R_0(\lambda),0,k,\varepsilon,\delta,\lambda\Big)>-\frac{1}{2}\mathrm{Var}(U_t(\lambda)). \tag{3}$$

It remains to compare the two terms in (3). Since $s_1^0=(1+R_f)^{-1}$, multiplying

(1) by $(s_1^0)^{1-k}$ yields

$$p_g(\lambda)c_g(\lambda)^{1-k} + p_m c_m(\lambda)^{1-k} + p_b(\lambda)c_b(\lambda)^{1-k} = 1. \tag{4}$$

Using (2) and $1-k<0$, the good and middle terms in (4) are $o(1)$, so

$$p_b(\lambda)c_b(\lambda)^{1-k} = 1 + o(1). \tag{5}$$

If $X_\lambda := v_k(\widetilde{c}_{t+1}(\lambda))$, then $U_t(\lambda) = v_k(1) + \beta X_\lambda$, so $\mathrm{Var}(U_t(\lambda)) = \beta^2 \mathrm{Var}(X_\lambda)$. Also, by (4),

$$\mathbb{E}_\lambda[X_\lambda] = \frac{1}{1-k} = -\frac{1}{k-1}.$$

For the second moment, the bad state dominates:

$$p_b(\lambda)v_k(c_b(\lambda))^2 = \frac{1}{(k-1)^2}\frac{(p_b(\lambda)c_b(\lambda)^{1-k})^2}{p_b(\lambda)} = \frac{\lambda}{(k-1)^2\varepsilon} + o(\lambda),$$

where we used (5) and $p_b(\lambda) = \varepsilon/(1+\lambda)$, while the good and middle contributions are $o(1)$ by (2). Hence

$$\mathrm{Var}(U_t(\lambda)) = \frac{\beta^2}{(k-1)^2\varepsilon}\lambda + o(\lambda). \tag{6}$$

Next, we have

$$\frac{\partial v_\gamma}{\partial \gamma}(c)\bigg|_{\gamma=k} = \frac{c^{1-k}\Big(1 + (k-1)\log c\Big)}{(k-1)^2}.$$

By (2), $c_m(\lambda)$ and $c_g(\lambda)$ are of order λ, so their contributions to F_γ are $O(\lambda^{1-k}\log\lambda) = o(1)$. From (5),

$$c_b(\lambda)^{1-k} = \frac{1+o(1)}{p_b(\lambda)} = \frac{1+\lambda}{\varepsilon}(1+o(1)),$$

and therefore

$$\log c_b(\lambda) = -\frac{1}{k-1}\log\lambda + O(1).$$

Substituting these relations into the bad-state term gives

$$p_b(\lambda)\,\frac{\partial v_\gamma}{\partial \gamma}(c_b(\lambda))\bigg|_{\gamma=k} = -\frac{1}{(k-1)^2}\log\lambda + o(\log\lambda).$$

Since $\frac{\partial v_\gamma}{\partial \gamma}(1)\Big|_{\gamma=k}$ is constant, it follows that

$$F_\gamma\Big(R_0(\lambda), 0, k, \varepsilon, \delta, \lambda\Big) = -\frac{\beta}{(k-1)^2}\log\lambda + o(\log\lambda). \tag{7}$$

Combining (6) and (7),

$$F_\gamma\Big(R_0(\lambda), 0, k, \varepsilon, \delta, \lambda\Big) + \frac{1}{2}\mathrm{Var}(U_t(\lambda)) = \frac{\beta^2}{2(k-1)^2\varepsilon}\lambda - \frac{\beta}{(k-1)^2}\log\lambda + o(\lambda),$$

and division by λ yields the strictly positive limit $\frac{\beta^2}{2(k-1)^2\varepsilon}$. Hence there exists a finite $\lambda^*(\varepsilon, \delta) \geq 1$ such that, for every $\lambda > \lambda^*(\varepsilon, \delta)$,

$$F_\gamma\Big(R_0(\lambda), 0, k, \varepsilon, \delta, \lambda\Big) > -\frac{1}{2}\mathrm{Var}(U_t(\lambda)).$$

By (3), $\Lambda_\alpha(0, \varepsilon, \delta, \lambda) > 0$ for all $\lambda > \lambda^*(\varepsilon, \delta)$, as claimed.

What's in a u?: Online Appendix

Larbi Alaoui Antonio Penta

Proof of Proposition 3

For each $\lambda \geq 1$, define

$$Q(R,\lambda) := p_g(\lambda)(1+R+\delta)^{1-k} + p_m(1+R)^{1-k} + p_b(\lambda)(1+R-\lambda\delta)^{1-k}.$$

At $(\alpha,\gamma) = (0,k)$, homogeneity of v_k reduces the risky problem to

$$\max_{s\in(0,y)} \left\{ \frac{(y-s)^{1-k}}{1-k} + \beta\frac{s^{1-k}}{1-k}Q(R,\lambda) \right\}, \quad \text{so} \quad s_2^*(R,0,k,\varepsilon,\delta,\lambda) = \frac{y(\beta Q(R,\lambda))^{1/k}}{1+(\beta Q(R,\lambda))^{1/k}}.$$

The corresponding indirect utility is

$$U_2(R,0,k,\varepsilon,\delta,\lambda) = \frac{y^{1-k}}{1-k}\left(1+(\beta Q(R,\lambda))^{1/k}\right)^k,$$

while,

$$U_1(0,k) = \frac{y^{1-k}}{1-k}\left(1+(\beta(1+R_f)^{1-k})^{1/k}\right)^k.$$

Hence

$$F(R,0,k,\varepsilon,\delta,\lambda) := U_2(R,0,k,\varepsilon,\delta,\lambda) - U_1(0,k) = 0 \iff Q(R,\lambda) = (1+R_f)^{1-k}. \tag{1}$$

For fixed λ, $Q(\cdot,\lambda)$ is continuous and strictly decreasing on $(\lambda\delta-1,\infty)$, with $Q(R,\lambda)\to+\infty$ as $R\downarrow\lambda\delta-1$ and $Q(R,\lambda)\to 0$ as $R\to\infty$. By the Intermediate Value Theorem, there is therefore a unique $R_0(\lambda)$ solving (1). Evaluating the formula for s_2^* at $R_0(\lambda)$ and using (1) yields

$$s_2^0(\lambda) := s_2^*(R_0(\lambda),0,k,\varepsilon,\delta,\lambda) = \frac{y(\beta(1+R_f)^{1-k})^{1/k}}{1+(\beta(1+R_f)^{1-k})^{1/k}} = s_1^0.$$

Thus risky saving at the EU corner is independent of λ and equal to certainty saving. In particular, under the normalisation $(c_t^f, c_{t+1}^f) = (1,1)$, first-period consumption remains equal to 1 along the benchmark path. Now define $x_\lambda :=$

$1 + R_0(\lambda) - \lambda\delta > 0$. Rewriting (1) at $R_0(\lambda)$ gives

$$\frac{\lambda\varepsilon}{1+\lambda}\Big((\lambda+1)\delta + x_\lambda\Big)^{1-k} + (1-\varepsilon)(\lambda\delta + x_\lambda)^{1-k} + \frac{\varepsilon}{1+\lambda}x_\lambda^{1-k} = (1+R_f)^{1-k}.$$

If $x_\lambda \nrightarrow 0$, some subsequence satisfies $x_{\lambda_n} \geq \eta > 0$. Since $1-k<0$, each term on the left then converges to 0, a contradiction. Hence $x_\lambda \to 0$. Therefore, writing

$$c_b(\lambda) := s_1^0 x_\lambda, \qquad c_m(\lambda) := s_1^0(\lambda\delta + x_\lambda), \qquad c_g(\lambda) := s_1^0((\lambda+1)\delta + x_\lambda),$$

we have

$$c_b(\lambda) \to 0, \qquad c_m(\lambda) \to \infty, \qquad c_g(\lambda) \to \infty. \tag{2}$$

By the Implicit Function Theorem,

$$\Lambda_\alpha(0,\varepsilon,\delta,\lambda) = \frac{F_\gamma - F_\alpha}{F_R},$$

with all derivatives evaluated at $(R,\alpha,\gamma) = (R_0(\lambda),0,k)$. By the envelope theorem,

$$F_R = \beta s_2^0(\lambda)\Big(p_g(\lambda)v_k'(c_g(\lambda)) + p_m v_k'(c_m(\lambda)) + p_b(\lambda)v_k'(c_b(\lambda))\Big) > 0.$$

Moreover, if $U_t(\lambda) := v_k(1) + \beta v_k(\widetilde{c}_{t+1}(\lambda))$ and $U^f := v_k(1) + \beta v_k(1)$, then $g_\alpha(z) = z - \frac{\alpha}{2}z^2 + o(\alpha)$ implies

$$F_\alpha = -\frac{1}{2}\mathbb{E}_\lambda[U_t(\lambda)^2] + \frac{1}{2}(U^f)^2.$$

Since $F(R_0(\lambda),0,k,\varepsilon,\delta,\lambda) = 0$, one has $\mathbb{E}_\lambda[U_t(\lambda)] = U^f$, hence

$$F_\alpha = -\frac{1}{2}\mathrm{Var}(U_t(\lambda)).$$

Finally,

$$F_\gamma = \beta\left(p_g(\lambda)\,\frac{\partial v_\gamma}{\partial\gamma}(c_g(\lambda))\bigg|_{\gamma=k} + p_m\,\frac{\partial v_\gamma}{\partial\gamma}(c_m(\lambda))\bigg|_{\gamma=k} + p_b(\lambda)\,\frac{\partial v_\gamma}{\partial\gamma}(c_b(\lambda))\bigg|_{\gamma=k} - \frac{\partial v_\gamma}{\partial\gamma}(1)\bigg|_{\gamma=k}\right),$$

because the first-period terms cancel. Therefore

$$\Lambda_\alpha(0,\varepsilon,\delta,\lambda) > 0 \iff F_\gamma\Big(R_0(\lambda),0,k,\varepsilon,\delta,\lambda\Big) > -\frac{1}{2}\mathrm{Var}(U_t(\lambda)). \tag{3}$$

It remains to compare the two terms in (3). Since $s_1^0 = (1+R_f)^{-1}$, multiplying

(1) by $(s_1^0)^{1-k}$ yields

$$p_g(\lambda)c_g(\lambda)^{1-k} + p_m c_m(\lambda)^{1-k} + p_b(\lambda)c_b(\lambda)^{1-k} = 1. \tag{4}$$

Using (2) and $1-k<0$, the good and middle terms in (4) are $o(1)$, so

$$p_b(\lambda)c_b(\lambda)^{1-k} = 1 + o(1). \tag{5}$$

If $X_\lambda := v_k(\widetilde{c}_{t+1}(\lambda))$, then $U_t(\lambda) = v_k(1) + \beta X_\lambda$, so $\mathrm{Var}(U_t(\lambda)) = \beta^2 \mathrm{Var}(X_\lambda)$. Also, by (4),

$$\mathbb{E}_\lambda[X_\lambda] = \frac{1}{1-k} = -\frac{1}{k-1}.$$

For the second moment, the bad state dominates:

$$p_b(\lambda)v_k(c_b(\lambda))^2 = \frac{1}{(k-1)^2}\frac{(p_b(\lambda)c_b(\lambda)^{1-k})^2}{p_b(\lambda)} = \frac{\lambda}{(k-1)^2\varepsilon} + o(\lambda),$$

where we used (5) and $p_b(\lambda) = \varepsilon/(1+\lambda)$, while the good and middle contributions are $o(1)$ by (2). Hence

$$\mathrm{Var}(U_t(\lambda)) = \frac{\beta^2}{(k-1)^2\varepsilon}\lambda + o(\lambda). \tag{6}$$

Next, we have

$$\left.\frac{\partial v_\gamma}{\partial \gamma}(c)\right|_{\gamma=k} = \frac{c^{1-k}\Big(1 + (k-1)\log c\Big)}{(k-1)^2}.$$

By (2), $c_m(\lambda)$ and $c_g(\lambda)$ are of order λ, so their contributions to F_γ are $O(\lambda^{1-k}\log\lambda) = o(1)$. From (5),

$$c_b(\lambda)^{1-k} = \frac{1+o(1)}{p_b(\lambda)} = \frac{1+\lambda}{\varepsilon}(1+o(1)),$$

and therefore

$$\log c_b(\lambda) = -\frac{1}{k-1}\log\lambda + O(1).$$

Substituting these relations into the bad-state term gives

$$p_b(\lambda)\left.\frac{\partial v_\gamma}{\partial \gamma}(c_b(\lambda))\right|_{\gamma=k} = -\frac{1}{(k-1)^2}\log\lambda + o(\log\lambda).$$

Since $\left.\frac{\partial v_\gamma}{\partial\gamma}(1)\right|_{\gamma=k}$ is constant, it follows that

$$F_\gamma\Big(R_0(\lambda), 0, k, \varepsilon, \delta, \lambda\Big) = -\frac{\beta}{(k-1)^2}\log\lambda + o(\log\lambda). \tag{7}$$

Combining (6) and (7),

$$F_\gamma\Big(R_0(\lambda), 0, k, \varepsilon, \delta, \lambda\Big) + \frac{1}{2}\mathrm{Var}(U_t(\lambda)) = \frac{\beta^2}{2(k-1)^2\varepsilon}\lambda - \frac{\beta}{(k-1)^2}\log\lambda + o(\lambda),$$

and division by λ yields the strictly positive limit $\frac{\beta^2}{2(k-1)^2\varepsilon}$. Hence there exists a finite $\lambda^*(\varepsilon,\delta) \geq 1$ such that, for every $\lambda > \lambda^*(\varepsilon,\delta)$,

$$F_\gamma\Big(R_0(\lambda), 0, k, \varepsilon, \delta, \lambda\Big) > -\frac{1}{2}\mathrm{Var}(U_t(\lambda)).$$

By (3), $\Lambda_\alpha(0,\varepsilon,\delta,\lambda) > 0$ for all $\lambda > \lambda^*(\varepsilon,\delta)$, as claimed.