Background Risk and Small-Stakes Risk Aversion*

Xiaosheng Mu† Luciano Pomatto‡ Philipp Strack§ Omer Tamuz¶

June 3, 2024

Abstract

Building on Pomatto, Strack, and Tamuz (2020), we identify a tight condition for when background risk can induce first-order stochastic dominance. Using this condition, we show that under plausible levels of background risk, no theory of choice under risk can simultaneously satisfy the following three economic postulates: (i) Decision makers are risk-averse over small gambles, (ii) their preferences respect stochastic dominance, and (iii) they account for background risk. This impossibility result applies to expected utility theory, prospect theory, rank dependent utility and many other models.

1 Introduction

How humans evaluate the trade-off between risks and rewards is one of the core questions in economics. In solving this trade-off many people exhibit *small-stakes risk aversion*: small, actuarially favorable gambles—such as a lottery where one loses \$10 or wins \$11 with equal probability—are often rejected. In this paper we study how background risk (stemming, for instance, from investments in the stock market or health conditions) affects risk attitudes towards small gambles. We argue that for plausible levels of background risk, no preference

*We want to thank the editor, Ted O'Donoghue, as well as Ned Augenblick, Nick Barberis, Sebastian Ebert, Stefano DellaVigna, Armin Falk, Christian Gollier, Paul Heidhues, Matthew Rabin, Todd Sarver, Dmitry Taubinsky, Charlie Sprenger and Juuso Välimäki for helpful comments and discussions.

†Princeton University. Email: xmu@princeton.edu. Xiaosheng Mu acknowledges the hospitality of Columbia University and the Cowles Foundation at Yale University, which hosted him during parts of this research.

‡Caltech. Email: luciano@caltech.edu.

§Yale University. Email: philipp.strack@yale.edu. Philipp Strack was supported by a Sloan fellowship.

¶Caltech. Email: tamuz@caltech.edu. Omer Tamuz was supported by a grant from the Simons Foundation (#419427), by a BSF grant (#2018397), by National Science Foundation CAREER award (DMS1944153), and by a Sloan fellowship.

can explain small-stakes risk aversion without violating the basic assumption that more money is preferred over less.

For example, we show that the aforementioned gamble, where one loses \$10 or wins \$11 with equal probability, must be accepted by any decision maker with preferences that are monotone in first-order stochastic dominance, provided that she faces an independent background risk with exponential tails and standard deviation larger than \$200. This insight is not specific to this particular example. Our main result, Theorem 1, establishes that for every large enough background risk, it is first-order stochastically dominant to accept any small enough favorable gamble. We define appropriate notions of "large enough" and "small enough" formally below.

This result builds on the main theorem in Pomatto et al. (2020), which implies that for every random variable X with positive expectation there is an independent random variable W such that $W + X$ first-order stochastically dominates W. Theorem 1 provides a tight condition on X and W that guarantees that this occurs. In order to prove it, we introduce a new technique for studying the effect of background risk, inspired by the literature on expected utility (Gollier and Schlesinger, 2003; Gollier, 2004): for a given background risk we define an associated *test utility function*, and show that when an expected utility agent possessing this utility and no background risk accepts a gamble at all wealth levels, then accepting this gamble is dominant under the original background risk. We then provide a simple bound on how large the background risk needs to be for this conclusion to hold for a given gamble. This bound implies that under reasonable levels of background risk it is stochastically dominant to accept many gambles that are frequently rejected. Thus, if a decision maker does display small stakes risk aversion and has monotone preferences, she must be engaging in narrow framing: namely, she must be considering the gamble in isolation, ignoring background risk.

To illustrate the result, consider a decision maker who is facing a binary gamble X under which she gains G dollars or loses L dollars with equal probability. If the gamble was taken in isolation, then the choice of whether or not to accept it would depend on her preferences between X and a sure outcome of 0. But if the decision maker is facing an independent background risk W regarding her wealth, then the relevant choice is between W, if the gamble is rejected, and $W + X$, if the gamble is accepted. Table 1 displays different levels of standard deviations and distributional assumptions for the background risk under which $W + X$ dominates W in first-order stochastic dominance. For example, in the case of the fifty-fifty gamble with gain $G = 11$ and loss $L = 10$, it is dominant to accept whenever the

Gamble	**StDeviation of Background Risk: σ**		
Gain/Loss	Laplace	Logistic	Normal*
\$11/\$10	$\sigma \geq$ \$156	$\sigma \geq$ \$200	$\sigma \geq$ \$3318
\$55/\$50	$\sigma \geq$ \$778	$\sigma \geq$ \$998	$\sigma \geq$ \$7421
\$110/\$100	$\sigma \geq$ \$1557	$\sigma \geq$ \$1997	$\sigma \geq$ \$10,498
\$550/\$500	$\sigma \geq$ \$7784	$\sigma \geq$ \$9983	$\sigma \geq$ \$23,525
\$1100/\$1000	$\sigma \geq$ \$15,568	$\sigma \geq$ \$19,967	$\sigma \geq$ \$33,361

Table 1: Standard deviation of background risk sufficient for $W + X$ to first-order stochastically dominate W, where X is a fifty-fifty gamble and W is an independent background risk, under different distributional assumptions on W. The numbers displayed are bounds derived from our Theorem 1 and Corollary 2.
*The normal distribution has mean \$100,000, and in this case the decision maker's wealth is bounded below by \$0, according to a limited liability assumption which we discuss in §2.

decision maker's wealth has a standard deviation higher than \$200 and is either Laplace or Logistically distributed.

The standard deviation of many real-life risks plausibly exceeds this threshold by a large margin. For example, an investor who has \$100,000 in an S&P 500 index fund reasonably faces a wealth risk with standard deviation \$1,000 (or 1%) for the value of her portfolio at the end of each *day*, and \$15,000 at the end of the year.[1] Nevertheless, small-stakes gambles are commonly rejected. Barberis, Huang, and Thaler (2006) find that among clients of a U.S. bank with median wealth exceeding \$10 million, the rejection rate of a hypothetical \$550/\$500 gamble is 71%. Table 1 suggests that this behavior is inconsistent with these investors taking into account even the short-term background risk they face.

Table 1 applies to all monotone preferences. Since monotonicity is, by itself, a very weak assumption, we need some conditions on the distribution of the background risk W to ensure that $W + X$ stochastically dominates W. In particular, W must have full support and the left tail of its distribution must be sufficiently thick. However, if the decision maker is protected by limited liability, so that her final wealth cannot go below a certain threshold, then no assumptions on the tails of the background risk are required. Moreover,

[1] See, e.g., Bardgett, Gourier, and Leippold (2019).

for particular preference specifications, our results continue to hold under smaller background risks. For example, in §A of the Appendix, we calculate the level of background risk needed for a decision maker with Cumulative Prospect Theory preferences to accept various small gambles, under the same parameter values calibrated by Tversky and Kahneman (1992). Compared to Table 1, the required levels of standard deviation are significantly smaller: a fifty-fifty gamble with gain $G = 11$ and loss $L = 10$ must be accepted whenever the background risk has a standard deviation higher than \$62, and is either Laplace, Logistic or Normally distributed.

Our analysis shows a tension between three natural requirements for any theory of choice under risk: (i) risk aversion over small gambles, (ii) monotonicity with respect to first-order stochastic dominance, and (iii) accounting for background risk. As (i) is commonly observed in real world choices, and relaxing (ii) is widely considered unappealing, our results suggest that theories that do not account for narrow framing—whereby independent sources of risk are evaluated separately by the decision maker—cannot explain commonly observed choices among risky alternatives.

1.1 Related Literature

Arrow (1970) and Pratt (1964) establish that under expected utility and a twice-differentiable utility function, a decision maker accepts any actuarially favorable gamble, provided that it is scaled to be small enough. Rabin (2000) shows that the degree of concavity necessary for expected utility theory to explain small-stakes risk aversion leads to implausible choices over large lotteries.[2] The literature has then suggested two different ways of explaining small-stakes risk aversion: (i) by considering more general preferences that allow for *loss aversion* or *first-order risk aversion* and (ii) by allowing for *narrow framing*.

A variety of alternatives to expected utility theory feature first-order risk aversion (Segal and Spivak, 1990; Ang et al., 2005; Khaw et al., 2020), i.e. non-vanishing risk aversion over small risks. Among the most prominent ones are prospect theory (Kahneman and Tversky, 1979), rank dependent utility (Quiggin, 1982), disappointment aversion (Gul, 1991) and expectations-based reference dependent preferences (Kőszegi and Rabin, 2007). While these theories can explain small-stakes risk aversion when such risks are evaluated in isolation,

[2]For example, given any risk-averse expected utility preference, if a gamble where one loses \$100 or wins \$110 with equal probability is rejected at all wealth levels below \$300,000, then a gamble where one loses \$2000 and wins \$12,000,000 with equal probability must also be rejected at wealth levels below \$290,000. See also Hansson (1988) for an early example illustrating this point. Zambrano (2020) provides further results on calibrating expected utility preferences.

subsequent work suggests that this is no longer true once background risk is taken into account.

In a dynamic context, Barberis, Huang, and Thaler (2006) show that in the presence of large background risk, Rabin's critique extends to recursive disappointment-averse preferences (see also Sarver, 2018). Kőszegi and Rabin (2007, page 1058) show in an example that a decision maker with reference dependent preference will for some gambles "approach risk neutrality [...] even for relatively limited amounts of background risk". Safra and Segal (2008) prove that for any risk-averse preference admitting a Gâteaux differentiable representation, a decision maker who rejects a small gamble with positive mean under all background risks must also reject highly favorable large gambles. Kőszegi and Rabin (2009) explore dynamic belief-dependent utility which sometimes predicts behavior consistent with narrow framing, and thus small stakes risk aversion, even in the presence of background risk. Our results imply that their decision-makers must violate first-order stochastic dominance when the background risk is large enough.

Recently, Tarsney (2018) and Pomatto, Strack, and Tamuz (2020) independently proved that a carefully chosen background risk can induce stochastic dominance between gambles. In particular, Pomatto, Strack, and Tamuz (2020) show that if X and Y are random variables with $\mathbb{E}[X] > \mathbb{E}[Y]$, then there exists an independent random variable W, tailored to X and Y, such that $X + W$ stochastically dominates $Y + W$. That result implies that no monotone preference can reject an actuarial favorable gamble under *all* background risks.

Our analysis differs from the existing literature in several ways. First, unlike the literature on first-order risk aversion, we do not focus on a specific class of preferences, but instead consider general monotone preferences. We also relax the assumption that the decision maker is globally risk-averse, which was crucial in the preceding literature. Second, our results focus on small-stakes risk aversion per se, and do not pertain to the decision maker's behavior with regard to large, and possibly hypothetical, gambles.

In addition, in this paper we go beyond proving the existence of a particular distribution of background risk under which accepting a given small gamble is dominant (Pomatto, Strack, and Tamuz, 2020), as our results hold uniformly over all background risks W that are sufficiently large compared to the gamble X. Thus, whenever a decision maker is observed to reject small gambles despite facing considerable background risks, an analyst can conclude that she either violates stochastic dominance or performs narrow framing, irrespective of the specific details of the gamble or of the background risk at hand.

Another innovation is that we provide explicit and tight lower bounds on the size of

the background risk necessary for the dominance relation to hold. Such bounds allow us to work with elementary families of distributions, such as Laplace or Logistic. This improved tractability is crucial for deriving practical estimates of the type displayed in Table 1. We derive these bounds using a new general method for studying the effect of background risk. This is based on the idea of a *test utility function*, which we discuss in §5.1.

To the best of our knowledge, our paper is the first to demonstrate theoretically that whenever a plausible level of background risk is taken into consideration, minimal assumptions on preferences lead to risk-neutral behavior for small gambles. Our analysis thus implies that narrow framing, which has been suggested as a possible explanation for small-stakes risk aversion,[3] is in fact an essential ingredient of any such explanation. Conversely, no departure from expected utility to another monotone preference can, on its own, suffice to explain small stakes risk aversion.

2 Model

A decision maker faces a choice between accepting or rejecting a gamble described by a bounded random variable X that takes negative values with positive probability. We thus rule out the trivial case where $X \geq 0$ and hence the gamble almost surely pays out a positive amount. The decision maker's wealth W is random and independent of X, and accepting the gamble leads to final wealth $W + X$. We interpret W as background risk the decision maker faces when considering whether or not to accept the gamble. We assume W is distributed according to a density $g\colon \mathbb{R} \to \mathbb{R}_+$ that has full support, is eventually decreasing, and is piece-wise continuously differentiable.[4] This is a weak technical assumption that holds for many common distributions, like the Normal, Logistic, or Laplace distributions.

Monotone Preferences. When $W + X$ dominates W with respect to first-order stochastic dominance, we say that accepting X is *dominant*. We make no assumptions on the decision maker's behavior except that she accepts dominant gambles.[5] Within the expected utility framework, a preference respects first-order stochastic dominance if and only if it is

[3] See Rabin and Thaler (2001), Cox and Sadiraj (2006), Rubinstein (2006) and Andersen et al. (2018). The idea that narrow framing affects decision makers' risk attitudes goes back to Benartzi and Thaler (1995), Gneezy and Potters (1997) and Thaler et al. (1997).

[4] Formally, g is continuous, and there exists a positive integer n and numbers $-\infty = a_0 < a_1 < \cdots < a_{n-1} < a_n = \infty$ such that g is continuously differentiable on each of the open intervals (a_{i-1}, a_i). Moreover, we require for each i, the limits $\lim_{x \nearrow a_i} g'(x)$ and $\lim_{x \searrow a_i} g'(x)$ exist and are finite.

[5] Recall that a random variable Y *first-order stochastically dominates* another random variable Z if for every $a \in \mathbb{R}$ it holds that $\mathbb{P}[Y \geq a] \geq \mathbb{P}[Z \geq a]$.

represented by an increasing utility function. More generally, a preference is monotone with respect to stochastic dominance if and only if it satisfies two conditions: (i) the preference between any two random variables depends only on their distributions, and (ii) Y is preferred to Z whenever $Y \geq Z$ almost surely.[6] Thus, the assumption that behavior is consistent with first-order stochastic dominance expresses the idea that the decision maker's choice between accepting or rejecting the gamble X at wealth W depends only on the distributions of $W+X$ and W, and that more money is preferred over less.

As discussed in the introduction, consistency with respect to stochastic dominance is a weak assumption satisfied by virtually all preference specifications studied in decision theory and behavioral economics.[7] In fact, the assumption is often satisfied even when choice behavior is not described by means of a single complete and transitive binary relation over wealth distributions, as in models of random expected utility.

Limited Liability. In our analysis the background risk is in general unbounded from below. However, in many contexts it is natural to assume that the decision maker is protected by limited liability, so that her wealth cannot go below a bound $\ell \in \mathbb{R}$. Given a random variable Z, we denote by $(Z)_\ell = \max\{Z, \ell\}$ the variable truncated at ℓ. Under limited liability, comparison between risky prospects boils down to a comparison between their truncated counterparts: The decision maker receives the amount $(W+X)_\ell$ when she accepts the gamble, and $(W)_\ell$ when she rejects it. Accordingly, we say that accepting X is *dominant* if $(W+X)_\ell$ dominates $(W)_\ell$ in first-order stochastic dominance. Thus, we assume the decision maker prefers more wealth than less, and prefers a lower probability of reaching the liability bound. Aside from this, we make no assumptions on what reaching the bound implies for the decision maker.

3 Main Results

Before stating the formal results we introduce two indices for quantifying the magnitude of the background risk and the riskiness of a gamble.

[6] This equivalent formulation is based on the well-known fact that if Y first-order stochastically dominates Z, then there exist two other random variables $\tilde{Y}, \tilde{Z}$ with the same distributions as Y and Z, respectively, and such that $\tilde{Y} \geq \tilde{Z}$ almost surely.

[7] An exception is the choice-acclimating equilibrium of Kőszegi and Rabin (2007), which violates monotonicity for some values of the loss aversion parameter.

Size of the Background Risk Given a background risk W with density g, we define its *exponential size* $\mathrm{S}(W)$ as

$$\mathrm{S}(W) = \left(\sup_a \frac{g'(a)}{g(a)} \right)^{-1},$$

where the supremum is taken over points a where g is differentiable. We say that W is *heavy left-tailed* if $\mathrm{S}(W) > 0$. This restriction includes common parametric distributions such as Logistic or Laplace, but excludes distributions with thin tails such as Normal. As we explain later, our analysis does apply to the Normal distribution, so long as the decision maker is protected by limited liability.

Intuitively, the exponential size is a measure of how likely large losses are, relative to small ones. The larger $\mathrm{S}(W)$, the more slowly the density g increases, and hence the thicker is the left tail of the distribution.[8] More formally, the density $g(-a)$ associated with a loss $a > 0$ must lie above the exponential function $g(0) \cdot \mathrm{e}^{-\frac{a}{\mathrm{S}(W)}}$ and thus cannot vanish faster than exponentially. As an example, the exponential size of a Laplace distribution, i.e. one following a density $g(a) = \frac{1}{2\lambda}\mathrm{e}^{-|a|/\lambda}$, is equal to its parameter λ.

Similar to the standard deviation, the exponential size is positive homogeneous and independent of the location of W (i.e. it satisfies $\mathrm{S}(tW) = t\mathrm{S}(W)$ for $t > 0$ and $\mathrm{S}(W) = \mathrm{S}(W + c)$ for any $c \in \mathbb{R}$).

Riskiness of the Gamble We quantify the riskiness of a gamble in terms of the Aumann-Serrano index. Given a gamble X with positive expectation and positive probability of being negative, Aumann and Serrano (2008) define its riskiness $\mathrm{R}(X)$ as the reciprocal of the level of absolute risk aversion at which a decision maker with CARA expected utility preferences is indifferent between accepting and rejecting X. Formally, $\mathrm{R}(X)$ is defined as the (unique) positive real number solving the equation

$$\mathbb{E}\left[\mathrm{e}^{-\frac{1}{\mathrm{R}(X)}X}\right] = 1\,.$$

Intuitively, a gamble that is assigned a higher index is riskier because it is accepted by a smaller pool of risk-averse CARA decision makers. For a gamble X that pays \$11 and $-$\$10 with equal probability, its riskiness index can be calculated to be $R(X) \approx 110$.

Aumann and Serrano (2008) provide an axiomatic foundation for the index. The riskiness index can also be related to simpler concepts such as the expectation of a gamble and the size

[8] The right tail is less constrained, since for large positive a the ratio $\frac{g'(a)}{g(a)}$ is negative by our assumption that the density $g(a)$ is eventually decreasing. The supremum is not affected by these a.

of its support. In particular, whenever a gamble X is supported in an interval $[-M, M]$ and has expectation $\epsilon > 0$, its riskiness index satisfies $\mathrm{R}(X) \leq \frac{M^2}{\epsilon}$, as we show in Proposition 2 in the Appendix. The index is also positively homogeneous: it satisfies $\mathrm{R}(tX) = t\mathrm{R}(X)$ for every gamble X and every $t > 0$.

The next theorem is the main technical result of the paper. It shows that it is dominant to accept any gamble that has positive expectation and whose riskiness is bounded by the exponential size of the background risk.

Theorem 1. *Let W be a background risk and X a bounded gamble with positive expectation and riskiness $\mathrm{R}(X) \leq \mathrm{S}(W)$. Then $W + X$ first-order stochastically dominates W.*

This result implies that accepting a gamble X is dominant provided it is not excessively risky compared to the size of the background risk W. While the complete proof is provided in §5.1, we illustrate here the main underlying ideas. Given a background risk W distributed according to a cdf G, we define its corresponding *test utility function* u_G, defined as

$$u_G(a) = -G(-a).$$

We show that the effect of the background risk can be understood by studying the choice behavior of an hypothetical expected utility decision maker with utility function u_G. A similar approach is used in the literature on expected utility (Gollier and Schlesinger, 2003; Gollier, 2004). The preference defined by the test utility function is neither globally risk-averse nor risk-loving, but it is monotone. The key observation is that X is dominant if and only if X is accepted under the test utility function at all wealth levels: as we show, accepting X is dominant if and only if

$$\mathbb{E}\left[u_G(X + a)\right] > \mathbb{E}\left[u_G(a)\right] \quad \text{for all } a \in \mathbb{R}. \tag{1}$$

Since X has positive expectation, in order for (1) to hold, the test utility u_G must be not too risk averse. Formally, its Arrow-Pratt coefficient of risk aversion must be below a certain threshold. In the proof, we establish that $\mathrm{R}(X) \leq \mathrm{S}(W)$ is satisfied exactly when the Arrow-Pratt coefficient of u_G is everywhere below the level at which u_G would reject the gamble X. Thus, the inequality $\mathrm{R}(X) \leq \mathrm{S}(W)$ is sufficient, and in fact necessary, for (1) to hold. This is the main step in the proof of the theorem.

By reasoning in terms of the test utility function, we can reformulate a statement about stochastic dominance into a statement about the risk attitude of u_G, and thus apply standard

tools from the literature on expected utility. The argument leading to Theorem 1 is entirely elementary, and can be easily extended, as we discuss in §4.

An important corollary of Theorem 1 is the following:

Corollary 1. *Let X be a bounded gamble with positive expectation. Then:*

1. *If the background risk W is heavy left-tailed, i.e. $\mathrm{S}(W) > 0$, then $W + tX$ first-order stochastically dominates W for all $t > 0$ small enough.*

2. *If the decision maker is protected by limited liability for some liability bound ℓ, then $(W + tX)_\ell$ first-order stochastically dominates $(W)_\ell$ for all $t > 0$ small enough.*

In the case of limited liability, the next result—another corollary of Theorem 1—establishes that a large enough Normal background risk suffices to make any actuarially favorable gamble dominant.

Corollary 2. *Consider a background risk W, distributed Normally with mean μ and standard deviation σ, and a bounded gamble X with positive expectation, riskiness $\mathrm{R}(X)$ and maximum $\max[X]$. Then $(W + X)_\ell$ first-order stochastically dominates $(W)_\ell$ if*

$$\sigma \geq \sqrt{\mathrm{R}(X) \cdot \max\{0,\ \mu - \ell + \max[X]\}}.$$

As we show in the Appendix, similar lower bounds can be derived for other distributions, and are especially simple to calculate for log-concave densities such as the Normal density.

Theorem 1 and its corollaries allow us to provide quantitative estimates for the level of background risk under which accepting a gamble is dominant. To illustrate, consider the gamble X that pays \$11 and $-$\$10 with equal probability and a background risk W having Laplace distribution, i.e. following a density $g(a) = \frac{1}{2\lambda}\mathrm{e}^{-|a|/\lambda}$. The parameter λ coincides with the exponential size of W, and the distribution has standard deviation $\sigma = \sqrt{2}\lambda$. It follows from Theorem 1 that it is dominant to accept the gamble X as long as the standard deviation of W satisfies (recalling that $R(X) \approx 110$)

$$\sigma \geq \sqrt{2}\mathrm{R}(X) \approx \$156.$$

If instead W follows a Logistic distribution with standard deviation σ, then its exponential size is $\mathrm{S}(W) = \frac{\sqrt{3}}{\pi}\sigma$. Thus accepting the gamble is dominant provided

$$\sigma \geq \frac{\pi}{\sqrt{3}}\mathrm{R}(X) \approx \$200.$$

For a decision maker with limited liability bound $\ell = 0$ and Normally distributed W with mean \$100,000, it is dominant to accept X if

$$\sigma \geq \sqrt{\mathrm{R}(X) \cdot 100{,}011} \approx \$3{,}319.$$

Table 1 is constructed from similar calculations.

4 Discussion

Full Support of the Background Risk. In our main result, Theorem 1, the background risk is assumed to have full support. To better interpret this assumption, suppose we were to require W to be bounded. Then, since the minimum of $W + X$ is smaller than that of W (as X takes negative values with positive probability), the random variable $W + X$ would not dominate W, regardless of the distributions of W and X.

This lack of continuity between what can be achieved with bounded versus unbounded background risk is due to the fact for a gamble to be first-order stochastically dominant, it must be attractive under every monotone preference, including preferences that are extremely risk averse. To elaborate, consider a preference over gambles that is defined by assigning to each bounded random variable Z a utility $U(Z)$ equal to the minimum of its support. This preference is infinitely risk averse, and ranks W as strictly better than $W + X$ whenever W is bounded, regardless of how small is the probability that X falls below zero.

Arguably, it is implausible that a decision maker would reject a very favorable small gamble because of the negligible effect it has on the support of her background risk. The assumption that W has unbounded support has the effect of excluding such preferences from stochastic dominance comparisons.[9] A different modelling approach would be to allow for bounded background risk while restricting attention to a suitable family of monotone preferences. We do not pursue this approach here, and instead focus on the standard notion of first-order stochastic dominance. This issue is mostly hypothetical as the unbounded support of the background risk plays no role in the practically relevant case, where the decision maker is protected by some level of limited liability,

[9]The discontinuity between bounded and unbounded support is reminiscent of the discontinuity in equilibrium predictions between finitely and infinitely repeated games. A common view is that an infinite repeated game is a better model for analyzing long-term interactions where players do not assign a special status to the last round of play. Here, the assumption that W has full support rules out preferences that assign special status to the support of a risk.

Heavy Tail and Constant Absolute Risk Aversion. As is well known, under CARA expected utility, a decision maker accepts or rejects a gamble independently of wealth levels, and thus also independently of background risk. This apparent contradiction to our main results is explained by the fact that CARA rules out heavy-tailed distributions if we additionally require expected utilities to be finite.[10]

Due to the fact that infinite expected utilities cannot be compared, there is no contradiction between the assumption that a CARA decision maker rejects a gamble X and our conclusion that she finds it dominant to accept X under a heavy-tailed background risk W. This technical issue notwithstanding, it is worth mentioning that heavy-tailed distributions have a long history in modeling risk and have seen a number of economic applications (see, e.g., Morris and Yildiz, 2019).

Extensions. In the appendix we apply our techniques to two more general settings. In §B we consider choices between two gambles. We show that given two gambles X and Y with $\mathbb{E}[X] > \mathbb{E}[Y]$, for any background risk W with sufficiently heavy tails both on the left and on the right, the resulting distribution of $X + W$ first-order stochastically dominates that of $Y + W$. In §C we consider decision makers whose preferences are monotone with respect to second-order stochastic dominance. This is a stronger assumption that is natural in the study of risk aversion. In this setting we prove a result that is analogous to our main Theorem 1: every actuarially favorable gamble is accepted when the left tail of the background risk is heavy enough. The measure of tail-heaviness is different in this case, requiring less background risk than the first-order stochastic dominance case.

5 Proofs of Main Results

5.1 Proof of Theorem 1

Let $s = \mathrm{S}(W)$, and let G denote the c.d.f. of W. The result is vacuous when $s = 0$, so we focus below on $s > 0$. A gamble X has positive expectation and Aumann-Serrano index $\mathrm{R}(X) \leq s$ if and only if it is accepted by a decision maker with constant absolute risk

[10] To be specific, consider a CARA decision maker with risk aversion level α, together with a gamble X that she *rejects* without any background risk. Then the riskiness index satisfies $\mathrm{R}(X) \geq 1/\alpha$. By Theorem 1, we can find a background risk W with size $\mathrm{S}(W) \geq \mathrm{R}(X)$, such that $W + X$ first-order stochastically dominates W. Nonetheless, since $\mathrm{S}(W) \geq \mathrm{R}(X) \geq 1/\alpha$, it can be shown that the CARA expected utilities of W and $W + X$ are both $-\infty$. For example, if W has Laplace distribution with density $g(x) = \frac{\beta}{2}\mathrm{e}^{-\beta|x|}$, then $\mathrm{S}(W) \geq 1/\alpha$ requires $\beta \leq \alpha$. Denoting by $u(x) = -\mathrm{e}^{-\alpha x}$ the CARA utility function, it is then easy to see $\int_{-\infty}^{\infty} u(x)g(x)\,\mathrm{d}x = -\infty$.

aversion of $1/s$, i.e. if and only if

$$\mathbb{E}\left[-\mathrm{e}^{-\frac{1}{s}X}\right] \geq -1. \tag{2}$$

On the other hand, accepting the gamble X is dominant if

$$\mathbb{P}\left[W + X \leq a\right] \leq \mathbb{P}\left[W \leq a\right] = G(a) \text{ for all } a \in \mathbb{R}\,.$$

Since the gamble X and the background risk W are independent, we have that

$$\mathbb{P}\left[W + X \leq a\right] = \mathbb{P}\left[W \leq a - X\right] = \mathbb{E}\left[G(a - X)\right].$$

Thus, accepting X is dominant if and only if

$$\mathbb{E}\left[G(a - X)\right] \leq G(a) \text{ for all } a \in \mathbb{R}\,. \tag{3}$$

Inequality (3) can be interpreted as saying that a decision maker with expected utility preferences and *test utility function* $u_G(x) = -G(-x)$ accepts the gamble X at every wealth level:

$$\mathbb{E}\left[u_G(X + a)\right] \geq u_G(a) \text{ for all } a \in \mathbb{R}\,. \tag{4}$$

This analogy is useful as it allows us to understand stochastic dominance through the behavior of a hypothetical expected utility decision maker whose utility coincides with u_G. The preference defined by G is in general neither globally risk-averse nor risk-loving. Since the expectation of X is positive, equation (4) is equivalent to imposing that the utility function u_G is "not too risk-averse." Below we formalize this intuition.

As shown by (2) and (4) above, to say that it is dominant to accept every gamble X with positive expectation and riskiness $\mathrm{R}(X) \leq s$ is equivalent to saying that for every gamble X,

$$\mathbb{E}\left[-\mathrm{e}^{-\frac{1}{s}X}\right] \geq -1 \implies \mathbb{E}\left[u_G(X + a)\right] \geq u_G(a) \text{ for all } a \in \mathbb{R}. \tag{5}$$

That is, any gamble X that is accepted by a decision maker with (risk-averse) CARA utility $U_s(a) = -\mathrm{e}^{-a/s}$ is also accepted by a decision maker with utility u_G, at all wealth levels. In other terms, we require u_G to be globally less risk-averse than U_s in the sense of Arrow-Pratt.

To show that this holds as long as $\mathbb{E}\left[X\right] > 0$ and $\mathrm{R}(X) \leq \mathrm{S}(W)$, let us first suppose for simplicity that the density g is everywhere differentiable. In this case, since the exponential size of W is s, we have $g'(a)/g(a) \leq 1/s$ for every a. Thus, the Arrow-Pratt index of absolute

risk aversion is everywhere lower for u_G than for U_s:

$$-\frac{u_G''(a)}{u_G'(a)} = \frac{G''(-a)}{G'(-a)} = \frac{g'(-a)}{g(-a)} \leq \frac{1}{s} = -\frac{U_s''(a)}{U_s'(a)} \quad \text{for all } a \in \mathbb{R}\,.$$

It follows that u_G is indeed less risk-averse than U_s.

For the general case where g is only piece-wise continuously differentiable, note that $U_s(a) = -\mathrm{e}^{-a/s}$ is a strictly increasing function, enabling us to write $u_G(a) = \phi\left(U_s(a)\right) = \phi(-\mathrm{e}^{-a/s})$ for some increasing function ϕ defined on $(-\infty, 0)$. Our goal is to show that ϕ is convex, so that u_G is globally less risk-averse than U_s. Note that

$$g(-a) = u_G'(a) = \phi'(-\mathrm{e}^{-\frac{a}{s}}) \cdot \frac{1}{s}\mathrm{e}^{-\frac{a}{s}}.$$

Thus ϕ is convex if and only if ϕ' is an increasing function, which in turn is equivalent to $g(-a) \cdot \mathrm{e}^{a/s}$ being increasing in a, or $g(a) \cdot \mathrm{e}^{-a/s}$ being decreasing in a. Since g is piece-wise continuously differentiable, we know that even if g is not differentiable at some point a, the left and right derivatives do exist, and they also satisfy $g'(a)/g(a) \leq 1/s$. From this it follows that $g(a) \cdot \mathrm{e}^{-a/s}$ is indeed decreasing, concluding the proof of Theorem 1.

We note that this proof establishes a stronger statement: given any background risk W and any number $s \geq 0$, it is dominant to accept every gamble X with positive expectation and riskiness $\mathrm{R}(X) \leq s$ *if and only if* the background risk satisfies $\mathrm{S}(W) \geq s$. In particular, the bound $\mathrm{R}(X) \leq \mathrm{S}(W)$ in Theorem 1 cannot be improved.

5.2 Proof of Corollaries 1 and 2

Proof of Corollary 1. Suppose W is heavy left-tailed, i.e. it satisfies $\mathrm{S}(W) > 0$. As shown by Aumann and Serrano (2008), the riskiness index R is positive homogeneous, i.e. it satisfies $\mathrm{R}(tX) = t\mathrm{R}(X)$ for all $t > 0$. Thus, the riskiness of tX is lower than the exponential size of W for all t small enough. It then follows from Theorem 1 that accepting tX is dominant.

Now suppose the decision maker is protected by limited liability. We prove the following analogue of Theorem 1, which will imply this part of Corollary 1 as well as Corollary 2.

Proposition 1. *Suppose the decision maker is protected by limited liability bound ℓ. Then under any background risk W, it is dominant to accept every gamble X with positive expectation and riskiness*

$$\mathrm{R}(X) \leq \left(\max\left\{0, \sup_{a \geq \ell - \max[X]} \frac{g'(a)}{g(a)}\right\}\right)^{-1}.$$

Compared to the definition of $\mathrm{S}(W)$, the supremum on the right-hand side above only considers those a with $a \geq \ell - \max[X]$. Because of this change, the supremum can now be negative if the density $g(a)$ is decreasing for $a \geq \ell - \max[X]$. If that is the case then Proposition 1 asserts that it is dominant to accept every gamble X with positive expectation.

We now show that Proposition 1 implies the limited liability case of Corollary 1. Since we assumed g to be strictly positive and piece-wise continuously differentiable, the ratio $g'(a)/g(a)$ is bounded on every compact interval. Moreover, as g is eventually decreasing, this ratio is bounded from above for points $a \geq \ell - 1$. Thus, for $t > 0$ small enough, $\mathrm{R}(tX) = t \cdot \mathrm{R}(X)$ is close to zero, while

$$\sup_{a \geq \ell - \max[tX]} \frac{g'(a)}{g(a)} \leq \sup_{a \geq \ell - 1} \frac{g'(a)}{g(a)}$$

is bounded from above. Applying Proposition 1 to the gamble tX yields Corollary 1. □

Proof of Proposition 1. We follow the proof of Theorem 1 in §5.1, and explain the necessary modifications. Under limited liability, accepting the gamble X is dominant if

$$\mathbb{P}[W + X \leq a] \leq \mathbb{P}[W \leq a] = G(a) \text{ for all } a \geq \ell\,.$$

Thus, instead of (3), we only need to check

$$\mathbb{E}[G(a - X)] \leq G(a) \text{ for all } a \geq \ell\,. \tag{6}$$

If $\sup_{a \geq \ell - \max[X]} \frac{g'(a)}{g(a)} \leq 0$, then $g'(a) \leq 0$ for all $a \geq \ell - \max[X]$ and thus $G(a)$ is increasing and concave for such a. In this case (6) holds immediately by Jensen's inequality: $\mathbb{E}[G(a - X)] \leq G(a - \mathbb{E}[X]) \leq G(a)$ whenever $\mathbb{E}[X] \geq 0$.

Otherwise let $s \geq 0$ denote $\left(\sup_{a \geq \ell - \max[X]} \frac{g'(a)}{g(a)}\right)^{-1}$, and consider any gamble X with positive expectation and $\mathrm{R}(X) \leq s$. We can without loss assume that $s > 0$. Then $\mathrm{R}(X) \leq s$ implies $\mathbb{E}\left[\mathrm{e}^{-X/s}\right] \leq 1$. So the CARA decision maker with utility function $U(a) = \mathrm{e}^{a/s}$ would reject the gamble $-X$.

On the other hand, the definition of s implies that the function G has weakly higher Arrow-Pratt index than U on the interval $[\ell - \max[X], \infty)$, so G is a concave transformation of U *on this interval*. Since the comparison between a and $a - X$ in (6) only involves wealth levels that are above $\ell - \max[X]$, we deduce that rejection of $-X$ by the utility function U implies rejection by G. Hence the proposition. □

Proof of Corollary 2. By Proposition 1, accepting X is dominant under limited liability if

$$\mathrm{R}(X) \leq \left(\max \left\{ 0, \sup_{a \geq \ell - \max[X]} \frac{g'(a)}{g(a)} \right\} \right)^{-1},$$

where $g(a) = \frac{1}{\sqrt{2\pi}\sigma} \mathrm{e}^{-\frac{(a-\mu)^2}{2\sigma^2}}$ denotes the density of the normal background risk. It is well known that g is log-concave, so that $g'(a)/g(a)$ is decreasing in a. Thus

$$\sup_{a \geq \ell - \max[X]} \frac{g'(a)}{g(a)} = \frac{g'(\ell - \max[X])}{g(\ell - \max[X])} = \frac{\mu - \ell + \max[X]}{\sigma^2}.$$

Thus the normal background risk makes X dominant whenever $\mathrm{R}(X) \leq \frac{\sigma^2}{\max\{0,\ \mu - \ell + \max[X]\}}$, or equivalently $\sigma \geq \sqrt{\mathrm{R}(X) \cdot \max\{0,\ \mu - \ell + \max[X]\}}$. □

References

S. Andersen, J. C. Cox, G. W. Harrison, M. I. Lau, E. E. Rutström, and V. Sadiraj. Asset integration and attitudes toward risk: Theory and evidence. *The Review of Economics and Statistics*, 100(5):816–830, 2018.

A. Ang, G. Bekaert, and J. Liu. Why stocks may disappoint. *Journal of Financial Economics*, 76:471–508, 2005.

K. J. Arrow. *Essays in the theory of risk-bearing*. Markham Publishing Company, 1970.

R. J. Aumann and R. Serrano. An economic index of riskiness. *Journal of Political Economy*, 116(5):810–836, 2008.

N. Barberis, M. Huang, and R. H. Thaler. Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American economic review*, 96(4): 1069–1090, 2006.

C. Bardgett, E. Gourier, and M. Leippold. Inferring volatility dynamics and risk premia from the s&p 500 and vix markets. *Journal of Financial Economics*, 131(3):593–618, 2019.

S. Benartzi and R. H. Thaler. Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, 110:73–92, 1995.

J. C. Cox and V. Sadiraj. Small- and large-stakes risk aversion: implications of concavity calibration for decision theory. *Games and Economic Behavior*, 56:45–60, 2006.

U. Gneezy and J. Potters. An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics*, 112(2):631–645, 1997.

C. Gollier. The economics of risk and time. *MIT Press*, 2004.

C. Gollier and H. Schlesinger. Preserving preference rankings under background risk. *Economics Letters*, 80(3):337–341, 2003.

F. Gul. A theory of disappointment aversion. *Econometrica*, 59(3):667–686, 1991.

B. Hansson. Risk aversion as a problem of conjoint measurement. In P. Gärdenfors and N.-E. Sahlin, editors, *Decisions, Probability and Utility*. Cambridge University Press, 1988.

D. Kahneman and A. Tversky. Prospect theory: an analysis of decision under risk. *Econometrica*, 47:263–291, 1979.

M. W. Khaw, Z. Li, and M. Woodford. Cognitive imprecision and small-stakes risk aversion. *Review of Economic Studies*, 2020. Forthcoming.

B. Kőszegi and M. Rabin. Reference-dependent risk attitudes. *American Economic Review*, 97(4):1047–1073, 2007.

B. Kőszegi and M. Rabin. Reference-dependent consumption plans. *American Economic Review*, 99(3):909–36, 2009.

S. Morris and M. Yildiz. Crises: Equilibrium shifts and large shocks. *American Economic Review*, 109(8):2823–54, 2019.

L. Pomatto, P. Strack, and O. Tamuz. Stochastic dominance under independent noise. *Journal of Political Economy*, 128(5):1877–1900, 2020.

J. W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1-2):122–136, 1964.

J. Quiggin. A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3:323–343, 1982.

M. Rabin. Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68(5):1281–1292, 2000.

M. Rabin and R. H. Thaler. Anomalies: Risk aversion. *Journal of Economic Perspectives*, 15(1):219–232, 2001.

A. Rubinstein. Dilemmas of an economic theorist. *Econometrica*, 74(4):865–883, 2006.

Z. Safra and U. Segal. Calibration results for non-expected utility theories. *Econometrica*, 76(5):1143–1166, 2008.

T. Sarver. Dynamic mixture-averse preferences. *Econometrica*, 86(4):1347–1382, 2018.

U. Segal and A. Spivak. First order versus second order risk aversion. *Journal of Economic Theory*, 51:111–125, 1990.

C. Tarsney. Exceeding expectations: Stochastic dominance as a general decision theory. *arXiv preprint arXiv:1807.10895*, 2018.

R. H. Thaler, A. Tversky, D. Kahneman, and A. Scharwtz. The effect of myopia and loss aversion on risk taking: an experimental test. *Quarterly Journal of Economics*, 112(2): 647–661, 1997.

A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323, 1992.

E. Zambrano. Risk attitudes over small and large stakes recalibrated. *Economics Letters*, 187:108899, 2020.

Background Risk and Small-Stakes Risk Aversion

Xiaosheng Mu Luciano Pomatto Philipp Strack Omer Tamuz

Online Appendix

A Background Risk for Cumulative Prospect Theory Preferences

Table 2 shows the levels of background risk needed to make a decision maker with cumulative prospect theory (CPT) preferences to accept various gambles. The specific CPT preference we consider has gain/loss probability weighting functions $w^+(p) = \frac{p^\gamma}{(p^\gamma+(1-p)^\gamma)^{1/\gamma}}$, $w^-(p) = \frac{p^\delta}{(p^\delta+(1-p)^\delta)^{1/\delta}}$ with $\gamma = 0.61, \delta = 0.69$, loss aversion parameter $\lambda = 2.25$ and value function $v(x) = x^{0.88}$ for $x \geq 0$ and $v(x) = -\lambda(-x)^{0.88}$ for $x < 0$. These parameter values are taken from Tversky and Kahneman (1992, pages 309–312). Given this choice of parameters, the table is constructed by computing numerically the utility of each gamble as a function of the standard deviation of the background risk.

Gamble	**StDeviation of Background Risk: σ**		
Gain/Loss	Laplace	Logistic	Normal
\$11/\$10	$\sigma \geq$ \$61	$\sigma \geq$ \$46	$\sigma \geq$ \$43
\$55/\$50	$\sigma \geq$ \$306	$\sigma \geq$ \$230	$\sigma \geq$ \$217
\$110/\$100	$\sigma \geq$ \$611	$\sigma \geq$ \$460	$\sigma \geq$ \$434
\$550/\$500	$\sigma \geq$ \$3057	$\sigma \geq$ \$2299	$\sigma \geq$ \$2169
\$1000/\$1100	$\sigma \geq$ \$6114	$\sigma \geq$ \$4598	$\sigma \geq$ \$4338

Table 2: Standard deviation of background risk sufficient for a CPT decision maker to accept various fifty-fifty gambles under different distributional assumptions on the background risk.

B Choice Between Two Gambles

In this section, we extend the analysis to situations where the decision maker faces a choice between two bounded gambles X and Y that have distinct distributions F_X and F_Y. We say it is *dominant to choose X over Y* under background risk W, if $W + X$ first-order stochastically dominates $W + Y$. A result similar to Theorem 1 can be obtained if we

consider background risks with heavy tails both on the left and on the right. For this we define the *two-sided exponential size*

$$\mathrm{S}^*(W) = \left(\sup_a \left| \frac{g'(a)}{g(a)} \right| \right)^{-1},$$

which is equal to $\min\{\mathrm{S}(W), \mathrm{S}(-W)\}$. Then we have:

Theorem 2. *The following are equivalent:*

(i) $\mathbb{E}[X] > \mathbb{E}[Y]$*;*

(ii) there exists $s \in (0, \infty)$ *such that under any background risk* W *with* $\mathrm{S}^*(W) \geq s$*, choosing* X *over* Y *is dominant.*

Proof of Theorem 2. We first show (ii) implies (i). Given any finite s, we can choose W to have a Laplace distribution with sufficiently large variance. Then W satisfies $\mathrm{S}^*(W) \geq s$, and by assumption $W + X$ must first-order stochastically dominate $W + Y$. Since such a W has finite expectation, we have $\mathbb{E}[W + X] \geq \mathbb{E}[W + Y]$, which implies $\mathbb{E}[X] \geq \mathbb{E}[Y]$. The inequality is in fact strict, for otherwise $W + X$ would have the same distribution as $W + Y$, and X would have the same distribution as Y. This last claim can be proved by considering the moment generating function in a neighborhood of 0. Since $\mathbb{E}[\mathrm{e}^{tW}]$ is finite for t close to 0, both $\mathbb{E}[\mathrm{e}^{t(W+X)}]$ and $\mathbb{E}[\mathrm{e}^{t(W+Y)}]$ are finite and are equal. It follows that $\mathbb{E}[\mathrm{e}^{tX}] = \mathbb{E}[\mathrm{e}^{tY}]$ for t in a neighborhood of 0, which implies X and Y have the same distribution.

To prove (i) implies (ii), we assume $\mathbb{E}[X] > \mathbb{E}[Y]$ and take s to be a large positive number (to be determined later). Consider any background risk W with $\mathrm{S}^*(W) \geq s$, i.e. the density g satisfies $|g'(a)/g(a)| \leq 1/s$ for all a. Let $h(a) = \ln g(a)$, then we can rewrite the condition as

$$|h'(a)| \leq \frac{1}{s} \text{ for all } a \in \mathbb{R}.$$

We now use this to show $\mathbb{P}[W + Y \leq a] \geq \mathbb{P}[W + X \leq a]$ for all a. Since W is independent from both X and Y, integration by parts shows this comparison is equivalent to

$$\int_{-M}^{M} g(a - z) \cdot F_Y(z)\, \mathrm{d}z \geq \int_{-M}^{M} g(a - z) \cdot F_X(z)\, \mathrm{d}z,$$

where M is a large number such that $[-M, M]$ contains the support of both X and Y. This

in turn is equivalent to

$$\int_{-M}^{M} \mathrm{e}^{h(a-z)} \cdot (F_Y(z) - F_X(z))\,\mathrm{d}z \geq 0.$$

Dividing both sides by $\mathrm{e}^{h(a)}$, we just need to show that for all a

$$\int_{-M}^{M} \mathrm{e}^{h(a-z)-h(a)} \cdot (F_Y(z) - F_X(z))\,\mathrm{d}z \geq 0.$$

Observe that since $|h'|$ is bounded above by $1/s$, we have $|h(a-z) - h(a)| \leq M/s$ for all $a \in \mathbb{R}$ and all $z \in [-M, M]$. Thus if s is chosen to be sufficiently large, then the above integral converges, uniformly across a, to the integral $\int_{-M}^{M}(F_Y(z) - F_X(z))\,\mathrm{d}z$. Since this limit integral evaluates, by integration by parts, to $\mathbb{E}[X] - \mathbb{E}[Y] > 0$, the result follows. □

If we only know that the background risk has a heavy left tail (as in Theorem 1), then the condition $\mathbb{E}[X] > \mathbb{E}[Y]$ is no longer sufficient to guarantee the dominance of X. Below we derive the suitable condition in this case. We say that X *strongly dominates* Y *in the convex order*, if $\max[X] > \max[Y]$ and

$$\int_{a}^{\infty} (F_Y(z) - F_X(z))\,\mathrm{d}z > 0 \text{ for all } a < \max[X]. \tag{7}$$

In particular, this requires $\mathbb{E}[X] > \mathbb{E}[Y]$ in the limit $a \to -\infty$.

To interpret this condition, note that X dominates Y in the convex order if and only if $-Y$ dominates $-X$ in second-order stochastic dominance. In other terms, X can be obtained from Y by a combination of mean-preserving spreads and right-ward mass shifts. Conversely, if X is obtained from Y by replacing *each* realization y of Y by a gamble with expectation *strictly greater* than y, then X *strongly* dominates Y in the convex order. This is a natural generalization of the case studied in the main text, where Y is a constant and X is any gamble with a higher expectation.

Theorem 3. *Suppose* $\max[X] \neq \max[Y]$. *Then the following are equivalent:*

(i) X *strongly dominates* Y *in the convex order;*

(ii) there exists $s \in (0, \infty)$ *such that under any background risk* W *with* $\mathrm{S}(W) \geq s$, *choosing* X *over* Y *is dominant.*

Proof of Theorem 3. As in the proof of Theorem 1, choosing X over Y is dominant if and only if

$$\mathbb{E}[G(a-X)] \leq \mathbb{E}[G(a-Y)] \text{ for all } a \in \mathbb{R}\,.$$

Since we want this to hold for all background risks G with exponential size $\geq s$, and since the exponential size is translation-invariant, it is without loss to restrict to the case of $a = 0$. That is, we seek to understand the conditions under which

$$\mathbb{E}[G(-X)] \leq \mathbb{E}[G(-Y)] \text{ for all } G \text{ with exponential size } \geq s.$$

As before, let $U(a) = \mathrm{e}^{\frac{a}{s}}$ denote a risk-loving CARA utility function. Then G has exponential size at least s if and only if $G(a) = \phi(U(a))$ for some increasing concave function ϕ.[11] Thus, the above comparison can be rewritten as

$$\mathbb{E}\left[\phi\left(\mathrm{e}^{\frac{-X}{s}}\right)\right] \leq \mathbb{E}\left[\phi\left(\mathrm{e}^{\frac{-Y}{s}}\right)\right] \text{ for all increasing concave functions } \phi.$$

In other terms, the random variable $\tilde{Y} = \mathrm{e}^{\frac{-Y}{s}}$ should dominate $\tilde{X} = \mathrm{e}^{\frac{-X}{s}}$ with respect to second-order stochastic dominance.

Let $\tilde{F}_X$ and $\tilde{F}_Y$ denote the c.d.f. of $\tilde{X}$ and $\tilde{Y}$, respectively. Then second-order stochastic dominance holds if and only if (noting that $\tilde{X}$ and $\tilde{Y}$ are both supported on $\mathbb{R}_+$):

$$\int_0^c (\tilde{F}_X(t) - \tilde{F}_Y(t))\,\mathrm{d}t \geq 0 \text{ for all } c > 0.$$

If we write $t = \mathrm{e}^{-\frac{z}{s}}$, then $\tilde{F}_X(t) = 1 - F_X(z)$, $\tilde{F}_Y(t) = 1 - F_Y(z)$. Changing variables in the above integral, and denoting $a = -s\ln(c)$, we obtain the following equivalent condition (modulo a factor of $1/s$):

$$\int_a^\infty (F_Y(z) - F_X(z)) \cdot \mathrm{e}^{-\frac{z}{s}}\,\mathrm{d}z \geq 0 \text{ for all } a \in \mathbb{R}. \tag{8}$$

Below we show that when the maxima of X and Y are different, the above condition holds for some positive s if and only if X strongly dominates Y in the convex order.

In one direction, suppose $\max[X] > \max[Y]$ and (7) holds. Then intuitively (8) would also

[11] To be fully rigorous, we also need $g(a) = \phi'(\mathrm{e}^{a/s}) \cdot \frac{1}{s}\mathrm{e}^{a/s}$ to be strictly positive, continuously differentiable, and eventually decreasing. These additional restrictions on ϕ do not affect the subsequent analysis because on any compact domain, any increasing concave function can be uniformly approximated by another increasing concave function with these additional properties.

hold if s is large, in which case the integrand $(F_Y(z) - F_X(z)) \cdot \mathrm{e}^{-\frac{z}{s}}$ is close to $F_Y(z) - F_X(z)$. This can be formalized by observing that we only need to prove (8) for a in the compact interval $\min[X] \leq a \leq \max[Y]$. As $s \to \infty$ the integral $\int_a^\infty (F_Y(z) - F_X(z)) \cdot \mathrm{e}^{-\frac{z}{s}} \, \mathrm{d}z$ converges uniformly to $\int_a^\infty (F_Y(z) - F_X(z)) \, \mathrm{d}z$ on this interval. Since this limit is a continuous function in a and strictly positive on this interval, it is bounded away from 0. Thus by uniform convergence, there exists some large s such that (8) holds.

For the converse, suppose (8) holds for some s. Then there cannot exist some a with $F_Y(a) < 1 = F_X(a)$, since otherwise (8) fails at this point a. It follows that $\max[X] \geq \max[Y]$, and the inequality is in fact strict by the assumption that $\max[X] \neq \max[Y]$. As a result, $F_Y(z) - F_X(z)$ is strictly positive for $z \in [\max[Y], \ \max[X])$, and (8) holds with strict inequality for a in the same interval. We now use this to prove (7). Observe that

$$\begin{aligned}
&\int_a^\infty (F_Y(z) - F_X(z)) \, \mathrm{d}z \\
&= \mathrm{e}^{\frac{a}{s}} \int_a^\infty (F_Y(z) - F_X(z)) \cdot \mathrm{e}^{-\frac{z}{s}} \, \mathrm{d}z + \int_a^\infty \left(\frac{\mathrm{e}^{\frac{c}{s}}}{s} \cdot \int_c^\infty (F_Y(z) - F_X(z)) \cdot \mathrm{e}^{-\frac{z}{s}} \, \mathrm{d}z \right) \mathrm{d}c.
\end{aligned}$$

So from (8), we must have $\int_a^\infty (F_Y(z) - F_X(z)) \, \mathrm{d}z \geq 0$. Moreover, the inequality is strict because in the double integral on the RHS above, the term $\int_c^\infty (F_Y(z) - F_X(z)) \cdot \mathrm{e}^{-\frac{z}{s}} \, \mathrm{d}z$ is strictly positive for any $c \in [\max[Y], \ \max[X])$. For any $a < \max[X]$, the mass of such $c > a$ is strictly positive. Hence (7) holds with strict inequality, completing the proof. □

C Second-Order Stochastic Dominance

Our analysis can also be extended to the smaller class of risk-averse preferences. We say that accepting X is *dominant for a risk-averse decision maker* if $W + X$ dominates W with respect to second-order stochastic dominance. We also introduce a modified version of the exponential size: for any background risk W with c.d.f. G, let

$$\mathrm{S}_2(W) = \left(\sup_{a \in \mathbb{R}} \frac{g(a)}{G(a)} \right)^{-1}.$$

It is easy to show that $\mathrm{S}_2(W) \geq \mathrm{S}(W)$.[12]

[12] If $\mathrm{S}(W) = 0$ then the result is trivial. If instead $\mathrm{S}(W) > 0$, then we have the inequality $g(x) \geq g(y) \cdot \mathrm{e}^{\frac{x-y}{\mathrm{S}(W)}}$. Note that $G(y) = \int_{-\infty}^y g(x) \, \mathrm{d}x \to 0$ as $y \to -\infty$. Using the previous inequality, we deduce that $g(y) \to 0$ as $y \to -\infty$. Hence, for each a, it holds that $\frac{g(a)}{G(a)} = \frac{\int_{-\infty}^a g'(x) \, \mathrm{d}x}{\int_{-\infty}^a g(x) \, \mathrm{d}x} \leq \sup_x \frac{g'(x)}{g(x)} = \frac{1}{\mathrm{S}(W)}$. As a result,

Theorem 4. *Under any given background risk W with finite expectation, it is dominant for a risk-averse decision maker to accept every gamble X with positive expectation and riskiness $\mathrm{R}(X) \leq \mathrm{S}_2(W)$.*

Proof of Theorem 4. Let $s = \mathrm{S}_2(W)$ and without loss focus on $s > 0$. By a well-known characterization of second-order stochastic dominance, it is dominant to accept X if and only if

$$\int_{-\infty}^{a} \mathbb{P}\left[W + X \leq t\right] \mathrm{d}t \leq \int_{-\infty}^{a} \mathbb{P}\left[W \leq t\right] \mathrm{d}t \text{ for all } a \in \mathbb{R}. \tag{9}$$

That the integrals in (9) are finite follows from the fact that W and $W + X$ have finite expectations. By Tonelli's Theorem, the quantity $\int_{-\infty}^{a} \mathbb{P}\left[W + X \leq t\right] \mathrm{d}t$ is equal to

$$\int_{-\infty}^{a} \mathbb{E}\left[G(t - X)\right] \mathrm{d}t = \mathbb{E}\left[\int_{-\infty}^{a} G(t - X)\, \mathrm{d}t\right] = \mathbb{E}\left[\int_{-\infty}^{a-X} G(t)\, \mathrm{d}t\right].$$

Hence, it is second-order dominant to accept a gamble X if and only if for every $a \in \mathbb{R}$

$$\mathbb{E}\left[u_G(a - X)\right] \leq u_G(a),$$

where $u_G(a) = \int_{-\infty}^{a} G(t)\, \mathrm{d}t$. Therefore, as in the proof of Theorem 1, we obtain that accepting X is dominant if

$$\mathbb{E}\left[\mathrm{e}^{-\frac{1}{s}X}\right] \leq 1 \implies \mathbb{E}\left[u_G(a - X)\right] \leq u_G(a) \text{ for all } a \in \mathbb{R}. \tag{10}$$

Equation (10) holds whenever u_G is globally more risk-averse than the CARA utility function $U(a) = \mathrm{e}^{\frac{a}{s}}$. The Arrow-Pratt index for u_G is $-g(a)/G(a)$, which by assumption is weakly larger than $-1/s$, the Arrow-Pratt index for U. Thus u_G is indeed more risk-averse than U, concluding the proof. □

D Additional Results

Proposition 2. *For any gamble X that is supported on $[-M, M]$ and has expectation $\epsilon > 0$, its riskiness index satisfies $\mathrm{R}(X) \leq \frac{M^2}{\epsilon}$.*

Proof of Proposition 2. Let $\lambda = \frac{\epsilon}{M^2}$. We first show that $\mathbb{E}[\mathrm{e}^{-\lambda X}] \leq 1$. Indeed, since $\epsilon = \mathbb{E}[X] \leq M$, we have $\lambda \leq \frac{1}{M}$. As $X \in [-M, M]$ with probability one, we have $-\lambda X \in [-1, 1]$.

$\mathrm{S}_2(W) \geq \mathrm{S}(W)$ again holds.

In this range, it always holds that $\mathrm{e}^{-\lambda X} \leq 1 - \lambda X + (\lambda X)^2$. Hence $\mathbb{E}[\mathrm{e}^{-\lambda X}] \leq 1 - \lambda \mathbb{E}[X] + \lambda^2 \mathbb{E}[X^2] \leq 1 - \lambda\epsilon + \lambda^2 M^2 = 1$.

Now consider the function $f(a) = \mathbb{E}[\mathrm{e}^{-aX}]$, defined for $a \geq 0$. It is easy to see that $f(0) = 1$ and f is strictly convex. Thus, $\frac{1}{\mathrm{R}(X)}$ is the unique number $c > 0$ such that $f(c) = 1$. Since we just proved that $f(\lambda) \leq 1$, convexity implies $c \geq \lambda$. In other words $\frac{1}{\mathrm{R}(X)} \geq \frac{\epsilon}{M^2}$. $\square$