Probability Pricing*

Eduardo Dávila † Cecilia Parlatore ‡ Ansgar Walther §

August 2026

Abstract

This paper develops probability pricing, extending cash-flow pricing to quantify the value of changes in probabilities. We show that the value of any marginal change in probabilities can be expressed as a standard asset-pricing formula with hypothetical cash flows derived from changes in the survival function. This equivalence between probability and cash-flow valuation allows us to construct hedging strategies and systematically decompose individual and aggregate valuations. We illustrate the usefulness of probability pricing in two applications: (i) valuing changes in the distribution of global temperature, which yields country-specific hedging strategies and predictions for which temperature markets should be more valuable and active, and (ii) revisiting the question of the value of public information, showing that the direct value of a more informative signal depends on whether it compresses or disperses the distribution of signal realizations.

JEL Classification: G12, D81, G14

Keywords: probability pricing; cash-flow pricing; hedging; value of information; welfare; climate change

*We would like to thank Andy Atkeson, Ravi Bansal, Philip Bond, John Campbell, Daniel Chen, Doug Diamond, John Geanakoplos, Sergei Glebkin, Piero Gottardi, Lars Hansen, Zhiguo He, John Heaton, Leonid Kogan, Guido Lorenzoni, Andrey Malenko, Alessandro Pavan, Adriano Rampini, and Marzena Rostek for helpful comments and discussions. We also would like to thank seminar participants at the Chicago-Minnesota Theory Conference, Essex-UCL Workshop, Rochester, McGill, UCLA, MIT, Chicago, Stanford, EIEF, LSE, Boston College, Frankfurt School, Harvard, Duke, Finance Theory Group meeting at UT Austin, Yale, SFS Cavalcade, St. Louis Fed, and Boston Fed. Abdelrahman Hassanein provided excellent research assistance. This material is based upon work supported by the National Science Foundation under Grant DGE-2237790.

†Yale University and NBER. Email: eduardo.davila@yale.edu

‡New York University, Stern School of Business, NBER, and CEPR. Email: cparlato@stern.nyu.edu

§University of Oxford and CEPR. Email: ansgar.walther@gmail.com

1 Introduction

Many economic problems involve valuing changes in probabilities. How much is an individual willing to pay to reduce the probability of a disaster? How much should a firm value access to a more precise technology? What are the private and social values of better information?[1] Despite the importance of these questions, the general tools for valuing changes in probabilities are less developed than those for valuing cash flows.

In this paper, we introduce *probability pricing*, an approach to value changes in the probabilities of different states. We initially develop probability pricing in a two-date environment in which the terminal consumption of an expected utility agent depends on a random state. We consider an arbitrary marginal perturbation to the probability distribution of the state, and define the *probability price* of that perturbation as the agent's valuation/willingness-to-pay for it.

Our main result shows that an agent's valuation for a marginal change in probabilities is equivalent to the agent's valuation for a hypothetical asset whose cash flows are the state-by-state consumption equivalent of the probability perturbation. That is, we establish a formal equivalence between changes in probabilities (probability pricing) and changes in consumption (cash-flow pricing). In each state, the equivalent cash flow is the product of (i) the change in the survival function, normalized by the density, and (ii) the sensitivity of consumption to the state.[2] The survival function is the relevant object because a perturbation that moves probability mass to the right (left) of a given state adds (removes) the marginal utility of consumption at that state. Importantly, equivalent cash flows depend only on the perturbation and on the consumption profile: preferences enter the probability price exclusively through the state prices at which cash flows are valued.

This equivalence contrasts with the way existing work values changes in probabilities or information. Prominent examples include Lucas (1987) on the cost of aggregate fluctuations, Barro (2009) and Martin and Pindyck (2015) on the probability of rare disasters, and Hirshleifer (1971), Morris and Shin (2002), and Pavan, Sundaresan, and Vives (2022) on the value of information. The standard approach solves a model with parametric utilities and distributions and then evaluates welfare directly as a function of the parameter of interest, such as the variance of consumption, the probability of a disaster, or the precision

[1] Any change in information, whether private or public, ultimately amounts to a change in the probabilities that agents assign to the different states of the world.

[2] The survival function of a probability distribution, $1 - F(s)$, is the complement of the cumulative distribution function, $F(s)$.

of a signal. Our approach is nonparametric, as it applies to general preferences, general distributions, and arbitrary perturbations of probabilities. It also expresses the value of any such perturbation in the language of asset pricing.

The equivalence between probability and cash-flow pricing also has two practical uses that have not been exploited before, both resting on the consumption-equivalent cash flows it produces. First, the equivalent cash flows define a hedging strategy. An investor who wants to be immunized against an (initially unanticipated) change in the probabilities of different scenarios can use them to identify the state-contingent payoffs that a hedging portfolio must replicate, and the probability price is the cost of that portfolio. This use of probability pricing for hedging parallels the role of Greeks in option pricing, such as vega.

Second, expressing changes in probabilities as equivalent cash flows enables us to implement different decompositions, as we illustrate in our applications. A stochastic decomposition can be used to decompose probability prices into changes in expected payoffs and risk compensation. In equilibrium settings with heterogeneous agents, a cross-sectional welfare decomposition can be used to separate the value due to changes in aggregate consumption (aggregate/production efficiency) from the part that reflects the reallocation of consumption across agents in particular states (risk sharing).

Further insights arise from parametric perturbations. When the distribution is parameterized by its mean and standard deviation and consumption equals the state, the probability price of increasing the mean equals the valuation for the risk-free asset, since a uniform rightward shift of all probabilities is worth exactly as much as an asset that raises consumption by one unit in every state. The probability price of an increase in the standard deviation is equivalent to valuing an asset with positive payoffs in good states and negative payoffs in bad states.

The probability pricing formula relates to the classic literature on preferences over monetary lotteries that follows Pratt (1964), Rothschild and Stiglitz (1970), and Arrow (1971). In fact, we show that first-order stochastically dominant perturbations have a positive probability price, and that mean-preserving spreads, which worsen the distribution in the sense of second-order stochastic dominance, have a negative probability price for a risk-averse agent. Probability pricing generalizes these results in two directions: it assigns a positive or negative price to any perturbation, whether or not it satisfies a dominance relation, and it quantifies that price in consumption units rather than only signing it.

Application: Physical Probabilities/Climate Change. The probability pricing approach has broad implications and multiple uses, which we illustrate in two applications.[3] Our first application uses probability pricing to value changes in the distribution of global mean temperature induced by climate change. We model the world as a collection of countries heterogeneously exposed to the same global state, and price a marginal shift from the status quo temperature distribution — taken nonparametrically from the Coupled Model Intercomparison Project Phase 6 (CMIP6) climate-model ensemble — toward a warmer mid-century scenario, with damage functions estimated from a cross-country growth panel.

In this setting, the consumption-equivalent cash flows that determine probability prices define country-specific hedging portfolios, which we use to characterize the demand for temperature-contingent securities, the volume and gains from trade such securities would generate, and the decomposition of the world's valuation into aggregate-efficiency and cross-country risk-sharing components.

We find that climate change reallocates large gross exposures across countries that sit on opposite sides of the growth-maximizing temperature, but that most of the world's exposure is an aggregate position that cross-country trade cannot net out. The tradable part is small at short horizons and grows as damages compound: between the one- and ten-year horizons, the predicted volume grows more than tenfold and the gains from trade a hundredfold. Clearing the cross-section of hedging demands as a double auction, one temperature state at a time, also separates where a temperature security would be most traded from where it would be most valuable: the predicted volume peaks near 14.9°C, just above the status quo mean of 14.52°C, whereas the gains from trade it generates peak at warm but not extreme states near 16°C, where fewer units trade but each crosses a wider gap in valuations.

Application: Information/Precision of Public Signal. A change in the information available to an agent is a change in probabilities. This makes the valuation of information a natural application of the probability pricing framework: changes in information are valued through the probability changes they induce. Our second application studies an economy with ex-ante identical individuals in which insurance markets open only after a public signal is realized, following Hirshleifer (1971). In this environment, more precise

[3]Earlier versions of the paper contained applications to a principal-agent problem and to a model with private information.

public information makes everyone worse off because it resolves uncertainty before agents can insure against it. Probability pricing yields three results.

First, we show that the valuation of a change in public information is determined by two distinct channels. The first is a *reallocation channel*: changes in public information modify equilibrium consumption allocations for each signal realization. In particular, as the signal becomes less informative, consumption becomes more similar across individuals for a given signal realization, a welfare-improving effect. This is the familiar mechanism emphasized in the literature. The second is a *probability channel*, a previously unexplored channel that probability pricing allows us to study. For a given mapping between consumption allocations and signal realizations, a change in information changes the probabilities of observing different signal realizations, which has an ambiguous impact on welfare.

Second, we characterize when this newly identified probability channel is welfare-improving or welfare-reducing. The probability channel is welfare-improving when a more precise signal *compresses* the distribution of signal realizations, reducing the likelihood of extreme realizations in which consumption allocations are most unequal. It is welfare-reducing when a more precise signal instead *disperses* that distribution. In a pure endowment economy, however, the reallocation channel always dominates, so more precise public information is always welfare-reducing.

Finally, when production decisions respond to the public signal, probability pricing quantifies separately the risk-sharing and production-efficiency consequences of changes in the precision of public information. This yields a measure of the relative welfare importance of the two channels.

Related Literature. Probability pricing is most related to the classic work characterizing notions of risk and risk aversion that follows Pratt (1964), Rothschild and Stiglitz (1970), and Arrow (1971), by now standard in graduate textbooks; see, e.g., Ingersoll (1987), Mas-Colell, Whinston, and Green (1995), Gollier (2001), or Campbell (2017). To the best of our knowledge, the probability pricing formula and its consumption-equivalent characterization of probability changes are new contributions to this literature. Two features of our approach explain why. First, we compute the valuation of a general perturbation given a preference specification, rather than deriving stochastic orders that rank specific perturbations for arbitrary preferences. Second, we focus on marginal perturbations, which is what allows

us to present our results in the language of cash-flow pricing.[4]

The value of changes in probabilities has been computed in specific settings. For instance, Barro (2009) computes the welfare cost of disaster probabilities and consumption volatility in a consumption-based asset pricing model; see also Martin and Pindyck (2015). Our climate application shares the spirit of this work, since it values a change in the distribution of a state that changes consumption growth, but does so nonparametrically and for a cross-section of heterogeneously exposed countries, and our general results value arbitrary perturbations with few assumptions on preferences or distributions. The construction of the consumption equivalent cash flows is unique to our approach. Probability pricing, which expresses the effect of a probability change as an asset price, is also distinct from the study of how probability changes affect asset prices themselves (e.g., Martin, 2013); we elaborate on this distinction in Appendix D.6.

Finally, whether information improves efficiency is a central question in the literature on the efficiency of stock markets (e.g., Grossman and Stiglitz, 1980; Hellwig, 1980; Diamond and Verrecchia, 1981; Vives, 2017), strategic trade (e.g., Kyle, 1989; Vives, 2011; Rostek and Weretka, 2012), and public disclosures (e.g., Hirshleifer, 1971; Diamond, 1985; Diamond and Verrecchia, 1991; Morris and Shin, 2002; Angeletos and Pavan, 2007; Goldstein and Yang, 2019). The conventional tools of valuation and welfare analysis are of limited use in these problems, since information changes the probabilities that agents attach to different states of the economy, in addition to prices and allocations. Perhaps for this reason, and despite the contributions of Morris and Shin (2002); Angeletos and Pavan (2007, 2009); Veldkamp (2011); Gottardi and Rahi (2014); Vives (2017); Kadan and Manela (2019); Bond and García (2022); Pavan, Sundaresan, and Vives (2022), the welfare analysis of models with imperfect information, in particular with risk-averse agents, remains understudied. Our information application illustrates how probability pricing separates the endogenous consumption adjustments that take place in equilibrium from the direct change in probabilities and, because it is based on consumption equivalents, can serve as the foundation for quantitative work in this area, connecting to Ai and Bansal (2018), Kadan and Manela (2019), and Veldkamp (2023), among others.

[4]See, e.g., Alvarez and Jermann (2004) and Hansen and Souganidis (2025) for related applications of marginal valuation methods.

2 Probability Pricing: Theory

This section introduces probability pricing in a two-date single-agent setting.

2.1 Environment

We consider a single-agent environment with two dates, $t \in \{0, 1\}$. At date 1, there is a continuum of possible states indexed by s with (possibly unbounded) support $[\underline{s}, \bar{s}]$. We denote the cumulative distribution function (cdf) of the state by $F(s) \in [0, 1]$, and its probability density function (pdf) by $f(s) > 0$.

The agent has standard expected utility preferences, given by

$$V = u(c_0) + \beta \int_{\underline{s}}^{\bar{s}} u(c_1(s)) f(s) \, ds, \tag{1}$$

where c_0 denotes date-0 consumption, $c_1(s)$ denotes date-1 consumption in state s, and $\beta \in [0, 1]$ is the time discount factor. The flow utility function $u(\cdot)$ is twice differentiable, strictly increasing, and concave. To maintain focus, we restrict attention to expected utility preferences, although the analysis can be extended to more general preference specifications, as shown in the Online Appendix.

2.2 Cash Flow Pricing

To fix ideas, it is useful to first consider the standard problem of cash-flow pricing. Suppose the agent can purchase q units of an asset with state-contingent cash flows $x(s)$, at a price p_x. The budget constraints are

$$\begin{aligned} c_0 &= \ldots - p_x q \\ c_1(s) &= \ldots + x(s) q, \end{aligned}$$

where the ellipses $(\ldots)$ represent other elements in the agent's budget constraints. The agent's valuation for a marginal unit of the asset satisfies the familiar asset pricing formula:

$$p_x = \int_{\underline{s}}^{\bar{s}} \omega(s) x(s) \, ds, \tag{2}$$

where $\omega(s) = \frac{\beta u'(c_1(s))}{u'(c_0)} f(s)$ defines a state price. One can also express p_x in terms of a stochastic discount factor $m(s) = \frac{\beta u'(c_1(s))}{u'(c_0)}$, which satisfies $\omega(s) = m(s) f(s)$. In this case, cash-flow pricing uncovers the date-0 valuation for changes in consumption at date 1 in different states induced by the asset's cash flows. Asset prices are higher for assets with higher payoffs $x(s)$, in particular in states with high state prices $\omega(s)$.

2.3 Probability Pricing

We now show that the logic behind cash-flow pricing extends to characterize the agent's date-0 valuation for changes in probabilities. We refer to this alternative thought experiment as *probability pricing*.

To formalize changes in probabilities, we introduce a perturbation parameter θ that shifts the cdf (and pdf) over states, assuming that $F(s;\theta)$ and $f(s;\theta)$ are differentiable functions of θ. Our goal is to characterize the probability price, p_θ: the agent's valuation for a marginal change $d\theta$, reflected in the budget constraint as

$$c_0 = \ldots - p_\theta \theta.$$

To isolate the new effects that arise from perturbing probabilities, we initially assume that the agent's consumption profile is predetermined and does not depend on θ. We relax this assumption below.

The agent's valuation for a marginal change in probabilities satisfies the formula

$$p_\theta = \int_{\underline{s}}^{\bar{s}} \frac{\beta u(c_1(s))}{u'(c_0)} \frac{df(s;\theta)}{d\theta} ds, \tag{3}$$

where it must be that $\int_{\underline{s}}^{\bar{s}} \frac{df(s;\theta)}{d\theta} ds = 0$.[5] This expression has an intuitive interpretation: the valuation for a change in probabilities is high, all else equal, if increases in density $\frac{df(s)}{d\theta}$ coincide with high-consumption states, i.e., those with large $u(c(s))$. However, directly

[5] Formally, this perturbation can be interpreted as taking a variational/Gateaux derivative (Luenberger, 1969). To ensure that the perturbed cdf remains valid without making parametric assumptions, $F(s;\theta)$ can always be formulated as

$$F(s;\theta) = (1-\theta)\underline{F}(s) + \theta \overline{F}(s),$$

where $\underline{F}(s)$ denotes the cdf of the "initial" distribution, and $\overline{F}(s)$ denotes the cdf of the "final" distribution, or equivalently the "direction" of the perturbation. In general, the parameter θ can be mapped to a parameter of a particular distribution; see e.g., Section 2.4. See Dávila and Walther (2023) for an application of variational derivatives to leverage regulation with distorted beliefs.

comparing p_x and p_θ is not possible because Equation (3) is expressed in terms of utility levels $u(\cdot)$, while (2) is expressed in terms of marginal utilities $u'(\cdot)$.[6]

Our main result, Proposition 1, transforms changes in probabilities into consumption equivalents. The resulting probability pricing formula parallels the traditional cash-flow pricing formula. While the proof of Proposition 1 relies on integration by parts and thus makes use of the continuity of the underlying distribution of s, the same logic applies with discrete states after relying on summation by parts, as shown in Appendix D.2.

Proposition 1 (Probability Pricing). *The valuation, or probability price, p_θ, for a marginal perturbation in probabilities indexed by θ is given by*

$$p_\theta = \int_{\underline{s}}^{\bar{s}} \omega(s)\, x_\theta(s)\, ds, \tag{4}$$

where $\omega(s) = \frac{\beta u'(c_1(s))}{u'(c_0)} f(s;\theta)$ defines a state price, and where

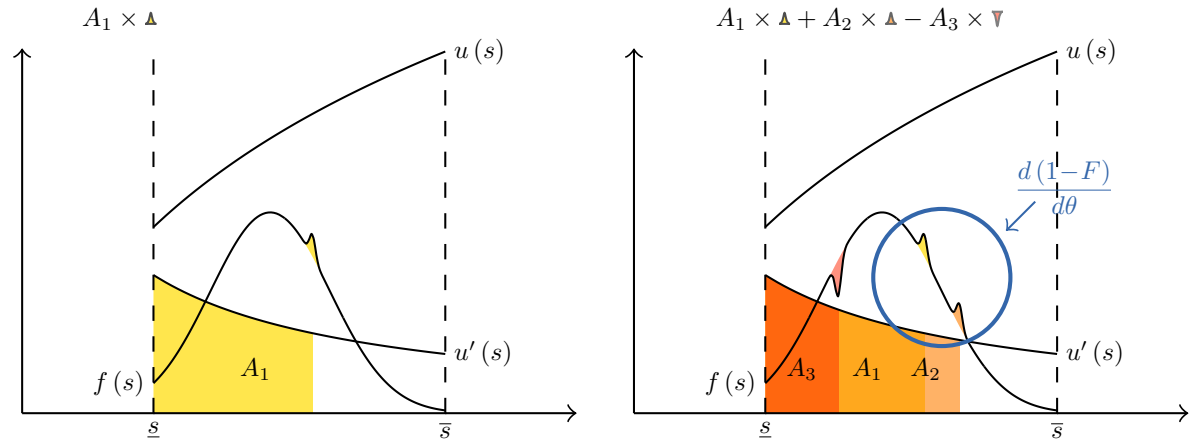
$$x_\theta(s) = \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f(s;\theta)} \frac{dc_1(s)}{ds} \tag{5}$$

defines a consumption-equivalent cash flow for state s.

Equation (4) shows that computing the agent's valuation for a marginal change in probabilities is equivalent to valuing a hypothetical asset that pays cash flows $x_\theta(s)$ in state s. These are the appropriate hypothetical cash flows that translate changes in probabilities into consumption-equivalent changes.

A basic economic intuition behind this transformation is as follows. Note that $\frac{d(1-F(s;\theta))}{d\theta}$ represents the change in the survival function, i.e., the total probability mass that the perturbation shifts from states to the left of s to states to the *right* of s. Dividing by the density $f(s;\theta)$ expresses this change relative to the likelihood of state s, yielding the *normalized survival change*, $\frac{\frac{d(1-F(s;\theta))}{d\theta}}{f(s;\theta)}$. Probability pricing says the cash flow equivalent of a perturbation in probabilities is large for an agent if probabilities are shifted towards higher realizations of s. Moreover, the equivalent cash flow is particularly large if those shifts occur in states in which the sensitivity $\frac{dc_1(s)}{ds}$ of consumption to the state is large.

[6] One may be tempted to view the ratio $\frac{\beta u(c_1(s))}{u'(c_0)}$ as analogous to the state price for cash-flow pricing. However, this ratio is not invariant to preference-preserving transformations — in particular, additive transformations of $u(\cdot)$ such as $u(\cdot) \to u(\cdot) + a$ — which makes it unsuitable as a foundation for a theory of valuation of changes in probabilities.



(a) Single State Perturbation

(b) General Perturbation

Figure 1: Illustrating Probability Pricing

Note: In this figure, $c_1(s) = s$, so $u(c_1(s)) \equiv u(s)$. The left panel illustrates the change in date-1 utility, $\int_{\underline{s}}^{\bar{s}} u(s) f(s) ds$, induced by perturbing the pdf in a particular state. In this case, the fundamental theorem of calculus allows us to express the utility change $u(s) \frac{df(s)}{d\theta}$ as $A_1 \times$, assuming without loss that $u(\underline{s}) = 0$. The right panel illustrates the change in date-1 utility induced by perturbing the pdf in multiple states. The probability price of the general perturbation, p_θ, corresponds to the sum over the perturbed states of the products of the areas A_1 to A_3 with the density changes:

$$p_\theta = A_1 \times + A_2 \times - A_3 \times ,$$

where we exploit the fact that $\int_{\underline{s}}^{\bar{s}} \frac{df(s)}{d\theta} ds = 0$. This sum can equivalently be calculated by adding up (integrating) at every state the product of marginal utility $u'(\cdot)$ with the change in densities to the right $\frac{d(1-F(\cdot))}{d\theta}$: this is exactly the probability pricing formula in Equation (4).

For a more detailed understanding of Equation (5), note that gaining the utility flow $u(\cdot)$ at a given state is equivalent to gaining the marginal utility $\frac{du(c_1(s))}{ds} = u'(c_1(s))\frac{dc_1(s)}{ds}$ in all states to the *left* of that state, by virtue of the fundamental theorem of calculus. Next, note that the normalized survival change $\frac{\frac{d(1-F(s;\theta))}{d\theta}}{f(s;\theta)}$ aggregates all the (net) changes in densities to the *right* of state s, where each of these changes induces a welfare gain valued at $u'(c_1(s))\frac{dc_1(s)}{ds}$. Aggregating these gains across all states yields the probability pricing formula in Equation (4). Figure 1 illustrates this logic.

In principle, both $\frac{d(1-F(s;\theta))}{d\theta}$ and $\frac{dc_1(s)}{ds}$ can take negative values, so the consumption-equivalent defined in Equation (5) can be negative, and so can the probability price p_θ. For instance, if the normalized survival change is positive, then the distribution of s becomes locally more "optimistic", since probability mass shifts from low to high realizations of s. A negative value of p_θ simply indicates that the perturbation makes the agent worse off, so the valuation for it is negative.[7]

Remark 1 (*Alternative Derivations*). Integration by parts of Equation (3) is a convenient way to obtain the result in Proposition 1, but it is not necessary.[8] Appendix D.1 presents two alternative derivations. The first changes variables from states to quantiles: at a fixed quantile, the perturbation moves the state, the chain rule produces the marginal utility directly, and the consumption-equivalent cash flow appears as the derivative of consumption with respect to θ holding the quantile fixed. The second exploits the fundamental theorem of calculus in the definition of expected utility, before differentiating. Both alternative derivations yield Equation (4) and extend to the case in which consumption responds to θ, as in Equation (7) below.

Remark 2 (*State Prices as the Pricing Basis for Probability Changes*). One might conjecture that pricing changes in probabilities requires a new basis distinct from state prices. Proposition 1 shows that it does not: state prices also serve as the basis for pricing changes

[7] Computing an agent's valuation for a marginal change in probabilities (probability pricing, p_θ) is different from computing how an agent's valuation for an asset changes in response to a marginal change in probabilities (comparative statics of cash-flow pricing, $\frac{dp_x}{d\theta}$). The answer to the latter question can be expressed as

$$\frac{dp_x}{d\theta} = \int_{\underline{s}}^{\bar{s}} \omega(s) \frac{\frac{df(s;\theta)}{d\theta}}{f(s;\theta)} x(s)\, ds.$$

Thus $\frac{dp_x}{d\theta}$ and p_θ differ. We further elaborate on this difference in Appendix D.6.

[8] Integration by parts, the key step behind our main result, is widely used in screening models and mechanism design (Mirrlees, 1971; Baron and Myerson, 1982; Pavan, Segal, and Toikka, 2014) — see Bolton and Dewatripont (2005) for a textbook treatment. Beyond this high-level relation, the focus of those papers and ours is different.

in probabilities.

Uses of Probability Pricing. The following paragraphs highlight the benefits of expressing probability prices in terms of marginal utilities and consumption equivalents, as in Equation (4), rather than in terms of utility levels, as in (3).

Hedging/Immunization against Changes in Probabilities. The consumption-equivalent cash flows defined in (5) offer a direct way to construct a trading strategy that would immunize an investor against (initially unanticipated) changes in probabilities. In particular, an investor who seeks insurance against changes in the probabilities of different scenarios can use Equation (5) to identify the cash flows that a hedging strategy must be designed to replicate. This strategy ensures that the investor's welfare is immunized against changes in probabilities. This application of the probability pricing formula for hedging/immunization is analogous to how Greeks — in particular, vega — are employed in option pricing. However, note that Equation (5) explicitly provides the replicating payoffs needed to immunize against any change in probabilities. The replicating portfolio depends only on changes in probabilities and consumption sensitivities $\frac{dc_1(s)}{ds}$, and is independent of the agent's preferences.

Stochastic and Cross-Sectional Decompositions. Equation (4) admits a useful stochastic decomposition into two components, an expected (equivalent) payoff component and a risk compensation component, as follows:

$$p_\theta = \underbrace{\frac{1}{1+r^f}\frac{d\mathbb{E}\left[c_1\left(s\right)\right]}{d\theta}}_{\text{Expected Payoff}} + \underbrace{\mathbb{C}ov\left[m\left(s\right), \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f\left(s;\theta\right)}\frac{dc_1\left(s\right)}{ds}\right]}_{\text{Risk Compensation}}, \tag{6}$$

where $1 + r^f = 1/\mathbb{E}\left[m\left(s\right)\right]$ denotes the risk-free rate and where, as shown in Appendix A, $\frac{d\mathbb{E}[c_1(s)]}{d\theta} = \mathbb{E}\left[\frac{\frac{d(1-F(s;\theta))}{d\theta}}{f(s;\theta)}\frac{dc_1(s)}{ds}\right]$. This decomposition clarifies whether a perturbation to probabilities is valuable because it changes expected consumption or because it reshuffles consumption across states with different valuations. We illustrate this use in our climate application (Section 3).

Cross-Sectional Welfare Decompositions. Probability pricing can also be useful in economies with heterogeneous agents and incomplete markets to compute cross-sectional welfare decompositions of the form introduced in Dávila and Schaab (2025). These decompositions require expressing the impact of perturbations in consumption-equivalent

units, which would not be possible without the probability pricing formula introduced in Equation (4). We illustrate this use in both of our applications.

Consumption Adjustments and Probability Pricing. In general, we can allow the agent's consumption profile at date 1 to vary with the perturbation parameter θ by writing $c_1(s;\theta)$; we could do the same with date-0 consumption. In particular, the mapping $c_1(s;\theta)$ can capture endogenous consumption adjustments that take place in equilibrium in response to changes in probabilities. In this case, the agent's valuation, expressed in date-0 consumption units, can be written as a combination of a cash-flow pricing term for consumption adjustments and a probability pricing term, as follows:

$$\frac{\frac{dV}{d\theta}}{u'(c_0)} = \int_{\underline{s}}^{\bar{s}} \omega(s) \left(\underbrace{\frac{\partial c_1(s;\theta)}{\partial \theta}}_{\text{Consumption}} + \underbrace{\frac{\frac{d(1-F(s;\theta))}{d\theta}}{f(s;\theta)} \frac{\partial c_1(s;\theta)}{\partial s}}_{\text{Probability}} \right) ds, \tag{7}$$

where $\omega(s)$ is the state price, as in (4). Equation (7) is agnostic about the exact mechanism through which the consumption mapping $c_1(s;\theta)$ depends on changes in probabilities. In our application to the value of public information (Section 4), consumption varies in response to a change in probabilities conditional on a state because agents' decisions endogenously adjust in competitive equilibrium. More generally, this consumption adjustment may also be due to changes in optimal contracts. Similar forces would also apply to economies with strategic trade (e.g., Kyle, 1989; Vives, 2011; Rostek and Weretka, 2012). Our climate application (Section 3) shows how probability pricing extends to dynamic economies.

Moreover, Equation (7) also implies that one can construct changes in consumption holding probabilities fixed and, symmetrically, changes in probabilities holding the consumption mapping fixed that yield equivalent valuation/welfare implications. Appendix D.7 elaborates on this equivalence, contrasting the approaches to evaluating the cost of business cycles in Lucas (1987) and Alvarez and Jermann (2004).

2.4 Probability Pricing for Specific Perturbations

We now illustrate the implications of the probability pricing result by examining specific perturbations to probability distributions. This is useful to better understand the form that normalized survival changes — and consequently consumption-equivalent cash flows

$x_\theta(s)$ and probability prices p_θ — take in practical scenarios.

2.4.1 Affine Perturbations

First, we highlight properties of probability prices by considering perturbations to a distribution that is characterized in terms of a mean parameter μ and a standard deviation parameter σ. Suppose that the state is defined in affine form, as in

$$s = \mu + \sigma n, \tag{8}$$

where μ and $\sigma \geq 0$ are parameters, and where n is a random variable distributed according to a cdf $H(n)$.[9] In this case, the normalized survival change $\frac{\frac{d(1-F(s))}{d\theta}}{f(s)}$ can be expressed in closed form:

(i) A marginal increase in the mean of s, $d\mu$, induces a normalized survival change given by

$$\frac{\frac{d(1-F(s))}{d\mu}}{f(s)} = 1. \tag{9}$$

(ii) A marginal increase in the standard deviation of s, $d\sigma$, induces a normalized survival change given by

$$\frac{\frac{d(1-F(s))}{d\sigma}}{f(s)} = \frac{s-\mu}{\sigma}. \tag{10}$$

Equation (9) shows that the hypothetical cash flow determining probability prices induced by a marginal increase in μ is simply the consumption sensitivity $\frac{dc_1(s)}{ds}$, so

$$p_\mu = \int_{\underline{s}}^{\bar{s}} \omega(s) \frac{dc_1(s)}{ds} ds \quad \underset{c_1(s)=s}{\Longrightarrow} \quad p_\mu = \int_{\underline{s}}^{\bar{s}} \omega(s)\, ds. \tag{11}$$

If $c_1(s) = s$, the distribution $F(s)$ is directly defined over consumption (s is a lottery), and Equation (11) implies that the probability price of a marginal increase in μ is the same as the price of the risk-free asset. Intuitively, the valuation for an asset that increases consumption unconditionally in all future states (risk-free asset) must be the same as the valuation for a perturbation that shifts all probabilities uniformly to the right when $c_1(s) = s$ (marginal increase in μ).

[9] If we further assume that $\mathbb{E}[n] = 0$ and $\mathbb{V}ar[n] = 1$, then $\mathbb{E}[s] = \mu$ and $\mathbb{V}ar[s] = \sigma^2$, but this is not necessary for Equations (9) and (10) to hold.

Similarly, Equation (10) implies that the hypothetical cash flow determining probability prices induced by a marginal increase in σ is given by $\left(\frac{s-\mu}{\sigma}\right)\frac{dc_1(s)}{ds}$, so

$$p_\sigma = \int_{\underline{s}}^{\bar{s}} \omega(s)\left(\frac{s-\mu}{\sigma}\right)\frac{dc_1(s)}{ds}ds \quad \underset{c_1(s)=s}{\Longrightarrow} \quad p_\sigma = \int_{\underline{s}}^{\bar{s}} \omega(s)\left(\frac{s-\mu}{\sigma}\right)ds. \tag{12}$$

A marginal increase in σ shifts mass to the tails, reducing survival probabilities in states below the mean ($s < \mu$), and increasing them in states above the mean ($s > \mu$). It thus follows from Equation (6) that, when $c_1(s) = s$, the probability price of a shift in σ is exclusively driven by risk compensation. Moreover, if $c_1(s) = s$, Equation (12) implies that the valuation for a forward contract with forward price μ must be the same as the valuation for a perturbation that increases the volatility of consumption σ, after normalizing by $1/\sigma$.

2.4.2 Mixture Distributions

A second useful scenario is one in which the state is defined as a mixture of two distributions. Formally, suppose that with probability $1-\theta$ the state follows a distribution with cdf $\underline{F}(s)$, and with probability θ the state follows a distribution with cdf $\overline{F}(s)$. Such mixtures are often used to model discrete jumps, large shocks, or disasters; in our climate application below, the mixture interpolates between a status quo and a counterfactual distribution.

In this case, a marginal increase in the mixing probability θ leads to a normalized survival change given by

$$\frac{\frac{d(1-F(s))}{d\theta}}{f(s)} = \frac{\underline{F}(s) - \overline{F}(s)}{(1-\theta)\,\underline{f}(s) + \theta\overline{f}(s)}. \tag{13}$$

The sign of the normalized survival change — and hence the hypothetical consumption-equivalent cash flow in state s — is entirely determined by the difference between cdf's $\underline{F}(s) - \overline{F}(s)$. Moreover, note that under this mixture structure p_θ is invariant to the level of θ. This is a property that we exploit in our climate application (Section 3).

2.4.3 Stochastic Dominance

Finally, we show how the probability pricing formula relates to the classic literature on preferences over lotteries.[10] Suppose that consumption is $c_1(s) = s$, so that the state s

[10] See, for example, Pratt (1964), Rothschild and Stiglitz (1970) and Arrow (1971). More recent work includes, among others, Aumann and Serrano (2008), Foster and Hart (2009), and, outside of the expected

and the distribution $F(s)$ define a lottery over consumption. In this case, probability prices satisfy the following properties:

(i) *First-order stochastic dominance:* The probability price for perturbations such that $\frac{dF(s)}{d\theta} \leq 0$ — equivalently, $\frac{d(1-F(s))}{d\theta} \geq 0$ — for all s satisfies $p_\theta \geq 0$.

(ii) *Second-order stochastic dominance/mean-preserving spreads:* The probability price for perturbations such that $\frac{d\mathbb{E}[s]}{d\theta} = 0$ and $\int_{\underline{s}}^{s} \frac{dF(t)}{d\theta} dt \geq 0$ for all s satisfies $p_\theta \leq 0$.

The two properties just highlighted confirm that probability prices can be used to derive well-known properties of the preferences of risk-averse agents. The first property shows that an agent is always willing to pay a positive price for a perturbation that implies "good news" in the sense of first-order stochastic dominance. The second property, as in Rothschild and Stiglitz (1970), shows that a risk-averse agent is never willing to pay a positive price for a perturbation that implies "more risk" in the sense of a mean-preserving spread. Note that these properties rely on assuming $c_1(s) = s$. Outside of the special case in which $\frac{dc_1(s)}{ds} = 1$, the sensitivity of consumption to the state $\frac{dc_1(s)}{ds}$ can have non-trivial implications that go beyond the classic literature on lotteries.

More broadly, Proposition 1 can be viewed as a generalization of this classical literature. Instead of insisting on deriving stochastic orders over the distributions of particular gambles, probability pricing provides the exact valuation — positive or negative — of an agent for arbitrary marginal changes in the distribution of outcomes, whether or not they satisfy a dominance relation.

3 Physical Probabilities: Probability Pricing of Climate Change

The remainder of the paper develops two applications of probability pricing. In the first, we quantify the gains and losses that a shift in the entire distribution of global temperature imposes on each country and on the world. Climate change is a natural setting for the framework of Section 2, since the object of interest is a change in physical probabilities — the probability of each level of warming — rather than a change in the cash flows paid in each state.

utility framework, Ai, Bansal, Guo, and Yaron (2024).

We model the world as a collection of countries exposed to a single global state, mean surface temperature, each through its own damage function, and we consider a nonparametric shift from the status quo temperature distribution toward an alternative distribution. We estimate country-specific growth responses to temperature from a cross-country panel and use them to quantify the country-specific and global impact of the shift. Probability pricing then allows us to identify hedging opportunities and predict which temperature-contingent markets would generate the most trade and the largest gains from trade.

Three lessons emerge. First, the shift in the temperature distribution reallocates large gross exposures across countries that sit on opposite sides of the growth-maximizing temperature, since additional warming raises growth in the coldest countries and lowers it in the hottest ones. Only a small part of that exposure is tradable, because most of it is an aggregate position that the world as a whole holds and cannot diversify away.

Second, the tradable part is small at short horizons and grows only as damages compound. At short horizons, country-specific valuations of the hedge are compressed and ordered by climate exposure, so the countries most exposed to warming are the ones that value the hedge the least.

Third, the object that identifies where a hedging market would be most valuable is the gains-from-trade surplus it would generate rather than the trading volume it would attract, and the two peak at different temperatures.

3.1 Environment

Time is discrete, with dates indexed by $t \in \{0, 1, 2, \ldots\}$. The world consists of countries indexed by $i \in \{1, 2, \ldots, I\}$, each populated by a representative agent with well-behaved preferences

$$V^i = u\left(c_0^i\right) + \mathbb{E}\left[\sum_{t \geq 1} \beta^t u\left(c_t^i\left(s\right)\right)\right],$$

where β is the discount factor, $u\left(\cdot\right)$ denotes flow utility, and $c_t^i\left(s\right)$ is country i's per-capita consumption at date t when the global state s — global mean surface temperature — realizes.

The status quo distribution of s, with cdf $\underline{F}\left(s\right)$, describes the current climate. We compare it with an alternative distribution $\overline{F}\left(s\right)$ associated with a counterfactual warming

scenario. The probability-pricing perturbation interpolates between the two,

$$F(s;\theta) = \theta \overline{F}(s) + (1-\theta)\,\underline{F}(s), \qquad \theta \in [0,1], \tag{14}$$

so that $\theta = 0$ recovers the status quo and $\theta = 1$ the alternative; a marginal experiment shifts the distribution by $d\theta$ toward the alternative. We write $f(s;\theta)$ for the corresponding density.

To incorporate dynamics while staying close to the static setting of Section 2, we consider permanent changes in the distribution of temperature; transitory changes are straightforward to accommodate. That is, the temperature state s realizes at date 1 and determines consumption growth at every subsequent date. Formally, country i's consumption depends on the state through a country-specific damage function that acts on its growth rate. Given an initial level c_0^i,

$$c_t^i(s)/c_0^i = e^{g_t^i(s)t}, \quad \text{where} \quad g_t^i(s) = \bar{g}_t^i + h^i\left(T^i(s)\right) - h^i\left(\bar{T}^i\right) \quad \text{and} \quad h^i(T) = \phi_1^i T + \phi_2^i T^2. \tag{15}$$

Here, $T^i(s)$ denotes country i's local temperature, a function of the global state s described below, $\bar{T}^i$ is its historical mean temperature, and $\bar{g}_t^i$ is the average baseline growth rate between dates 0 and t, the rate at which the country would grow absent any temperature effect, that is, if its local temperature stayed at $\bar{T}^i$. The temperature component of growth, $h^i(T^i(s)) - h^i\left(\bar{T}^i\right)$, is measured relative to the country's own climate and vanishes at its mean temperature. This is consistent with the way we identify (ϕ_1^i, ϕ_2^i) below: the regression identifies the slope and curvature of $h^i(\cdot)$ from within-country temperature variation, while its level is absorbed by country fixed effects.[11]

3.2 Probability Pricing

Proposition 1 prices a marginal change in the distribution of global temperature. In this context, the probability price p_θ^i is the amount of date-0 consumption that country i would pay — or would need to receive, when negative — for a marginal shift from the status

[11] This specification rules out adaptation: a country's response to temperature does not change over time.

quo toward the alternative. It aggregates a sequence of horizon-specific valuations,

$$\frac{dV^i/d\theta}{u'\left(c_0^i\right)} = p_\theta^i = \sum_{t\geq 1} \beta^t p_{t,\theta}^i, \qquad \text{where} \qquad p_{t,\theta}^i = \int_{\underline{s}}^{\overline{s}} \omega_t^i\left(s\right) x_{t,\theta}^i\left(s\right) ds,$$

with the consumption-equivalent cash flow $x_{t,\theta}^i\left(s\right)$ and the state-price density $\omega_t^i\left(s\right)$ given by

$$x_{t,\theta}^i\left(s\right) = \frac{\underline{F}\left(s\right) - \overline{F}\left(s\right)}{f\left(s;\theta\right)} \frac{dc_t^i\left(s\right)}{ds}, \qquad \omega_t^i\left(s\right) = \frac{u'\left(c_t^i\left(s\right)\right)}{u'\left(c_0^i\right)} f\left(s;\theta\right). \tag{16}$$

Pricing the marginal change in probabilities is therefore equivalent to pricing a hypothetical asset that pays $x_{t,\theta}^i\left(s\right)$ in state s at each horizon t.[12] As Section 2 explains, the equivalent cash flows $x_{t,\theta}^i\left(s\right)$ have the interpretation of a *hedging strategy*: they are the payoffs an investor must replicate to immunize welfare against the change in probabilities. Preferences do not enter them; they depend only on the cdf gap $\underline{F}\left(s\right) - \overline{F}\left(s\right)$ and on the sensitivity of consumption to temperature. Under the mixture (14), the cdf gap is the same for every θ, so $x_{t,\theta}^i\left(s\right)$ depends on θ only through the density $f\left(s;\theta\right)$ at which the position is evaluated.

3.3 Measurement and Calibration

Besides preferences, probability prices depend on two inputs that we measure directly:[13]

(i) the status quo and alternative distributions of global temperature, and

(ii) the country-specific mapping from temperature to consumption, determined by the response functions $h^i\left(\cdot\right)$ and the baseline growth paths $\bar{g}^i\left(\cdot\right)$.

3.3.1 Nonparametric Temperature Distributions

First, we identify the distributions of the global state with the cross-model dispersion of the Coupled Model Intercomparison Project Phase 6 (CMIP6) multi-model archive (Eyring et al., 2016). The data come from the Intergovernmental Panel on Climate Change

[12] Differentiating (15), $\frac{dc_t^i(s)}{ds} = c_t^i\left(s\right) \frac{dg_t^i(s)}{ds} t = c_t^i\left(s\right) t \frac{dh^i\left(T^i(s)\right)}{dT} \frac{dT^i(s)}{ds}$.

[13] Both inputs are produced by a two-stage replication package. The first stage downloads the raw data and estimates every input — the per-model CMIP6 temperatures and their kernel distributions, the country-year growth–climate panel, the country damage coefficients, and the 2023 country levels and 1990–2019 growth rates — and saves each as a comma-separated file. The second stage reads those files, computes consumption paths, hedging positions, state prices, and the double-auction clearing of Section 3.6, and produces every figure and number reported in this section.

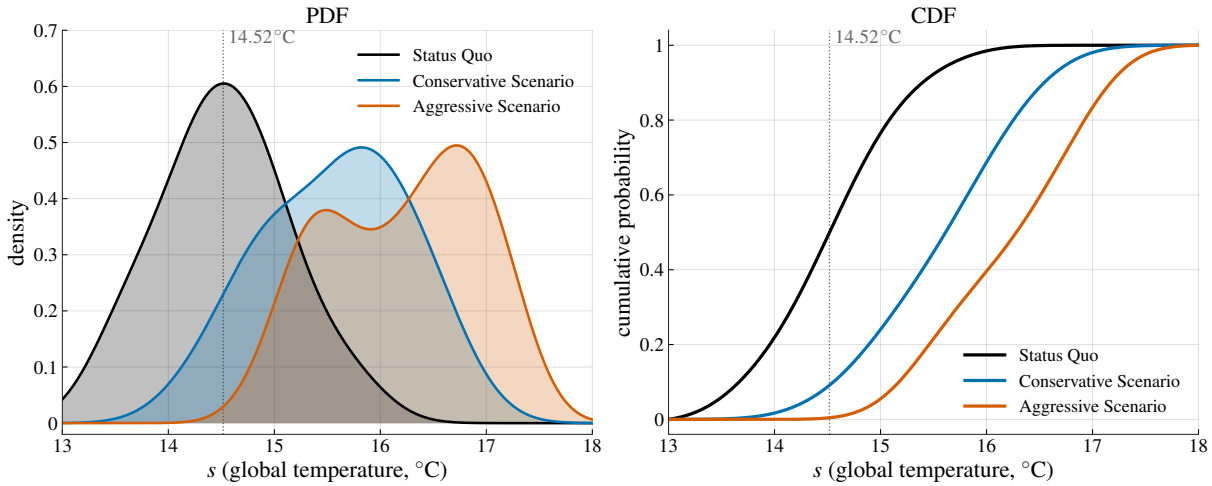


Figure 2: Nonparametric Distributions of Global Temperature

Note: The left panel plots the kernel-smoothed probability density and the right panel the cumulative distribution function of global mean surface temperature. The status quo is estimated from the 1995–2014 averages of 35 CMIP6 models; the conservative and aggressive scenarios are estimated from the 2041–2060 averages of 32 and 34 models under SSP1-2.6 and SSP5-8.5, respectively. Each density uses a Gaussian kernel with the Sheather–Jones bandwidth, truncated to $[13°\text{C}, 18°\text{C}]$ and renormalized. The dotted vertical line marks the cross-model status quo mean of 14.52°C.

(IPCC) Sixth Assessment Report Working Group I Atlas repository (Gutiérrez et al., 2021; Iturbide et al., 2022), which provides, for each CMIP6 model run, the annual mean near-surface air temperature averaged over the globe, land and sea included. For each model, we average annual global mean temperature over 1995–2014 to obtain a status quo realization (35 models), and over 2041–2060 to obtain a realization under each of two Shared Socioeconomic Pathways (SSP): SSP1-2.6, a conservative emissions scenario (32 models), and SSP5-8.5, an aggressive one (34 models). Each distribution therefore describes the disagreement among climate models about a given world: the status quo dispersion is disagreement about the recent past, and each alternative is disagreement about mid-century under a fixed emissions path.

To make the status quo and the alternatives absolutely continuous on a common support, we smooth each cross-model sample with a Gaussian kernel, using the plug-in bandwidth of Sheather and Jones (1991), truncate the result to the interval $[13°\text{C}, 18°\text{C}]$ — wide enough to enclose every observation with a margin — and renormalize it to integrate to one. The resulting density and cdf are stored on a grid of 501 points spaced 0.01°C apart, and every integral over s below is a trapezoidal sum on that grid. Figure 2 plots the densities and cdfs of the status quo and the two alternatives.

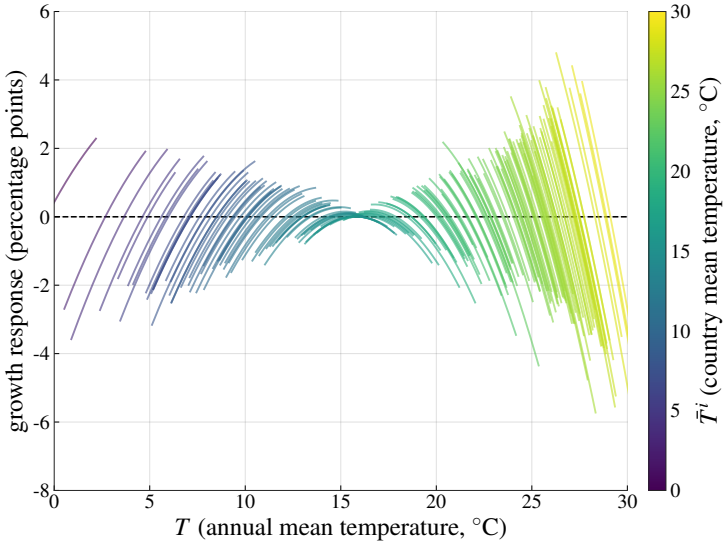


Figure 3: Country-Specific Growth Responses

Note: Each curve plots one country's growth response to temperature, $h^i(T) - h^i(\bar{T}^i)$ with $h^i(T) = \phi_1^i T + \phi_2^i T^2$, in percentage points per year, over the range of annual mean temperatures the country experienced in 1961–2010 (plus a margin of 0.5°C on each side), colored by its mean temperature $\bar{T}^i$. Cold-country curves slope upward at the country's own climate, so additional warming raises growth; hot-country curves slope downward, so additional warming reduces it.

The two alternatives have means of 15.6°C and 16.2°C, which differ by about 0.6°C, less than their within-scenario standard deviations of 0.65 to 0.70°C; the status quo has mean 14.52°C and standard deviation 0.59°C. Both alternatives place more probability on warmer outcomes than the status quo at every threshold and therefore first-order stochastically dominate it: $\underline{F}(s) - \overline{F}(s) \geq 0$ for every s. Probability pricing uses the entire cdf gap, so it distinguishes the two scenarios by how much probability mass each shifts toward warmer outcomes at every temperature, without reducing either to its mean.

3.3.2 Growth–Temperature Mapping

Next, we estimate the temperature–growth mapping $h^i(\cdot)$ from a cross-country panel regression of the growth of real gross domestic product (GDP) per capita on a quadratic in annual mean temperature, following the specification of Burke, Hsiang, and Miguel (2015), hereafter BHM.

To allow the impact of warming to differ across countries while preserving statistical power, we let the damage coefficients vary linearly with a country's typical climate. Formally, letting $\bar{T}^i$ be country i's mean temperature over the estimation sample and

$Z^i = \bar{T}^i - \frac{1}{I}\sum_j \bar{T}^j$ its cross-country demeaned value, we estimate

$$g_t^i = \left(\phi_{10} + \phi_{11} Z^i\right) T_t^i + \left(\phi_{20} + \phi_{21} Z^i\right) \left(T_t^i\right)^2 + \varsigma_1 P_t^i + \varsigma_2 \left(P_t^i\right)^2 + \alpha^i + \zeta_t + f^i(t) + \varepsilon_t^i, \quad (17)$$

where g_t^i is the growth of real GDP per capita of country i in year t, T_t^i and P_t^i are annual mean temperature and annual total precipitation, α^i are country fixed effects, ζ_t are year fixed effects, and $f^i(t)$ are country-specific quadratic time trends. The implied country-level damage coefficients are $\phi_1^i = \phi_{10} + \phi_{11} Z^i$ and $\phi_2^i = \phi_{20} + \phi_{21} Z^i$, and the pooled specification of BHM corresponds to $\phi_{11} = \phi_{21} = 0$. Since the regression includes country fixed effects and country-specific trends, the temperature coefficients are identified from within-country year-to-year deviations of annual temperature from its trend, that is, from weather rather than from climate.

We rebuild the BHM panel from its original sources and extend it. Annual growth of real GDP per capita comes from the World Bank World Development Indicators (WDI). Temperature and precipitation come from version 4.08 of the Climatic Research Unit gridded Time Series (CRU TS) dataset (Harris, Osborn, Jones, and Lister, 2020), which reports monthly values on a 0.5° latitude–longitude grid; for each grid cell, we average the monthly temperatures into an annual mean and sum the monthly precipitation into an annual total. To aggregate grid cells into countries, we assign each cell to a country using the Natural Earth administrative boundaries and weight cells by population from the WorldPop year-2000 raster (Tatem, 2017), coarsened to the same 0.5° grid, so that a country's temperature is the temperature its population experiences rather than the average over its land area. The weights are fixed at their year-2000 values, so changes in a country's temperature reflect weather rather than the movement of people.

Ordinary least squares (OLS) on the 1961–2010 panel, which covers the 168 countries with at least twenty usable country-years, with standard errors clustered by country, yields a negative curvature intercept, $\phi_{20} < 0$, so the average country's response is concave, and a negative interaction, $\phi_{21} < 0$, so hotter countries are more concave: the marginal effect of an additional degree on growth declines faster the warmer a country already is. Both signs hold country by country. Every implied response is concave ($\phi_2^i < 0$) and upward-sloping at zero degrees ($\phi_1^i > 0$), so each country has an interior growth-maximizing temperature, $-\phi_1^i / (2\phi_2^i)$.

Figure 3 plots the estimated responses of all 168 countries, each relative to its own mean temperature $\bar{T}^i$. Cold-country curves slope upward, so additional warming raises

growth; hot-country curves bend downward more strongly, so additional warming reduces it. This dependence on a country's typical climate drives the hedging analysis below: cool countries, where $dh^i/dT > 0$ locally, and hot countries, where $dh^i/dT < 0$ locally, are affected in opposite directions by the same global temperature shock and therefore gain from trading a temperature-contingent security.

3.3.3 Additional Calibration Inputs

Preferences. We assume common isoelastic preferences, $u(c) = c^{1-\gamma}/(1-\gamma)$, with a relative risk-aversion coefficient $\gamma = 2$, and an annual discount factor $\beta = 0.96$.

Global-to-local temperature mapping. To map the global temperature s into country-level temperatures, we assume that each country's temperature moves one-for-one with the global realization:

$$T^i(s) = \bar{T}^i + s - \bar{s},$$

where $\bar{T}^i$ is country i's historical mean temperature and $\bar{s} = 14.52°\text{C}$ is the cross-model mean of the status quo distribution. That is, a global realization one degree above the status quo mean raises every country's temperature one degree above its own historical mean, so $dT^i(s)/ds = 1$. This mapping abstracts from regional differences in the rate of warming.

Population shares and baseline consumption. When aggregating across countries, we weight by population shares $\chi^i = \text{pop}^i / \sum_j \text{pop}^j$, and 2023 GDP per capita at purchasing power parity (PPP) sets baseline consumption levels. Both come from the WDI: total population and GDP at PPP in current international dollars for 2023. The analysis sample consists of the 162 countries with complete inputs — 2023 population, 2023 PPP GDP, and 1990–2019 growth — and population shares are renormalized to sum to one over this set. We normalize per-capita consumption so that the population-weighted world average equals one at $t = 0$: with $\text{gdppc}^i = \text{gdp}^i_{\text{ppp}}/\text{pop}^i$ and $\overline{\text{gdppc}} = \sum_i \chi^i \text{gdppc}^i \approx 23,474$ PPP dollars,

$$c_0^i = \frac{\text{gdppc}^i}{\overline{\text{gdppc}}}, \qquad \text{so that} \qquad \sum_i \chi^i c_0^i = 1.$$

That is, c_0^i measures country i's 2023 PPP GDP per capita as a multiple of the world per-capita average, which places every quantity that follows on a scale directly comparable

across countries.

Baseline growth. The last input is the growth path each country would follow absent any temperature effect. To account for growth trends while allowing for convergence, we start each country at its mean 1990–2019 per-capita growth rate in the WDI, $\bar{g}^i$, and let growth converge linearly to a common long-run rate:

$$\bar{g}^i(t) = \begin{cases} \bar{g}^i + \dfrac{t}{T_c}\left(\bar{g}^{LR} - \bar{g}^i\right), & 0 \le t \le T_c, \\ \bar{g}^{LR}, & t > T_c, \end{cases}$$

with $\bar{g}^{LR} = 2.5\%$ and $T_c = 10$ years. The baseline rate that enters (15) is the average of this path up to date t, $\bar{g}^i_t = \frac{1}{t}\int_0^t \bar{g}^i(\tau)\, d\tau$.

3.4 Hedging Portfolio

Section 2 identifies the consumption-equivalent cash flow $x_\theta(s)$ as a hedging portfolio: the state-contingent payoff that immunizes an agent's welfare against the change in probabilities. With heterogeneous exposure the portfolio is country-specific — country i would replicate $x^i_\theta(s;\theta) = \left(\underline{F}(s) - \overline{F}(s)\right)\frac{dc^i_t(s)}{ds}/f(s;\theta)$, where the second argument indicates the density at which the position is evaluated — and because countries sit on opposite sides of the growth-maximizing temperature, their hedging demands point in opposite directions. One country's desired hedge is another's natural supply, so the economic question is how much of the world's climate exposure these offsetting demands can net out. This subsection traces, state by state and horizon by horizon, the cross-country hedging demand that the perturbation (14) induces across the 162 countries of the analysis sample under the aggressive scenario, in which $\underline{F}$ is the status quo kernel cdf and $\overline{F}$ the SSP5-8.5 kernel cdf, at horizons $t \in \{1, 5, 10, 20, 50, 100\}$ years.

It is convenient to work with the density-weighted cash flow

$$\tilde{x}^i_\theta(s) = x^i_\theta(s;\theta)\, f(s;\theta) = \left(\underline{F}(s) - \overline{F}(s)\right)\frac{dc^i_t(s)}{ds}, \tag{18}$$

the product of the two inputs measured in Section 3.3: the cdf gap $\underline{F} - \overline{F} \ge 0$ and the consumption derivative dc^i_t/ds implied by (15). The density cancels in (18), so the integrand — unlike the position x^i_θ itself — does not depend on the density at which it is evaluated. The population-weighted sum $\sum_i \chi^i \tilde{x}^i_\theta$ integrates over s to the marginal effect

of the perturbation on expected world consumption per capita at horizon t; weighting each country's position by its own state price before summing yields the value counterpart studied in Section 3.4.2. Since $\underline{F} - \overline{F} \geq 0$, the sign of each country's hedging demand is the sign of its consumption response, $g^{i\prime}(s) = \phi_1^i + 2\phi_2^i T^i(s)$: cold countries ($g^{i\prime} > 0$) are on one side of the market and hot countries ($g^{i\prime} < 0$) on the other.

3.4.1 The Cross-Section of Hedging Demand

Country by country, the hedging integrand (18) has a simple structure. Cold countries such as Germany and Russia have positive integrands at every state and horizon; hot countries such as Brazil, Indonesia, and India have negative ones; and the vertical scale grows with t, because both the consumption level c_t^i and the explicit horizon factor in dc_t^i/ds inflate the consumption derivative. Scaling each country by its population share χ^i — the contribution whose sum across countries is the world marginal effect — changes the picture in two ways. First, populous middle- and low-income countries (India, Indonesia, Nigeria, China, Brazil), which are mostly hot, gain weight relative to rich but small Northern countries, so the negative side dominates. Second, the compounding of growth at long horizons scales up high-growth countries disproportionately, so by $t = 100$ India alone reaches $|\chi^i \tilde{x}_\theta^i| \approx 1.2$.

Because world consumption per capita is the population-weighted sum of country consumption, the marginal effect of the perturbation on expected world consumption at any state is the algebraic sum $\sum_i \chi^i \tilde{x}_\theta^i(s)$, which can be read off a stacked bar at each s. Figure 4 stacks the contributions of all 162 countries, binned in 0.1°C intervals of s: gainers ($g^{i\prime} > 0$) above zero, losers ($g^{i\prime} < 0$) below, colored by mean temperature — cold countries in blue, hot countries in red — with the net $\sum_i \chi^i \tilde{x}_\theta^i$ drawn as a black line.

The stacked figure reveals three features that country-level plots hide. First, the spread between winners and losers is two to six times larger than the net, depending on the horizon: there are large gross gains from cross-country risk sharing. Second, at short and medium horizons ($t \leq 50$) the net is negative almost everywhere; once weighted by population, hot countries dominate. Third, computing the same stack at $t = 100$, the compounding of cold countries' warming gains turns the net positive on the warm side of the support — above roughly 14.8°C, with a peak near +1.8 around 15.7°C — while it remains strongly negative on the cold side, with a trough of about −2.55 near 13.9°C, a peak-to-trough difference of more than four. The gap between the gross stack heights and

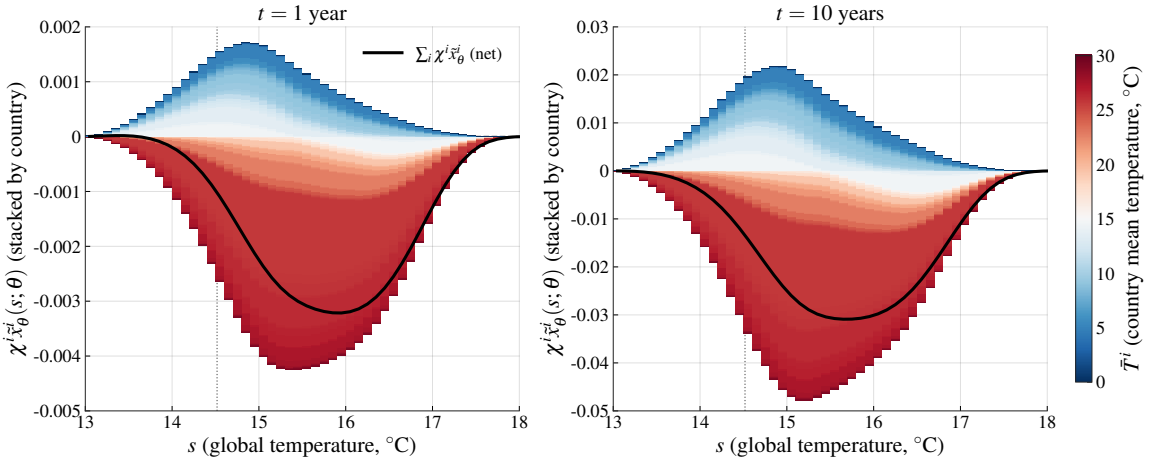


Figure 4: Stacked Hedging Integrand Across Countries

Note: Population-weighted hedging integrand $\chi^i \tilde{x}^i_\theta(s)$ stacked across the 162 countries of the analysis sample at $t = 1$ year (left panel) and $t = 10$ years (right panel), under the aggressive scenario (SSP5-8.5). Each colored segment is one country's contribution in one 0.1°C bin of s; segments stack above zero for gainers ($g^{i\prime} > 0$) and below zero for losers ($g^{i\prime} < 0$). Color encodes mean temperature $\bar{T}^i$; the black line is the net $\sum_i \chi^i \tilde{x}^i_\theta$. The dotted vertical line marks the status quo mean of 14.52°C.

the net line is the key object in the figure. It bounds cross-country reallocation from above: it is the θ-exposure that gainers would shed and losers absorb if a temperature-contingent security were traded, and any wedge between gross and net measures gains from trade that autarky leaves unrealized.

The stacked integrand treats every temperature symmetrically, but most of the support carries little probability under either scenario. Reweighting each contribution by the midpoint density, $f\left(s; \frac{1}{2}\right) \chi^i \tilde{x}^i_\theta(s)$, concentrates the picture on the states that occur: the central region $s \in [14, 17]^\circ$ C dominates, the tails are strongly damped, and the sign pattern is preserved, since $f > 0$ everywhere. Once we integrate over states, the density weighting flips the sign of the $t = 100$ net, from -0.0089 unweighted to $+0.14$: the negative cold-side mass, which offsets the positive warm-side mass in the unweighted net, receives less weight under the midpoint density. The gross gains from trade should therefore be read as the probability-weighted reallocation that would take place: most of the hedging value arises where uncertainty is largest in absolute terms rather than in the extreme tails.

3.4.2 The Value of Hedging

The positions so far are quantities; weighting each country's position by its state price expresses them in value. For country i at horizon t, the normalized Arrow–Debreu

state-price density is the bond-normalized counterpart of $\omega_t^i(s)$ in (16),

$$\omega_t^i(s) = \beta^t \frac{u'(c_t^i(s))\, f(s;\theta)}{\int u'(c_t^i(\sigma))\, f(\sigma;\theta)\, d\sigma}, \qquad u'(c) = c^{-\gamma},$$

with $\gamma = 2$ and $\beta = 0.96$, priced under the status quo density $f(s;0) = \underline{f}(s)$. Relative to the state price in (16), which is normalized by $u'(c_0^i)$, we normalize by each country's own expected marginal utility. Under isoelastic utility, both the consumption level c_0^i and the baseline growth term $e^{\bar{g}_t^i t}$ cancel from ω_t^i, so cross-country differences in state prices isolate exposure to the temperature lottery rather than differences in income levels or trend growth; the normalization also prices every country's t-period discount bond at β^t, which puts reservation valuations on a common scale. The value integrand $\omega_t^i \chi^i x_\theta^i$ integrates over s to a bond-normalized counterpart, at horizon t, of country i's contribution to the world probability price: it equals the corresponding term of $\chi^i p_\theta^i$ scaled by the country-specific factor $u'(c_0^i) / \int u'(c_t^i(\sigma))\, f(\sigma;\theta)\, d\sigma$ that converts between the two normalizations, so all cross-country comparisons below are made in common reservation-value units. Since the state price carries f while the position carries $1/f$, the integrand is free of tail blow-up: in terms of the hedging integrand, $\omega_t^i \chi^i x_\theta^i = m_t^i \chi^i \tilde{x}_\theta^i$, where $m_t^i(s) = \omega_t^i(s) / \underline{f}(s)$ is country i's per-probability reservation valuation of consumption in state s. Summing the positive and negative sides separately gives, at each state,

$$P(s) = \sum_i \left(\omega_t^i \chi^i x_\theta^i\right)^+, \qquad N(s) = \sum_i \left(-\omega_t^i \chi^i x_\theta^i\right)^+, \qquad \mathcal{V}^{\$}(s) = \min\{P(s), N(s)\},$$

where cold countries populate the positive side P and hot countries the negative side N, and the cleared value $\mathcal{V}^{\$}$ is bounded by the short side of the market.

3.5 Aggregate Efficiency and Risk Sharing

At each state, the value integrand admits the cross-sectional decomposition of Remark 3: the population-average position priced at the population-average state price, an aggregate-efficiency (AE) component, plus the cross-country covariance between state prices and positions, a risk-sharing (RS) component. The cross-country covariance identity with population weights χ^i, which sum to one over the sample, splits the value integrand at

each state as

$$\sum_i \chi^i \omega_t^i x_\theta^i = \underbrace{\left(\sum_i \chi^i \omega_t^i\right)\left(\sum_i \chi^i x_\theta^i\right)}_{\mathrm{AE}(s)} + \underbrace{\mathbb{C}ov_\chi\left[\omega_t^i, x_\theta^i\right]}_{\mathrm{RS}(s)}.$$

AE prices the population-average position at the population-average state price; RS captures the cross-country comovement between state prices and positions. RS would vanish identically if all countries shared one stochastic discount factor. It is nonzero precisely because the ω_t^i differ across countries — the signature of incomplete risk sharing — and it is the cross-sectional welfare object that probability pricing makes computable in consumption-equivalent units.

Figure 5 reports both components and their sum, which reproduces the net value integrand to machine precision. At short and medium horizons AE carries most of the net and RS is a small correction: at $t = 1$, $\int \mathrm{AE}\, ds = -6.82 \times 10^{-3}$ against $\int \mathrm{RS}\, ds = -2.18 \times 10^{-4}$. At $t = 10$, the two are -6.2×10^{-2} and -1.7×10^{-2}. The risk-sharing term grows in relative magnitude with the horizon as the cross-country dispersion in marginal utility widens; when the same integrals are computed at horizons beyond those in the figure, by $t = 50$ the two are of the same order (-0.179 and -0.138). At $t = 100$ the decomposition becomes a large near-cancellation ($\int \mathrm{AE}\, ds = 13.1$, $\int \mathrm{RS}\, ds = -13.2$, net -0.175) dominated by the warm tail, where the pricing density is negligible; there the net remains the well-behaved object, and the AE/RS split should be read with care.

3.6 Valuable Hedging Markets

The AE/RS identity has an equivalent reading as a market, which makes the hedging interpretation concrete. Fix a state s^* and a horizon t and treat the cross-section $\{(x_\theta^i(s^*), \omega_t^i(s^*))\}_i$ as a limit-order book for the Arrow–Debreu claim that pays in state s^*. Each country is one order: its *side* is the sign of x_θ^i — a buy order, $x_\theta^i(s^*) < 0$, if the country is under-provisioned in s^*, and a sell order otherwise — its *quantity* is $\chi^i |\tilde{x}_\theta^i(s^*)|$, and its *limit price* is the per-probability reservation valuation $m_t^i(s^*) = \omega_t^i(s^*)/\underline{f}(s^*)$. Stacking buy orders from the highest reservation valuation down gives the demand schedule, and stacking sell orders from the lowest up gives the supply schedule. Clearing the book is a double auction: we match the highest-valuation buyer with the lowest-valuation seller and continue, in order, as long as the buyer's reservation valuation is at least the seller's. Demand and supply cross at a valuation m^*, the matched quantity is the predicted volume $\mathcal{V}^*$, and the area between the two schedules up to $\mathcal{V}^*$ is the gains-from-trade surplus $\mathcal{S} \geq 0$.

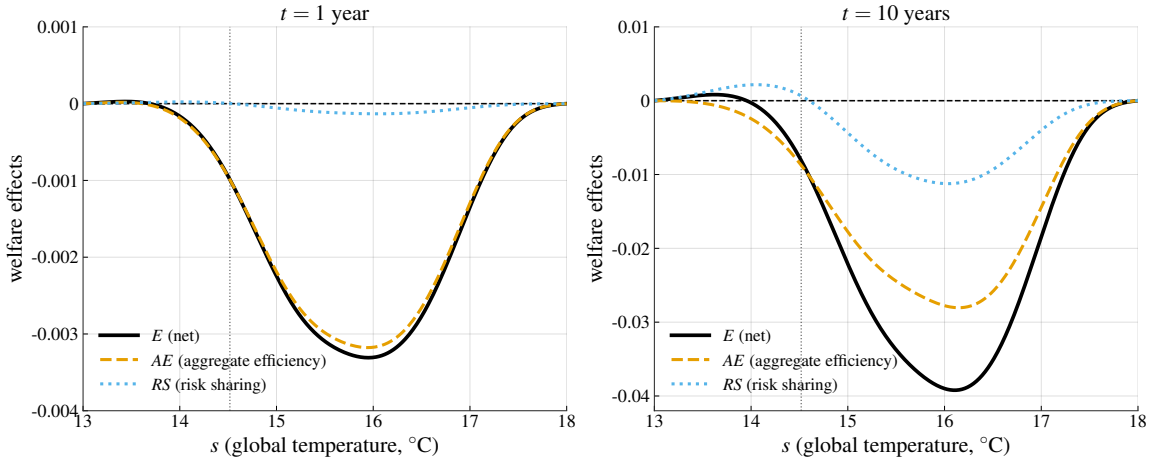


Figure 5: Aggregate-Efficiency and Risk-Sharing Decomposition

Note: Population-weighted decomposition of the value integrand at $t = 1$ year (left panel) and $t = 10$ years (right panel), under the aggressive scenario (SSP5-8.5). Orange dashed: aggregate efficiency, $\mathrm{AE}(s) = \left(\sum_i \chi^i \omega_t^i\right)\left(\sum_i \chi^i x_\theta^i\right)$; blue dotted: risk sharing, $\mathrm{RS}(s) = \mathbb{C}ov_\chi\left[\omega_t^i, x_\theta^i\right]$; black solid: the net, $\mathrm{AE} + \mathrm{RS}$. The dotted vertical line marks the status quo mean of 14.52°C.

The quantity bound $\mathcal{V} = \min\left\{\sum_i (\chi^i \tilde{x}_\theta^i)^+, \sum_i (-\chi^i \tilde{x}_\theta^i)^+\right\}$ is an upper bound on volume; dispersion in valuations can only reduce trade, so $\mathcal{V}^* \leq \mathcal{V}$, and the net $\sum_i \chi^i \tilde{x}_\theta^i$ — the AE residual — is untradable, since it finds no counterparty at any price.

Consider first the book at the status quo mean, $s^* = \bar{s} \approx 14.52$°C, where 102 of the 162 countries are buyers (damaged by warming) and 60 are sellers. Two features stand out. First, at short horizons there are *no gains from trade*: at $t = 1$ and $t = 5$ the demand schedule lies entirely below the supply schedule and $\mathcal{V}^* = 0$. At $s^* = \bar{s}$ the global-to-local mapping leaves every country at its own historical mean temperature, so the damage term in $u'(c_t^i(s^*))$ vanishes and (because the level and trend factors cancel against the normalizer) reservation valuations differ only through each country's expected marginal utility under the status quo, $m_t^i(s^*) = \beta^t / \mathbb{E}_{\underline{F}}\left[e^{-\gamma t\left(h^i\left(T^i(s)\right) - h^i\left(\bar{T}^i\right)\right)}\right]$. A country heavily exposed to the temperature lottery has high expected marginal utility and hence a low reservation valuation for consumption at the undamaged mean state: for it, s^* is a relatively good state. Exposure is ordered by climate — the cross-sectional correlation between m_t^i and $\bar{T}^i$ is about -0.95 at $t = 1$ and weakens with the horizon — so the hot, under-provisioned buyers value the claim systematically below the cool sellers, and the schedules do not cross. Trade emerges only as compounding widens the dispersion of valuations within each side, beginning with the mildly exposed countries near their growth optimum, whose valuations sit near β^t on both sides of the book: $\mathcal{V}^*$ rises from 4.2×10^{-4}

at $t = 10$ to 0.142 at $t = 100$. Second, trade is dwarfed by aggregate exposure: the net is large and negative at every horizon (-1.29 at $t = 100$), so the world as a whole is a net buyer of the s^*-claim, under-provisioned and unable to import resources because the economy is closed. Most of each country's position is undiversifiable aggregate exposure that the security reallocates but cannot remove; the tradable, welfare-relevant part is the small portion that crosses.

Holding $t = 50$ fixed and sweeping the evaluation state from 14.5 to 18°C (Figure 6, which displays four of these states) shows where a temperature-sharing security generates the most trade. As s^* warms, buyers increase from 102 to 127 and sellers fall from 60 to 35 by 17.5°C; from about 15°C onward, the damaged buyers value the claim so far above the few cool sellers that the entire short side clears and $\mathcal{V}^* = \mathcal{V}$. Volume and surplus peak at *different* states. Predicted volume $\mathcal{V}^*$ is largest between 15 and 15.5°C and falls thereafter as the warm tail thins, whereas the surplus $\mathcal{S}$ peaks near 16.5°C (0.196, against 0.015 at 15°C): there, fewer units trade, but each crosses a wide valuation gap between intensely exposed buyers and cool sellers. Since the welfare-relevant object is the surplus rather than the volume, the security is most valuable to introduce at warm but not extreme states around 16.5°C, well above where trading volume is highest. At 18°C both scenarios assign the state negligible mass, $\left(\underline{F} - \overline{F}\right)(s^*) \to 0$, and the book is empty: there is nothing to hedge.

Sweeping s^* over a fine grid on $[13, 18]^\circ$ C — 141 states, about 0.036°C apart — and clearing the double auction at each state traces the predicted volume $\mathcal{V}^*(s^*)$ and the gains-from-trade surplus $\mathcal{S}(s^*)$ as curves (Figure 7). Three features stand out. Predicted volume is hump-shaped and peaks just above the status quo mean, near 14.9°C at every horizon. It falls to zero at both ends of the support: the cdf gap $\left(\underline{F} - \overline{F}\right)(s^*)$ vanishes at both edges, which empties the book, and on the cool side the schedules additionally fail to cross — buyers still outnumber sellers there, but the same ordering of reservation valuations that prevents trade at the status quo mean at short horizons keeps the demand schedule below supply, so $\mathcal{V}^* = 0$ even where the bound $\mathcal{V}$ is positive. The level scales steeply with the horizon, from a peak of about 1.7×10^{-3} at $t = 1$ to about 2.2×10^{-2} at $t = 10$. The gap between the quantity bound $\mathcal{V}$ and the realized $\mathcal{V}^*$ is concentrated on the cool side, where no trade clears at all; the two curves coincide on the warm side, where the intensely exposed buyers out-value the few cool sellers and the entire short side clears. The right panel shows the surplus $\mathcal{S}(s^*)$. It peaks near 16°C at every horizon, moving slightly warmer as the horizon lengthens, about 1°C warmer than the volume peak, and

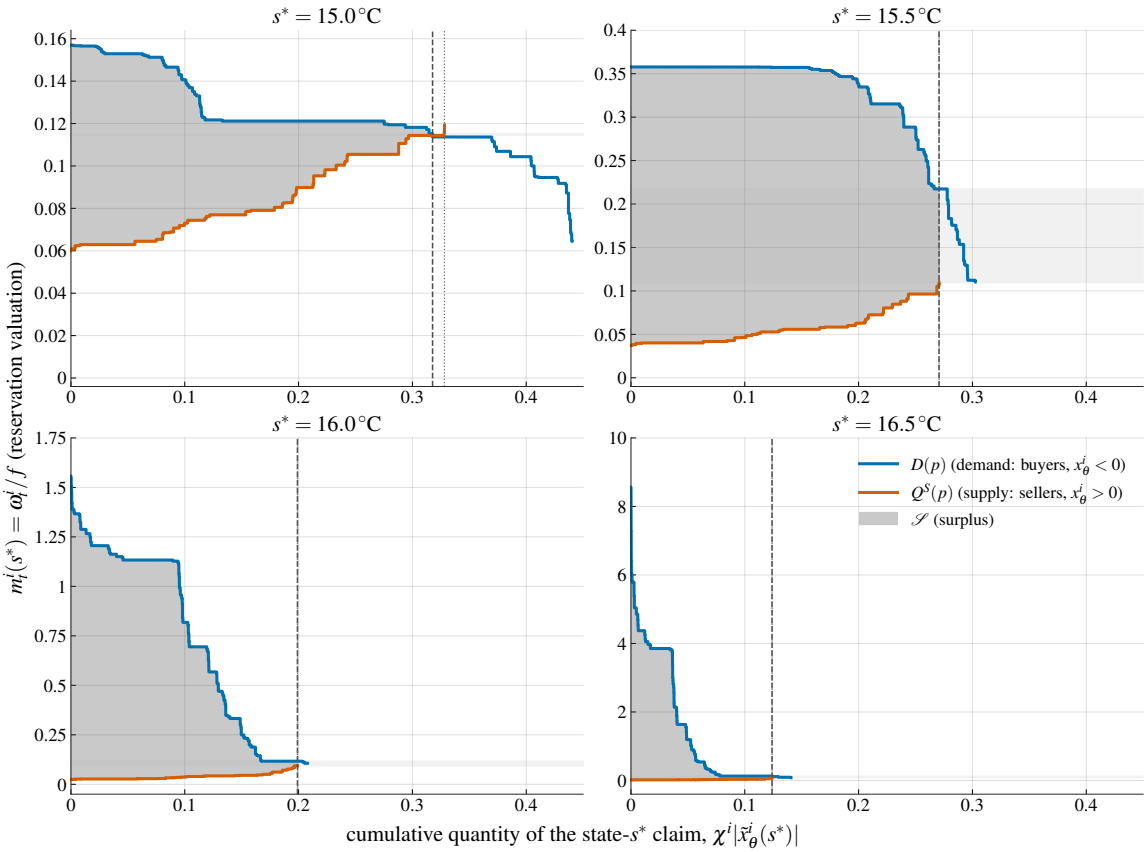


Figure 6: Supply and Demand for State Securities

Note: Supply and demand schedules for the s^*-claim across the 162 countries at four states $s^* \in \{15.0, 15.5, 16.0, 16.5\}^\circ$ C and the fixed horizon $t = 50$ years, under the aggressive scenario (SSP5-8.5). Each country is one limit order: its side is the sign of x_θ^i, its quantity is $\chi^i \left|\tilde{x}_\theta^i\right|$, and its limit price is m_t^i. Blue: demand; red: supply; shaded area: surplus $\mathcal{S}$ up to the predicted volume $\mathcal{V}^*$ (dashed vertical line); dotted vertical line: quantity bound $\mathcal{V}$.

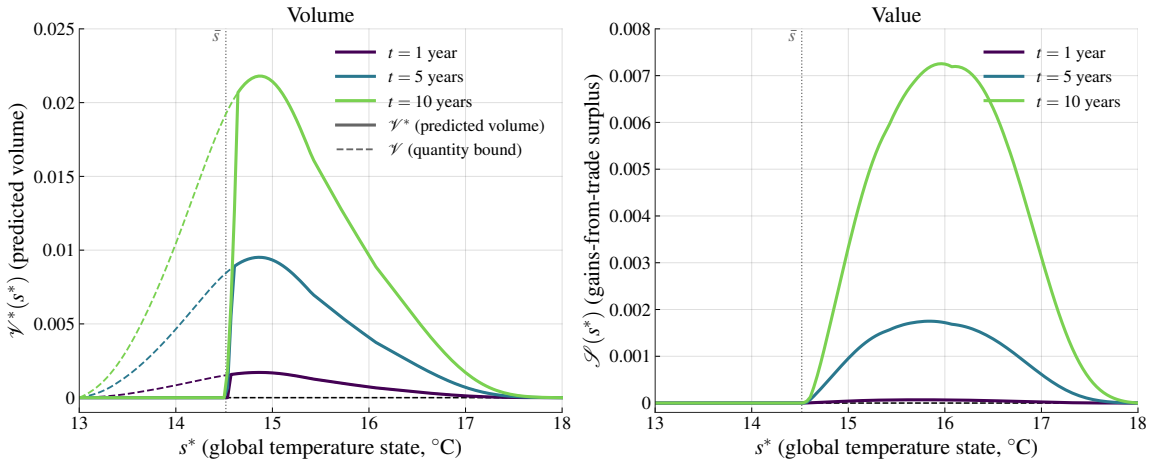


Figure 7: Predicted Hedging Volume and Surplus

Note: Left panel: predicted hedging volume $\mathcal{V}^*(s^*)$ (solid) and quantity bound $\mathcal{V}(s^*)$ (dashed); right panel: gains-from-trade surplus $\mathcal{S}(s^*)$. Both are functions of the global temperature state, swept over 141 states on $[13, 18]^\circ$ C at horizons $t \in \{1, 5, 10\}$ years, under the aggressive scenario (SSP5-8.5), with $\gamma = 2$ and $\beta = 0.96$. The cross-section is cleared as a double auction at each s^*; only matched volume is retained. The surplus is the area between the demand and supply schedules up to the predicted volume. Dotted vertical line: status quo mean $\bar{s} \approx 14.52°$C.

scales with the horizon in the same way (from about 7×10^{-5} at $t = 1$ to about 7×10^{-3} at $t = 10$): fewer units trade at warm states, but each crosses a wider valuation gap, so the state where the security is most valuable to introduce is warmer than the state where it trades most.

Remark 3 (*Lessons from the Hedging Portfolio*). The hedging analysis yields three lessons. First, climate change reallocates large gross exposure across countries that sit on opposite sides of the growth-maximizing temperature, yet only a small part of that exposure is tradable: most of it is an undiversifiable aggregate position that the world holds collectively. Second, the tradable part is small at short horizons — where reservation valuations are compressed toward β^t and ordered by climate exposure, so the damaged buyers value the claim below the cool sellers — and grows only as damages compound. Third, the welfare-relevant object is the gains-from-trade surplus, which peaks at warm but not extreme states, rather than the volume, which peaks near the status quo mean.

Three features of the analysis rely on probability pricing. First, the perturbation (14) is a nonparametric shift of the entire distribution, so the valuation uses the full cdf gap $\underline{F} - \overline{F}$ rather than a change in the mean or in the parameters of a given family; two scenarios with similar means but different tails receive different prices. Second, the equivalence

between probability pricing and cash-flow pricing converts a change in the probabilities of warming into a state-contingent payoff, $x_\theta^i(s)$, that does not depend on preferences. That payoff is what makes hedging demand, trading volume, and gains from trade computable: without it, a change in probabilities has no counterpart in the space of securities. Third, because the same payoff is priced with each country's state prices, the world valuation decomposes into an aggregate-efficiency and a risk-sharing component, which identifies the part of the world's valuation that cross-country trade could reallocate.

4 Information: Probability Pricing of Public Information

A change in the information available to an agent is a change in probabilities. This makes the valuation of information a natural application of the probability pricing framework: changes in information are valued through the probability changes they induce. Our second application studies an economy with ex-ante identical individuals in which insurance markets open only after a public signal is realized, following Hirshleifer (1971). In this environment, more precise public information makes everyone worse off because it resolves uncertainty before agents can insure against it. Probability pricing yields three results.

First, we show that the valuation of a change in public information is determined by two distinct channels. The first is a *reallocation channel*: changes in public information modify equilibrium consumption allocations for each signal realization. In particular, as the signal becomes less informative, consumption becomes more similar across individuals for a given signal realization, a welfare-improving effect. This is the familiar mechanism emphasized in the literature. The second is a *probability channel*, a previously unexplored channel that probability pricing allows us to study. For a given mapping between consumption allocations and signal realizations, a change in information changes the probabilities of observing different signal realizations, which has an ambiguous impact on welfare.

Second, we characterize when this newly identified probability channel is welfare-improving or welfare-reducing. The probability channel is welfare-improving when a more precise signal *compresses* the distribution of signal realizations, reducing the likelihood of extreme realizations in which consumption allocations are most unequal. It is welfare-reducing when a more precise signal instead *disperses* that distribution. In a pure endowment economy, however, the reallocation channel always dominates, so more precise

public information is always welfare-reducing.

Finally, when production decisions respond to the public signal, probability pricing quantifies separately the risk-sharing and production-efficiency consequences of changes in the precision of public information. This yields a measure of the relative welfare importance of the two channels.

4.1 Endowment Economy

Environment. We consider an economy with three dates, $t \in \{0, 1, 2\}$, a single consumption good, and two (types of) ex-ante identical individuals, indexed by $i \in \{A, B\}$. Individuals consume at date 2, after a binary state $s \in \{a, b\}$ is realized with equal prior probabilities: $\pi(s) = \frac{1}{2}$. The aggregate endowment is fixed at $n_2(s) = 1$. Individual A's endowment shares are

$$\chi^A(a) = \frac{1+\delta}{2} \quad \text{and} \quad \chi^A(b) = \frac{1-\delta}{2},$$

where $\chi^A(s) + \chi^B(s) = 1$. Therefore, there is no aggregate risk and $\delta \in (0, 1)$ modulates idiosyncratic risk.

At date 1, before trade takes place, a public signal ξ about the state s is realized, with a conditional cdf and pdf respectively denoted by $F(\xi \mid s)$ and $f(\xi \mid s)$. The signal takes the form

$$\xi \mid s = a \sim \mathcal{N}(\nu, \sigma^2) \quad \text{and} \quad \xi \mid s = b \sim \mathcal{N}(-\nu, \sigma^2),$$

where we refer to the parameter $\nu \geq 0$ as the *strength* of the signal and $\sigma > 0$ as the *noise* of the signal. The strength-to-noise ratio $\rho \equiv \nu/\sigma$ summarizes the informativeness of the signal, with informativeness increasing in ρ.[14]

Individual i has expected utility preferences that can be expressed recursively as

$$V^i = \int V_1^i(\xi) f(\xi; \theta) d\xi \quad \text{where} \quad V_1^i(\xi) = \sum_s u\left(c_2^i(s, \xi)\right) \Pr(s \mid \xi),$$

and $f(\xi; \theta)$ denotes the marginal density of the signal and $\Pr(s \mid \xi)$ the posterior distribution of the state s given the signal ξ. Individual i's expected utility V^i is given by the expected continuation value $V_1^i(\xi)$ at date 1 when the public signal is ξ. In turn, $V_1^i(\xi)$

[14] To see this, note that the posterior log-likelihood ratio $\log \frac{\Pr(s=a|\xi)}{\Pr(s=b|\xi)} = 2\rho z$, after standardizing $z \equiv \xi/\sigma$. Moreover, if $\rho_1 < \rho_2$, it is easy to show that the ρ_1 experiment is a garbling of the ρ_2 experiment in the Blackwell (1953) sense.

is computed as the individual's expected utility from consumption $c_2^i(s,\xi)$ at date 2 when the state is s and the date-1 signal is ξ.[15] We assume isoelastic preferences throughout:

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}.$$

As in Hirshleifer (1971), there are no markets at date 0, while at date 1, after ξ is observed, there are complete Arrow–Debreu markets against the state s. Markets are therefore complete from date 1 onward, but incomplete between dates 0 and 1, because individuals cannot insure against the realization of the signal itself. Hence, individual i faces a date-1 budget constraint given by

$$\sum_s q_1(s \mid \xi)\, x_1^i(s \mid \xi) = 0, \ \forall \xi,$$

where $q_1(s \mid \xi)$ denotes the price of an Arrow–Debreu security that pays at date 2 in state s and $x_1^i(s \mid \xi)$ denotes type-i's holdings of that security when the signal realization is ξ. At date 2, consumption is given by

$$c_2^i(s,\xi) = \chi^i(s)\, n_2(s) + x_1^i(s \mid \xi), \ \forall (s,\xi),$$

where market clearing requires that $\sum_i x_1^i(s \mid \xi) = 0, \forall (s,\xi)$.

Equilibrium. The equilibrium definition is standard. Because there is no aggregate risk and beliefs are common, the date-1 markets ensure full insurance conditional on the signal: equilibrium consumption is independent of s given ξ, state prices are proportional to posterior probabilities, and each individual consumes the conditional expectation of their endowment, $c_2^i(s,\xi) = c_2^i(\xi) = \mathbb{E}[\chi^i(s) \mid \xi]$, so

$$c_2^A(\xi) = \frac{1}{2} + \frac{\delta}{2}(2p(\xi) - 1), \tag{19}$$

where we denote the posterior probability of state a by

$$p(\xi) = \Pr(s = a \mid \xi) = \left(1 + \exp\{-2\nu\xi/\sigma^2\}\right)^{-1}.$$

[15] This formulation also implies the ex-ante expected utility $V^i = \int \sum_s u\left(c_2^i(s,\xi)\right) \tilde{f}(s,\xi)\, d\xi$, where $\tilde{f}(s,\xi)$ is the joint distribution of s and ξ. Expected utility rules out preferences for early or late resolution of uncertainty and therefore excludes direct "psychic" value of the timing of information. There is scope to further explore these additional sources of value in future work.

4.2 Reallocation vs. Probability Channels

We can now study the welfare impact of a perturbation to either the strength or the noise of the signal, $\theta \in \{\nu, \sigma\}$. We write $\partial/\partial\theta$ for derivatives with respect to θ in this section, since the signal distribution $F(\xi \mid s; \theta)$ also depends on the signal and the state. Since individuals only consume at date 2, we express welfare changes in units of uncontingent date-2 consumption, normalizing utility gains by $\lambda^i = \int \sum_s u'\left(c_2^i(s,\xi)\right) f(\xi \mid s;\theta)\, \pi(s)\, d\xi$.[16] Following Equation (7), the impact of the perturbation on individual i is

$$\frac{\frac{dV^i}{d\theta}}{\lambda^i} = \int \sum_s \omega_2^i(s,\xi) \left(\underbrace{\frac{\partial c_2^i(\xi;\theta)}{\partial\theta}}_{\text{Reallocation}} + \underbrace{\frac{\frac{\partial(1-F(\xi|s;\theta))}{\partial\theta}}{f(\xi \mid s;\theta)} \frac{\partial c_2^i(\xi;\theta)}{\partial\xi}}_{\text{Probability}} \right) d\xi, \tag{20}$$

where $\omega_2^i(s,\xi) = \frac{\pi(s) f(\xi|s;\theta) u'\left(c_2^i(\xi)\right)}{\lambda^i}$ is individual i's state price at any history with signal ξ. Summing over the two ex-ante identical individuals to compute efficiency gains yields

$$\Xi^E \equiv \sum_i \frac{\frac{dV^i}{d\theta}}{\lambda^i} = \Xi_c + \Xi_\xi,$$

where

$$\Xi_c = \sum_i \int \sum_s \omega_2^i(s,\xi) \frac{\partial c_2^i(\xi;\theta)}{\partial\theta} d\xi \qquad \text{(Reallocation)} \tag{21}$$

$$\Xi_\xi = \sum_i \int \sum_s \omega_2^i(s,\xi) \frac{\frac{\partial(1-F(\xi|s;\theta))}{\partial\theta}}{f(\xi \mid s;\theta)} \frac{\partial c_2^i(\xi;\theta)}{\partial\xi} d\xi. \qquad \text{(Probability)} \tag{22}$$

Throughout, each channel is signed per unit of the stated perturbation; for a reduction in noise, $d\sigma < 0$, we report $\Xi^{-\sigma} \equiv -\Xi^\sigma$ (Appendix C), so that on both margins the channels are signed in the direction of more informative signals. The reallocation term Ξ_c captures the endogenous adjustment of consumption allocations conditional on the signal realization. The probability term Ξ_ξ captures the effect of changes in the distribution of signal realizations at a given signal-to-consumption mapping, using the consumption-equivalent cash flows of Proposition 1.

[16] Because individuals are ex-ante identical, $\lambda^i = \lambda$, $\forall i$.

Reallocation channel. We now show that increases in signal informativeness, either through an increase in signal strength, $d\nu > 0$, or a reduction in signal noise, $d\sigma < 0$, are associated with a negative reallocation channel. Because the aggregate endowment is fixed, $c_2^A(\xi) + c_2^B(\xi) = n_2$, so $\frac{\partial c_2^B(\xi)}{\partial \theta} = -\frac{\partial c_2^A(\xi)}{\partial \theta}$ and the reallocation term can be expressed

$$\Xi_c = \int \left[\omega_2^A(\xi) - \omega_2^B(\xi)\right] \frac{\partial c_2^A(\xi;\theta)}{\partial \theta} d\xi, \tag{23}$$

where $\omega_2^i(\xi) = \frac{f(\xi;\theta)u'\left(c_2^i(\xi)\right)}{\lambda^i}$. This channel is entirely driven by the difference between the two individuals' state prices, which is nonzero exactly when the signal is informative: individual A's marginal utility is low precisely after signals that are favorable to A. Proposition 2 establishes that this term is always negative when information becomes more precise, either by increasing the strength of the signal or reducing its noise.

Proposition 2 (**Sign of the Reallocation Channel**). *For $\nu > 0$, an increase in signal informativeness, either through an increase in signal strength, $d\nu > 0$, or a reduction in signal noise, $d\sigma < 0$, is associated with a negative reallocation channel:*

$$\Xi_c^{\nu} < 0 \qquad \textit{and} \qquad \Xi_c^{-\sigma} < 0.$$

More precise information makes the equilibrium allocation more dispersed given a signal realization, hurting in particular the already least favored individuals, so the reallocation channel is negative. This channel is a manifestation of *constrained inefficiency* — as in Geanakoplos and Polemarchakis (1986) — since agents face incomplete markets against the realization of the signal: if complete markets against the signal opened at date 0, state prices would be identical and the reallocation channel of changes in public information would be zero.

Probability channel. The sign of the probability channel is more subtle. Because aggregate consumption is fixed, the probability term can be expressed as

$$\Xi_\xi = \int \left[\omega_2^A(\xi) - \omega_2^B(\xi)\right] \frac{\partial c_2^A(\xi;\theta)}{\partial \xi} \frac{\frac{\partial(1-F(\xi;\theta))}{\partial \theta}}{f(\xi;\theta)} d\xi. \tag{24}$$

The probability channel depends once again on the difference in state prices, $\omega_2^A(\xi) - \omega_2^B(\xi)$, which inherits the sign of $-\xi$. When $\xi = 0$, $c_2^A = c_2^B = 1/2$ and the state prices coincide. Above zero, A does relatively better, so the difference $\omega_2^A(\xi) - \omega_2^B(\xi)$ is negative; below

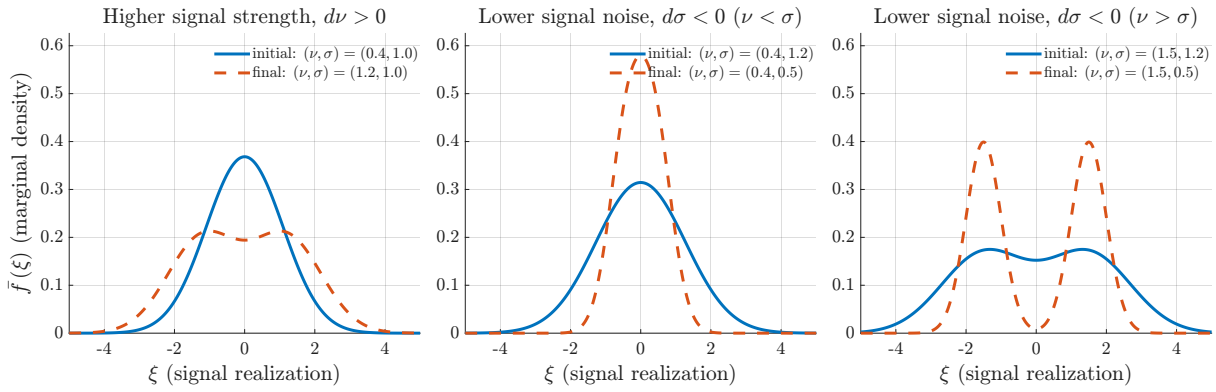


Figure 8: Marginal Signal Distribution

Note: The solid blue curve shows the initial marginal density, $f(\xi;\theta) = \frac{1}{2\sigma}\left[\phi\left(\frac{\xi-\nu}{\sigma}\right) + \phi\left(\frac{\xi+\nu}{\sigma}\right)\right]$, where $\phi(\cdot)$ denotes the standard normal density, and the dashed orange curve shows the density after the indicated increase in informativeness. The left panel increases the signal strength from $\nu = 0.4$ to $\nu = 1.2$ at $\sigma = 1$: probability mass moves outward and extreme signals become more likely, so $\Xi_\xi < 0$. The middle panel reduces the signal noise from $\sigma = 1.2$ to $\sigma = 0.5$ at $\nu = 0.4$, in the unimodal regime $(\nu < \sigma)$: probability mass concentrates near $\xi = 0$, so $\Xi_\xi > 0$. The right panel applies the same noise reduction at $\nu = 1.5$, in the bimodal regime $(\nu > \sigma)$: the center of the distribution hollows out toward the conditional means $\pm\nu$, so $\Xi_\xi < 0$ in this case.

zero it is positive. Given $c_2^A(\xi)$ in (19), a higher signal shifts consumption toward A since $\frac{\partial c_2^A}{\partial \xi} \geq 0$. This derivative is largest at $\xi = 0$ and vanishes as $|\xi| \to \infty$. Because $\frac{\partial c_2^A}{\partial \xi}$ is positive, the sign of Ξ_ξ is determined by the normalized survival change of the marginal distribution of the signal.

An increase in the signal strength, $d\nu > 0$, spreads the conditional means apart and shifts the marginal signal distribution toward the tails. Hence, $\frac{\partial(1-F(\xi))}{\partial \nu}$ has the sign of ξ, opposite the state-price difference, so this shift raises the probability of the extreme signals, making Ξ_ξ negative. A reduction in the signal noise $d\sigma < 0$ shrinks the conditional distributions toward $\pm\nu$. In this case, there are two possible scenarios. First, if the conditionals overlap $(\nu < \sigma)$ — the unimodal regime — a reduction in σ concentrates probability near zero, making Ξ_ξ positive. Second, if they are separated $(\nu > \sigma)$ — the bimodal regime — a reduction in σ hollows the center of the marginal distribution of the signal toward the conditional means $\pm\nu$, turning Ξ_ξ negative once the strength-to-noise ratio $\rho = \nu/\sigma$ is sufficiently large.

Figure 8 illustrates these three configurations. The left panel increases signal strength, $d\nu > 0$, which spreads the conditional means apart and disperses the marginal distribution of signals. The middle panel reduces noise, $d\sigma < 0$, while $\nu < \sigma$, which compresses the

marginal distribution toward zero. The right panel reduces noise, $d\sigma < 0$, while $\nu > \sigma$, which hollows the center and concentrates probability around $\pm\nu$. Together, the three panels show that greater informativeness does not by itself determine the sign of the probability channel: the sign depends on both the margin through which information improves and the regime from which the perturbation begins, and its ultimate impact is determined by the normalized survival change of the distribution of the signal, $\frac{\frac{\partial(1-F(\xi;\theta))}{\partial\theta}}{f(\xi;\theta)}$.

Proposition 3 establishes that the sign of the probability channel depends on whether informativeness changes the marginal distribution of the signal through changes in the strength or the noise of the signal.

Proposition 3 (**Sign of the Probability Channel**). *For $\nu > 0$, the sign of the probability channel depends on the source of the increase in signal informativeness:*

(i) An increase in signal strength, $d\nu > 0$, makes the marginal distribution more dispersed by making signal realizations more extreme, so $\frac{\partial(1-F(\xi))}{\partial\nu}$ has the sign of ξ. The probability channel is then negative, $\Xi_\xi^\nu < 0$.

(ii) A reduction in signal noise, $d\sigma < 0$, makes the marginal distribution more concentrated around points that depend on the regime. In the unimodal regime ($\nu < \sigma$) signal realizations concentrate around central signals, compressing the marginal distribution of the signal, and the probability channel is then positive, $\Xi_\xi^{-\sigma} > 0$. In the bimodal regime ($\nu > \sigma$) signal realizations concentrate around $\pm\nu$, hollowing the center of the marginal distribution, so $\Xi_\xi^{-\sigma}$ can take either sign: it remains positive when the strength-to-noise ratio $\rho = \nu/\sigma$ is sufficiently close to one, and becomes negative when $\rho = \nu/\sigma$ is larger than a threshold $\rho^*(\gamma, \delta) > 1$, which we characterize in Appendix C.

The case $\Xi_\xi > 0$ is worth emphasizing. There, the probability channel value of more informative signals is positive and works *against* the Hirshleifer effect. Intuitively, a more compressed marginal distribution of signals is welfare-improving, as extreme signal realizations — which are associated with unequal consumption patterns — are less likely.

Results. We now quantify the two channels. Table 1 reports the baseline parameterization, which requires only three parameter choices: the risk-aversion coefficient γ, the idiosyncratic endowment risk δ, and the strength-to-noise ratio ρ.[17] We set $\gamma = 2$, a

[17] Under the change of variables $z = \xi/\sigma$, the equilibrium allocation is $c_2^A(z) = \frac{1}{2} + \frac{\delta}{2}\tanh(\rho z)$, implying that the allocation depends on (ν, σ) only through the strength-to-noise ratio $\rho = \nu/\sigma$. This allows us to

Description	Symbol	Value
Relative risk aversion	γ	2
Idiosyncratic endowment risk	δ	0.4
Strength-to-noise ratio	$\rho = \frac{\nu}{\sigma}$	0.75

Table 1: Endowment Economy

Note: Baseline parameters for the endowment economy. The aggregate endowment $n_2 = 1$ and the signal noise $\sigma = 1$ are normalizations, so the implied signal strength is $\nu = \rho\sigma = 0.75$. The implied discrimination of the public signal is AUC $= \Phi\left(\sqrt{2}\rho\right) = 0.86$, and the sign threshold on the noise margin, $\rho^*(\gamma, \delta) \approx 1.22$, solves $\Xi_\xi = 0$ given (γ, δ).

standard risk aversion value, and $\delta = 0.4$, a value in line with the cross-sectional dispersion of consumption once the state is revealed. The strength-to-noise ratio ρ maps one-to-one into the discrimination of the signal, measured by the area under its receiver operating characteristic (ROC) curve (AUC): AUC $= \Phi\left(\sqrt{2}\rho\right)$, where $\Phi(\cdot)$ denotes the standard normal cdf. We thus choose $\rho = 0.75$ to match the AUC $= 0.86$ estimated by Fuster, Goldsmith-Pinkham, Ramadorai, and Walther (2022) for machine-learning screening models. In this baseline calibration, the two channels carry opposite signs on the noise margin: the probability channel Ξ_ξ reverses sign only above $\rho^* \approx 1.22$.

Figure 9 illustrates the welfare effect of a marginal reduction in signal informativeness along two margins, consistent with Propositions 2 and 3. The left panel lowers signal strength, $d\nu < 0$, while holding the signal noise σ fixed. The right panel increases the signal noise, $d\sigma > 0$, while holding the signal strength fixed. In both panels, a positive value of Ξ^E or its components means that less informative signals are welfare-improving, so the Hirshleifer effect reads $\Xi^E > 0$.

Consistent with Proposition 2, the reallocation channel Ξ_c is positive on both margins: worse information, whether through weaker or noisier signals, improves risk sharing in a constrained-inefficient market. In line with Proposition 3, the probability channel Ξ_ξ instead depends on the margin. On the strength margin, it is positive throughout: a lower ν compresses the marginal distribution of the signal, making the extreme signal realizations — which are associated with unequal consumption — less likely. On the noise margin, its sign is regime-dependent: additional noise first fills in the hollow center of the bimodal signal distribution and then disperses it, so Ξ_ξ switches from positive to negative at $\sigma = \nu/\rho^* \approx 0.62$, inside the bimodal regime ($\sigma < \nu$) and before the distribution becomes

normalize $\sigma = 1$.

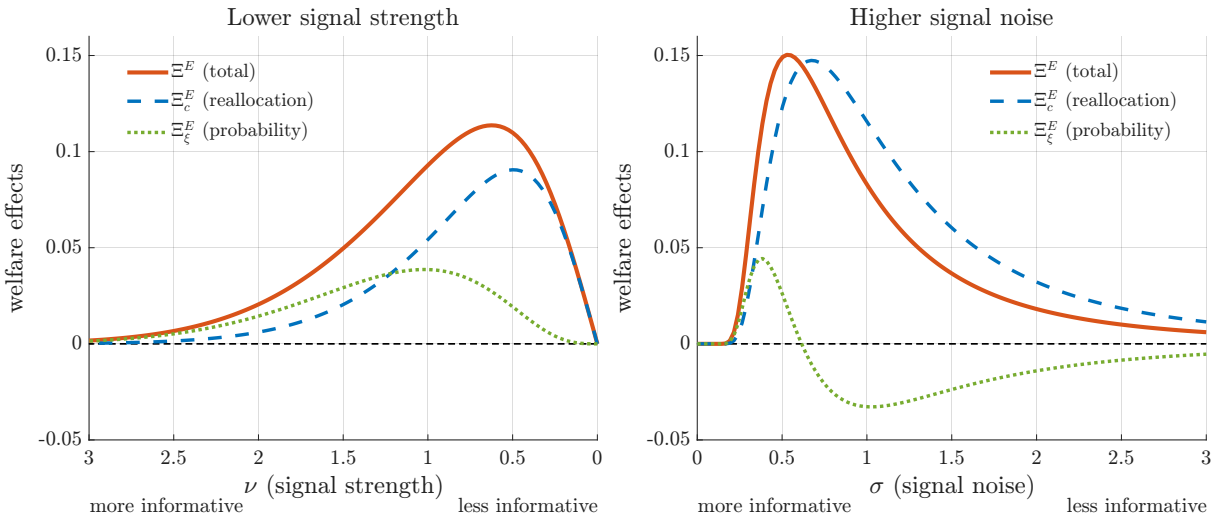


Figure 9: Probability Pricing of Public Information: Endowment Economy

Note: Welfare effects of a marginal worsening of information, where $\xi \mid a \sim \mathcal{N}\left(\nu, \sigma^2\right)$, $\xi \mid b \sim \mathcal{N}\left(-\nu, \sigma^2\right)$, $\delta = 0.4$, $\gamma = 2$, and aggregate endowment $n_2 = 1$. The left panel reports the effect of $d\nu < 0$ over $\nu \in [0, 3]$ while holding $\sigma = 1$. The right panel reports the effect of $d\sigma > 0$ over $\sigma \in [0, 3]$ while holding $\nu = 0.75$.

unimodal ($\sigma > \nu$). In line with Hirshleifer (1971), the total effect Ξ^E is positive at every ν and σ, always operating through the reallocation channel, though the probability channel counteracts it when less informative signals are associated with a more dispersed marginal distribution of signals.

4.3 Production Economy

Environment. We now allow production decisions to respond to the public signal. The environment is that of Section 4.1, except that each individual, after observing the signal, operates a technology that yields $R(s)\, k(\xi) - \frac{\kappa}{2} k(\xi)^2$ units of date-2 consumption, where $R(a) = e^{\mu}$ and $R(b) = e^{-\mu}$, the parameter $\mu \geq 0$ measures the productivity spread across states, and $\kappa > 0$ parameterizes adjustment costs; $\mu = 0$ recovers the endowment economy. Since markets against s are complete once ξ is realized, equilibrium consumption features full insurance,[18] and investment $k(\xi)$ satisfies the first-order condition:

$$\sum_s u'\left(c_2^i(s,\xi)\right)\left(R(s) - \kappa k(\xi)\right) f(\xi \mid s)\, \pi(s) = 0, \ \forall \xi. \tag{25}$$

[18] In this case, $c_2^i(s,\xi) = \alpha^i(\xi)\, c_2(s,\xi)$, where $c_2(s,\xi) = n_2 + 2\left[R(s)\, k(\xi) - \frac{\kappa}{2} k(\xi)^2\right]$ denotes aggregate consumption and $\alpha^i(\xi)$ denotes the value share of individual i, with $\sum_i \alpha^i(\xi) = 1$.

Because markets are complete after the signal, both individuals agree on the investment, which is strictly increasing in the signal for $\mu > 0$, so $k'(\xi) > 0$.

Because aggregate consumption now responds to the signal, we split each channel of Equation (20) into a production-efficiency (PE) component, which captures the response of aggregate consumption, and a risk-sharing (RS) component, which captures the cross-sectional comovement between state prices and consumption responses, following Dávila and Schaab (2025), and as explained in Appendix C:

$$\Xi^E = \underbrace{\Xi_c^{PE} + \Xi_c^{RS}}_{\Xi_c} + \underbrace{\Xi_\xi^{PE} + \Xi_\xi^{RS}}_{\Xi_\xi},$$

where $\Xi^{PE} = \Xi_c^{PE} + \Xi_\xi^{PE}$ and $\Xi^{RS} = \Xi_c^{RS} + \Xi_\xi^{RS}$.

Production Efficiency vs. Risk Sharing. Proposition 4 shows that the production margin operates exclusively through the probability channel, and that its welfare contribution is unambiguously signed.

Proposition 4 (**Production Efficiency**). *In the production economy:*

(i) The production-efficiency component of the reallocation channel is zero: $\Xi_c^{PE} = 0$, so $\Xi^{PE} = \Xi_\xi^{PE}$. Hence, the reallocation channel exclusively contributes to risk sharing.

(ii) For $\mu > 0$, the production-efficiency component of the probability channel is positive (negative) for increases (decreases) in signal informativeness, for both changes in signal strength ($d\nu$) and noise ($d\sigma$).

The first part of Proposition 4 is an envelope property: each individual chooses the investment optimally, so investment responses carry no first-order welfare value at given probabilities, and every $\frac{\partial k}{\partial \theta}$ term vanishes from the reallocation channel. The second part reflects that the classic allocative value of information exclusively operates through the probability channel. The risk-sharing mechanisms explored in the endowment economy also operate when there is production.

Results. Table 2 reports the two additional parameters required by the production economy; all others are unchanged from Table 1. The productivity spread μ determines both the aggregate risk and the productive value of the signal — each vanishes at $\mu = 0$, which recovers the endowment economy — and the adjustment cost κ sets the scale

Description	Symbol	Value
Productivity spread across states	μ	0.4
Adjustment cost	κ	0.5

Table 2: Production Economy

Note: Parameters that the production economy adds to the endowment calibration of Table 1. The two free parameters are the productivity spread μ and the adjustment cost κ; $\mu = 0$ recovers the endowment economy.

of investment. We set $\mu = 0.4$, which implies that aggregate output has mean 3.32 and standard deviation 1.48 at the baseline calibration, against 3.67 and 1.78 at the full-information benchmark. We set $\kappa = 0.5$, at which production supplies about 70% of aggregate output and optimal investment has mean 1.69 and standard deviation 0.32 across signal realizations.

Figure 10 shows the resulting decomposition along the same two margins as Figure 9, now allowing for production. As the first part of Proposition 4 predicts, the production-efficiency component Ξ^{PE} is due entirely to the probability channel and, in line with the second part, it is negative on both margins: worse information, whether through lower strength or higher noise, misaligns investment with the realized state. The risk-sharing component Ξ^{RS} collects the reallocation channel Ξ_c^{RS} and the risk-sharing part of the probability channel Ξ_ξ^{RS}, which the figure plots separately: Ξ_c^{RS} remains positive on both margins, as in Proposition 2, while Ξ_ξ^{RS} is positive on the strength margin and switches sign on the noise margin, as in Proposition 3. In this parameterization, the production-efficiency component dominates at every ν and σ, so the total is negative throughout: better public information raises welfare even though the Hirshleifer effect remains fully active. Which force dominates is a quantitative question: as $\mu \to 0$ or $\kappa \to \infty$ the production margin shuts down and the endowment-economy results reemerge, so overturning the Hirshleifer result requires a sufficiently strong allocative role of information relative to the underlying idiosyncratic risk.

Remark 4 (*Role of Probability Pricing*). Separating the production-efficiency and risk-sharing components of a change in the precision of the public signal requires expressing the impact of changes in information in consumption equivalents, which relies on the probability pricing formulation of Proposition 1. Therefore, while it was well understood that changes in the precision of public information impact production efficiency and risk

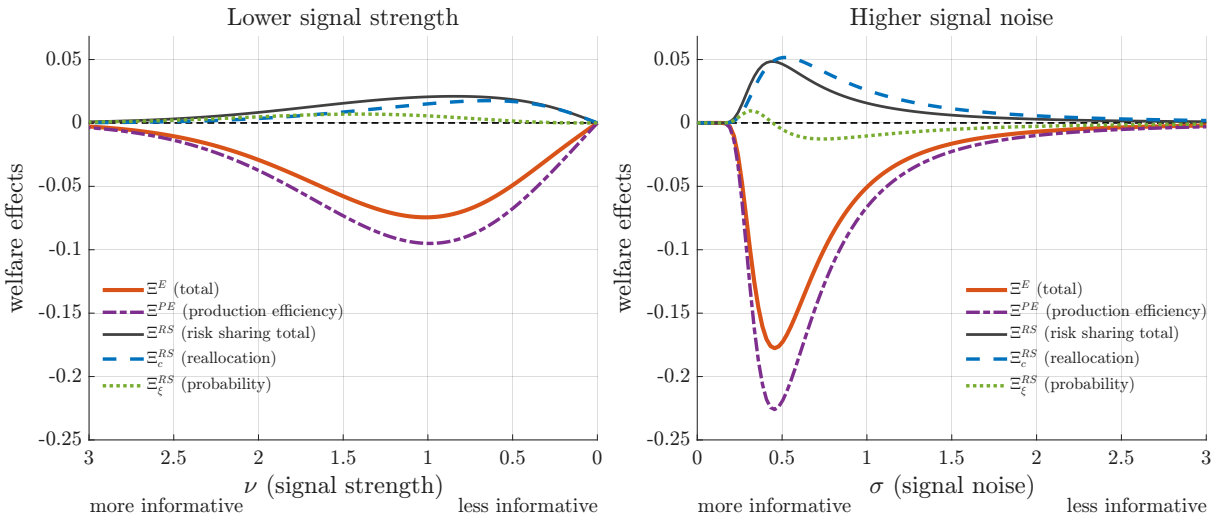


Figure 10: Probability Pricing of Public Information: Production Economy

Note: Welfare effect of a marginal worsening of information, $\Xi^E = \Xi^{PE} + \Xi^{RS}$, in the production economy with parameters $\mu = 0.4$ and $\kappa = 0.5$; all other parameters as in Figure 9. The left panel plots the welfare effect of a reduction in signal strength, $d\nu < 0$, as a function of ν, holding $\sigma = 1$; the right panel plots the welfare effect of an increase in signal noise, $d\sigma > 0$, as a function of σ, holding $\nu = 0.75$.

sharing, probability pricing separates the two quantitatively, as illustrated in Figure 10.

5 Conclusion

This paper extends traditional cash-flow pricing to analyze the valuation of changes in probabilities, that is, probability pricing. It shows that an agent's valuation for a marginal change in probabilities is equivalent to the agent's valuation for a hypothetical asset whose cash flows are the product of the change in the normalized survival function and the sensitivity of consumption to the state. That is, probability pricing establishes a formal equivalence between changes in probabilities and changes in consumption (cash-flow pricing). This equivalence makes it possible to construct hedging strategies and stochastic and cross-sectional decompositions for general preferences, distributions, and perturbations.

We present two direct applications of probability pricing to changes in physical probabilities, in the context of climate change, and to changes in information, in the context of changes in the precision of a public signal. In our climate application, the consumption-equivalent cash flows define country-specific portfolios that hedge shifts in the distribution of global mean temperature. Climate change reallocates large gross exposures

across countries that sit on opposite sides of the growth-maximizing temperature, but most of the world's exposure is an aggregate position that cross-country trade cannot net out. Probability pricing predicts that if countries wanted to hedge their asymmetric exposure to temperature changes, markets for temperatures around 14.9°C would see the most trading, while those around 16°C would generate the largest gains from trade.

In our information application, probability pricing identifies a new channel — the probability channel — by which public information impacts efficiency and welfare in a Hirshleifer (1971)-style environment, whose sign depends on whether more precise signals compress or disperse the distribution of signal realizations. Once information conditions real decisions, probability pricing provides a new approach to quantify the production efficiency and risk-sharing effects of changes in public information, opening the door to further quantitative explorations of these forces.

We hope that probability pricing provides a useful framework for future quantitative work in asset pricing, climate economics, information economics, and welfare analysis.

References

Ai, H., and R. Bansal (2018): "Risk Preferences and the Macroeconomic Announcement Premium," *Econometrica*, 86(4), 1383–1430.

Ai, H., R. Bansal, H. Guo, and A. Yaron (2024): "An Arrow-Pratt Theory of Preference for Early Resolution of Uncertainty," *Working Paper*.

Alvarez, F., and U. J. Jermann (2004): "Using Asset Prices to Measure the Cost of Business Cycles," *Journal of Political Economy*, 112(6), 1223–1256.

Angeletos, G.-M., and A. Pavan (2007): "Efficient Use of Information and Social Value of Information," *Econometrica*, 75(4), 1103–1142.

——— (2009): "Policy with Dispersed Information," *Journal of the European Economic Association*, 7(1), 11–60.

Arrow, K. J. (1971): "The Theory of Risk Aversion," in *Essays in the Theory of Risk-Bearing*, pp. 90–120. Markham, Chicago.

Aumann, R. J., and R. Serrano (2008): "An Economic Index of Riskiness," *Journal of Political Economy*, 116(5), 810–836.

Baron, D. P., and R. B. Myerson (1982): "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50(4), 911–930.

Barro, R. J. (2009): "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 99(1), 243–264.

Blackwell, D. (1953): "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics*, 24(2), 265–272.

Bolton, P., and M. Dewatripont (2005): *Contract Theory*. MIT Press.

Bond, P., and D. García (2022): "The Equilibrium Consequences of Indexing," *Review of Financial Studies*, 35(7), 3175–3230.

Burke, M., S. M. Hsiang, and E. Miguel (2015): "Global Non-Linear Effect of Temperature on Economic Production," *Nature*, 527(7577), 235–239.

Campbell, J. Y. (2017): *Financial Decisions and Markets: A Course in Asset Pricing*. Princeton University Press.

Courant, R., and F. John (1989): *Introduction to Calculus and Analysis, Volume I*. Springer-Verlag, New York, revised edn.

Dávila, E., and A. Schaab (2025): "Welfare Assessments with Heterogeneous Individuals," *Journal of Political Economy*, 133(9), 2918–2961.

Dávila, E., and A. Walther (2023): "Prudential Policy with Distorted Beliefs,"

American Economic Review, 113(7), 1967–2006.

Diamond, D. W. (1985): "Optimal Release of Information by Firms," *Journal of Finance*, 40(4), 1071–1094.

Diamond, D. W., and R. E. Verrecchia (1981): "Information Aggregation in a Noisy Rational Expectations Economy," *Journal of Financial Economics*, 9(3), 221–235.

——— (1991): "Disclosure, Liquidity, and the Cost of Capital," *Journal of Finance*, 46(4), 1325–1359.

Eyring, V., et al. (2016): "Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) Experimental Design and Organization," *Geoscientific Model Development*, 9(5), 1937–1958.

Foster, D. P., and S. Hart (2009): "An Operational Measure of Riskiness," *Journal of Political Economy*, 117(5), 785–814.

Fuster, A., P. Goldsmith-Pinkham, T. Ramadorai, and A. Walther (2022): "Predictably Unequal? The Effects of Machine Learning on Credit Markets," *Journal of Finance*, 77(1), 5–47.

Geanakoplos, J., and H. Polemarchakis (1986): "Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market Is Incomplete," in *Uncertainty, Information and Communication: Essays in Honor of Kenneth J. Arrow, Volume III*, ed. by W. P. Heller, R. M. Starr, and D. A. Starrett, pp. 65–95. Cambridge University Press.

Goldstein, I., and L. Yang (2019): "Good Disclosure, Bad Disclosure," *Journal of Financial Economics*, 131(1), 118–138.

Gollier, C. (2001): *The Economics of Risk and Time*. MIT Press.

Gottardi, P., and R. Rahi (2014): "Value of Information in Competitive Economies with Incomplete Markets," *International Economic Review*, 55(1), 57–81.

Grossman, S. J., and J. E. Stiglitz (1980): "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70(3), 393–408.

Gutiérrez, J. M., et al. (2021): "Atlas," in *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, ed. by V. Masson-Delmotte, et al., pp. 1927–2058. Cambridge University Press.

Hansen, L. P., and P. E. Souganidis (2025): "Stochastic Responses and Marginal Valuation," *Proceedings of the National Academy of Sciences*, 122(48).

Harris, I., T. J. Osborn, P. Jones, and D. Lister (2020): "Version 4 of the CRU

TS Monthly High-Resolution Gridded Multivariate Climate Dataset," *Scientific Data*, 7(1).

Hellwig, M. F. (1980): "On the Aggregation of Information in Competitive Markets," *Journal of Economic Theory*, 22(3), 477–498.

Hirshleifer, J. (1971): "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61(4), 561–574.

Ingersoll, J. E. (1987): *Theory of Financial Decision Making*. Rowman & Littlefield.

Iturbide, M., et al. (2022): "Implementation of FAIR Principles in the IPCC: The WGI AR6 Atlas Repository," *Scientific Data*, 9(1).

Kadan, O., and A. Manela (2019): "Estimating the Value of Information," *Review of Financial Studies*, 32(3), 951–991.

Kyle, A. S. (1989): "Informed Speculation with Imperfect Competition," *Review of Economic Studies*, 56(3), 317–355.

Lucas, R. E. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46(6), 1429–1445.

——— (1987): *Models of Business Cycles*. Basil Blackwell.

Luenberger, D. G. (1969): *Optimization by Vector Space Methods*. Wiley.

Martin, I. (2013): "Consumption-Based Asset Pricing with Higher Cumulants," *Review of Economic Studies*, 80(2), 745–773.

Martin, I., and R. S. Pindyck (2015): "Averting Catastrophes: The Strange Economics of Scylla and Charybdis," *American Economic Review*, 105(10), 2947–2985.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995): *Microeconomic Theory*. Oxford University Press.

Mirrlees, J. A. (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38(2), 175–208.

Morris, S., and H. S. Shin (2002): "Social Value of Public Information," *American Economic Review*, 92(5), 1521–1534.

Pavan, A., I. Segal, and J. Toikka (2014): "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica*, 82(2), 601–653.

Pavan, A., S. Sundaresan, and X. Vives (2022): "(In)Efficiency in Information Acquisition and Aggregation Through Prices," *Working Paper*.

Pratt, J. W. (1964): "Risk Aversion in the Small and in the Large," *Econometrica*, 32(1/2), 122–136.

Rostek, M., and M. Weretka (2012): "Price Inference in Small Markets,"

Econometrica, 80(2), 687–711.

Rothschild, M., and J. E. Stiglitz (1970): "Increasing Risk: I. A Definition," *Journal of Economic Theory*, 2(3), 225–243.

Sheather, S. J., and M. C. Jones (1991): "A Reliable Data-Based Bandwidth Selection Method for Kernel Density Estimation," *Journal of the Royal Statistical Society: Series B*, 53(3), 683–690.

Tatem, A. J. (2017): "WorldPop, Open Data for Spatial Demography," *Scientific Data*, 4(1).

Valli, A. (2023): "Weak Derivatives and Sobolev Spaces," in *A Compact Course on Linear PDEs*, pp. 43–55. Springer, 2nd edn.

Veldkamp, L. (2011): *Information Choice in Macroeconomics and Finance*. Princeton University Press.

——— (2023): "Valuing Data as an Asset," *Review of Finance*, 27(5), 1545–1562.

Vives, X. (2011): "Strategic Supply Function Competition with Private Information," *Econometrica*, 79(6), 1919–1966.

——— (2017): "Endogenous Public Information and Welfare in Market Games," *Review of Economic Studies*, 84(2), 935–963.

Online Appendix

A Section 2: Supplementary Material

Proof of Proposition 1 (Probability Pricing).

Proof. Let $U(s) = \beta u(c_1(s))$, so

$$\frac{dU(s)}{ds} = \beta u'(c_1(s)) \frac{dc_1(s)}{ds}.$$

Using integration by parts, Equation (3) can be expressed as

$$\begin{aligned} p_\theta &= \frac{1}{u'(c_0)} \int_{\underline{s}}^{\bar{s}} U(s) \frac{df(s;\theta)}{d\theta} ds = \frac{1}{u'(c_0)} \left[U(s) \frac{dF(s;\theta)}{d\theta} \Big|_{\underline{s}}^{\bar{s}} - \int_{\underline{s}}^{\bar{s}} \frac{dU(s)}{ds} \frac{dF(s;\theta)}{d\theta} ds \right] \\ &= \frac{1}{u'(c_0)} \left[U(\bar{s}) \underbrace{\frac{dF(\bar{s};\theta)}{d\theta}}_{=0} - U(\underline{s}) \underbrace{\frac{dF(\underline{s};\theta)}{d\theta}}_{=0} - \int_{\underline{s}}^{\bar{s}} \frac{dU(s)}{ds} \frac{dF(s;\theta)}{d\theta} ds \right] \\ &= \frac{1}{u'(c_0)} \left[\int_{\underline{s}}^{\bar{s}} \frac{dU(s)}{ds} \frac{d(1 - F(s;\theta))}{d\theta} ds \right] = \int_{\underline{s}}^{\bar{s}} \frac{\beta u'(c_1(s))}{u'(c_0)} \frac{dc_1(s)}{ds} \frac{d(1 - F(s;\theta))}{d\theta} ds, \end{aligned}$$

which corresponds to Equation (4) in the text. We use the fact that for all θ, we have $F(\bar{s};\theta) = 1$ and $F(\underline{s};\theta) = 0$. If the support is unbounded, we assume in addition that $\lim u(c_1(s)) \frac{dF(s;\theta)}{d\theta} = 0$ at both ends of the support, so that the boundary term vanishes.[19] □

Additional Results. Note that defining expected consumption as

$$\mathbb{E}[c_1(s)] = \int_{\underline{s}}^{\bar{s}} c_1(s) f(s;\theta) ds,$$

[19] On a more technical note, the hypothetical cash flows in our analysis are related to *weak derivatives*, and the associated probability price is often well-defined even when standard regularity conditions fail; see, for example, Valli (2023).

we have

$$\frac{d\mathbb{E}\left[c_1\left(s\right)\right]}{d\theta} = \int_{\underline{s}}^{\bar{s}} c_1\left(s\right) \frac{df\left(s;\theta\right)}{d\theta} ds = \int_{\underline{s}}^{\bar{s}} \frac{dc_1\left(s\right)}{ds} \frac{d\left(1-F\left(s;\theta\right)\right)}{d\theta} ds$$
$$= \int_{\underline{s}}^{\bar{s}} \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f\left(s;\theta\right)} \frac{dc_1\left(s\right)}{ds} f\left(s;\theta\right) ds.$$

Note also that

$$p_\theta = \int_{\underline{s}}^{\bar{s}} m\left(s\right) \frac{dc_1\left(s\right)}{ds} \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f\left(s;\theta\right)} f\left(s;\theta\right) ds$$
$$= \int_{\underline{s}}^{\bar{s}} m\left(s\right) f\left(s;\theta\right) ds \int_{\underline{s}}^{\bar{s}} \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f\left(s;\theta\right)} \frac{dc_1\left(s\right)}{ds} f\left(s;\theta\right) ds + \mathbb{C}ov\left[m\left(s\right), \frac{dc_1\left(s\right)}{ds} \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f\left(s;\theta\right)}\right],$$

which corresponds to Equation (6) in the text.

Affine Perturbations. We assume that $s = \mu + \sigma n$. Therefore, the cdf $F\left(s\right)$ is given by

$$F\left(s\right) = H\left(\frac{s-\mu}{\sigma}\right),$$

where $H\left(\cdot\right)$ denotes the cdf over n. Therefore, the pdf $f\left(s\right)$ is given by

$$f\left(s\right) = \frac{d}{ds} H\left(\frac{s-\mu}{\sigma}\right) = \frac{1}{\sigma} h\left(\frac{s-\mu}{\sigma}\right).$$

A marginal increase in μ implies that

$$\frac{dF\left(s\right)}{d\mu} = \frac{d}{d\mu} H\left(\frac{s-\mu}{\sigma}\right) = -\frac{1}{\sigma} h\left(\frac{s-\mu}{\sigma}\right), \quad \text{so} \quad \frac{\frac{d(1-F(s))}{d\mu}}{f\left(s\right)} = 1.$$

A marginal increase in σ implies that

$$\frac{dF\left(s\right)}{d\sigma} = \frac{d}{d\sigma} H\left(\frac{s-\mu}{\sigma}\right) = -\frac{s-\mu}{\sigma} \frac{1}{\sigma} h\left(\frac{s-\mu}{\sigma}\right), \quad \text{so} \quad \frac{\frac{d(1-F(s))}{d\sigma}}{f\left(s\right)} = \frac{s-\mu}{\sigma}.$$

Mixture Distributions. In this case, note that $f\left(s;\theta\right) = \left(1-\theta\right) \underline{f}\left(s\right) + \theta \overline{f}\left(s\right)$ and $F\left(s;\theta\right) = \left(1-\theta\right) \underline{F}\left(s\right) + \theta \overline{F}\left(s\right)$, so we can express the survival change as

$$\frac{d\left(1-F\left(s;\theta\right)\right)}{d\theta} = \underline{F}\left(s\right) - \overline{F}\left(s\right),$$

implying Equation (13) in the text. Note that

$$p_\theta = \int_{\underline{s}}^{\bar{s}} \frac{\beta u'\left(c_1\left(s\right)\right)}{u'\left(c_0\right)} \frac{dc_1\left(s\right)}{ds} \left(\underline{F}\left(s\right) - \overline{F}\left(s\right)\right) ds,$$

which is invariant to the level of θ, as stated in the text.

Stochastic Dominance. First, we consider a perturbation that satisfies first-order stochastic dominance. In the case of a lottery, $c_1\left(s\right) = s$, and $\frac{dc_1(s)}{ds} = 1$, so Equation (4) simplifies as follows

$$p_\theta = \int_{\underline{s}}^{\bar{s}} \frac{\beta u'\left(c_1\left(s\right)\right)}{u'\left(c_0\right)} \frac{d\left(1 - F\left(s;\theta\right)\right)}{d\theta} ds.$$

It is immediate that any first-order stochastic dominance shift, with $\frac{d(1-F(s;\theta))}{d\theta} \geq 0$, leads to $p_\theta \geq 0$, since utility is increasing in consumption with $u'\left(c_1\left(s\right)\right) > 0$.

Second, we consider a perturbation that satisfies second-order stochastic dominance. Define the cumulative perturbation $H\left(s;\theta\right) = \int_{\underline{s}}^{s} \frac{dF(t;\theta)}{d\theta} dt$, where we have assumed that $H\left(s;\theta\right) \geq 0$ for all s. Notice that, integrating by parts, the change in the expected value of s as a result of the perturbation can be written as

$$\begin{aligned}
\frac{d\mathbb{E}\left[s\right]}{d\theta} &= \int_{\underline{s}}^{\bar{s}} s \frac{df\left(s;\theta\right)}{d\theta} ds = \bar{s} \frac{dF\left(\bar{s};\theta\right)}{d\theta} - \underline{s} \frac{dF\left(\underline{s};\theta\right)}{d\theta} - \int_{\underline{s}}^{\bar{s}} \frac{dF\left(s;\theta\right)}{d\theta} ds \\
&= - \int_{\underline{s}}^{\bar{s}} \frac{dF\left(s;\theta\right)}{d\theta} ds = -H\left(\bar{s};\theta\right),
\end{aligned}$$

so that, by assumption, we have $H\left(\bar{s};\theta\right) = 0$. Now integrating by parts again, we have

$$\begin{aligned}
p_\theta &= \int_{\underline{s}}^{\bar{s}} \frac{\beta u'\left(c_1\left(s\right)\right)}{u'\left(c_0\right)} \frac{d\left(1 - F\left(s;\theta\right)\right)}{d\theta} ds \\
&= - \frac{\beta}{u'\left(c_0\right)} \left\{ u'\left(\bar{s}\right) \underbrace{H\left(\bar{s};\theta\right)}_{=0} - u'\left(\underline{s}\right) \underbrace{H\left(\underline{s};\theta\right)}_{=0} - \int_{\underline{s}}^{\bar{s}} u''\left(s\right) H\left(s;\theta\right) ds \right\} \\
&= \frac{\beta}{u'\left(c_0\right)} \int_{\underline{s}}^{\bar{s}} u''\left(s\right) H\left(s;\theta\right) ds.
\end{aligned}$$

Then, since utility is strictly concave with $u''\left(s\right) < 0$, we obtain $p_\theta \leq 0$.

B Section 3: Supplementary Material

This appendix collects the derivations behind Section 3 and documents the construction of its inputs. The first part derives the probability price for a permanent shift in the temperature distribution, the consumption sensitivity, the state-price normalization, the aggregate-efficiency/risk-sharing identity, and the double auction used to predict hedging volume. The second part describes the data sources, the estimation sample, and the regression estimates, and reports the country-level hedging integrands that underlie the stacked figures in the text.

Probability Prices with Permanent Shifts. Country i's welfare under the perturbed distribution (14) is

$$V^i(\theta) = u\left(c_0^i\right) + \sum_{t\geq 1}\beta^t \int_{\underline{s}}^{\overline{s}} u\left(c_t^i(s)\right) f(s;\theta)\, ds,$$

where consumption at every date is predetermined given s by (15) and does not depend on θ. Under the mixture, $f(s;\theta) = \theta\overline{f}(s) + (1-\theta)\underline{f}(s)$, so $\frac{\partial f(s;\theta)}{\partial\theta} = \overline{f}(s) - \underline{f}(s)$ and $\frac{\partial F(s;\theta)}{\partial\theta} = \overline{F}(s) - \underline{F}(s)$, neither of which depends on θ. Differentiating and integrating by parts horizon by horizon, as in the proof of Proposition 1,

$$\begin{aligned}
\frac{dV^i}{d\theta} &= \sum_{t\geq 1}\beta^t \int_{\underline{s}}^{\overline{s}} u\left(c_t^i(s)\right)\left(\overline{f}(s) - \underline{f}(s)\right) ds \\
&= \sum_{t\geq 1}\beta^t \left[u\left(c_t^i(s)\right)\left(\overline{F}(s) - \underline{F}(s)\right)\Big|_{\underline{s}}^{\overline{s}} - \int_{\underline{s}}^{\overline{s}} u'\left(c_t^i(s)\right)\frac{dc_t^i(s)}{ds}\left(\overline{F}(s) - \underline{F}(s)\right) ds\right] \\
&= \sum_{t\geq 1}\beta^t \int_{\underline{s}}^{\overline{s}} u'\left(c_t^i(s)\right)\frac{dc_t^i(s)}{ds}\left(\underline{F}(s) - \overline{F}(s)\right) ds,
\end{aligned}$$

where the boundary term vanishes because both distributions have the same support, $\underline{F}(\underline{s}) = \overline{F}(\underline{s}) = 0$ and $\underline{F}(\overline{s}) = \overline{F}(\overline{s}) = 1$. Dividing by $u'(c_0^i)$ and multiplying and dividing the integrand by $f(s;\theta)$ gives the horizon-specific valuations in the text,

$$p_{t,\theta}^i = \int_{\underline{s}}^{\overline{s}} \underbrace{\frac{u'\left(c_t^i(s)\right)}{u'\left(c_0^i\right)}}_{\omega_t^i(s)} f(s;\theta) \underbrace{\frac{\underline{F}(s) - \overline{F}(s)}{f(s;\theta)}\frac{dc_t^i(s)}{ds}}_{x_{t,\theta}^i(s)} ds,$$

which is (16), and $p^i_\theta = \sum_{t\geq 1} \beta^t p^i_{t,\theta}$. The state price and the equivalent cash flow each depend on θ through $f(s;\theta)$, but their product does not: for a mixture, the probability price is the same at every point of the path between the two distributions, as noted in Section 2 for the static case. The permanent-shift assumption enters only through the fact that a single realization s determines consumption at every horizon; a transitory shift would replace $c^i_t(s)$ by the consumption path that depends on s only at the affected dates.

Consumption Sensitivity. Taking logarithms in (15), $\log c^i_t(s) = \log c^i_0 + \bar{g}^i_t t + t\left[h^i\left(T^i(s)\right) - h^i\left(\bar{T}^i\right)\right]$, so with $dT^i(s)/ds = 1$,

$$\frac{dc^i_t(s)}{ds} = c^i_t(s)\; t\, g^{i\prime}(s), \qquad g^{i\prime}(s) \equiv \frac{dh^i\left(T^i(s)\right)}{dT} = \phi^i_1 + 2\phi^i_2 T^i(s).$$

Since every estimated response is concave ($\phi^i_2 < 0$), $g^{i\prime}(s) > 0$ if and only if the country's local temperature lies below its growth-maximizing temperature $T^{i*} = -\phi^i_1/(2\phi^i_2)$, that is, if and only if the global state lies below $s^i_0 \equiv \bar{s} + T^{i*} - \bar{T}^i$. The hedging integrand (18), $\tilde{x}^i_\theta(s) = \left(\underline{F}(s) - \overline{F}(s)\right) c^i_t(s)\; t\, g^{i\prime}(s)$, therefore has the sign of $s^i_0 - s$: a country whose s^i_0 lies above the support $[13, 18]^\circ$ C gains from warming at every state, one whose s^i_0 lies below it loses at every state, and a country with s^i_0 inside the support switches from gainer to loser as s crosses s^i_0. Table OA.1 reports s^i_0 for the nine countries in Figure OA.1. The factor $c^i_t(s)\; t$ explains why the vertical scale of the integrand grows with the horizon: the consumption level compounds at the baseline rate and the growth effect accumulates linearly in t.

State Prices and Normalizations. With isoelastic utility, $u'\left(c^i_t(s)\right) = \left(c^i_0 e^{\bar{g}^i_t t}\right)^{-\gamma} e^{-\gamma t \Delta h^i(s)}$, where $\Delta h^i(s) \equiv h^i\left(T^i(s)\right) - h^i\left(\bar{T}^i\right)$. The bond-normalized state price of Section 3.4.2, evaluated under the status quo density, is therefore

$$\omega^i_t(s) = \beta^t \frac{e^{-\gamma t \Delta h^i(s)} \underline{f}(s)}{\int e^{-\gamma t \Delta h^i(\sigma)} \underline{f}(\sigma)\, d\sigma} = \beta^t \frac{e^{-\gamma t \Delta h^i(s)}}{\mathbb{E}_{\underline{F}}\left[e^{-\gamma t \Delta h^i(s)}\right]} \underline{f}(s),$$

from which the consumption level c^i_0 and the baseline growth term $e^{\bar{g}^i_t t}$ have cancelled. It integrates to β^t, the price of a t-period discount bond, and the per-probability reservation valuation is $m^i_t(s) = \omega^i_t(s)/\underline{f}(s)$. At the status quo mean, $T^i(\bar{s}) = \bar{T}^i$ and $\Delta h^i(\bar{s}) = 0$, so $m^i_t(\bar{s}) = \beta^t / \mathbb{E}_{\underline{F}}\left[e^{-\gamma t \Delta h^i(s)}\right]$, the expression used in Section 3.6: a country whose growth responds more strongly to temperature has a more dispersed $\Delta h^i(s)$, a larger expected

marginal utility by convexity of the exponential, and hence a lower reservation valuation for consumption at the undamaged mean state.

The state price in (16), normalized by $u'(c_0^i)$, is $\widehat{\omega}_t^i(s) = u'(c_t^i(s)) f(s;\theta)/u'(c_0^i)$. The two are proportional at each horizon, so the value integrand of the text and the corresponding term of the probability price are related by

$$\int \omega_t^i(s) x_{t,\theta}^i(s) ds = \frac{u'(c_0^i)}{\int u'(c_t^i(\sigma)) \underline{f}(\sigma) d\sigma} \beta^t p_{t,\theta}^i = \frac{e^{\gamma \bar{g}_t^i t}}{\mathbb{E}_{\underline{F}}\left[e^{-\gamma t \Delta h^i(s)}\right]} \beta^t p_{t,\theta}^i,$$

evaluated at $\theta = 0$. The conversion factor is country-specific: it rescales each country's contribution by its trend growth and by its exposure, which is why cross-country comparisons in the text are made in bond-normalized units rather than in units of date-0 consumption. Because ω_t^i carries $\underline{f}$ and $x_{t,\theta}^i$ carries $1/\underline{f}$, the value integrand equals $m_t^i(s)\tilde{x}_\theta^i(s)$ and is bounded on the support even where the density is small.

Aggregate Efficiency and Risk Sharing. For weights $\chi^i \geq 0$ with $\sum_i \chi^i = 1$ and any two cross-sections $\{a^i\}$ and $\{b^i\}$,

$$\sum_i \chi^i a^i b^i = \left(\sum_i \chi^i a^i\right)\left(\sum_i \chi^i b^i\right) + \sum_i \chi^i \left(a^i - \sum_j \chi^j a^j\right)\left(b^i - \sum_j \chi^j b^j\right),$$

which follows by expanding the last sum. Setting $a^i = \omega_t^i(s)$ and $b^i = x_\theta^i(s)$ at each state gives the decomposition of Section 3.5, with $\mathrm{AE}(s)$ the first term and $\mathrm{RS}(s) = \mathbb{C}ov_\chi[\omega_t^i, x_\theta^i]$ the second. The same identity holds with $a^i = m_t^i(s)$ and $b^i = \tilde{x}_\theta^i(s)$, since the density factors out of both sides; this is the form in which the figures are computed. If all countries shared one stochastic discount factor, $\omega_t^i(s)$ would be the same across i at every state, the covariance would vanish identically, and the value integrand would reduce to the population-average position priced at the common state price.

Double Auction. Fix a state s^* and a horizon t, and let $q^i = \chi^i |\tilde{x}_\theta^i(s^*)|$ and $m^i = m_t^i(s^*)$. Buyers are the countries with $x_\theta^i(s^*) < 0$ (the set $\mathcal{B}$) and sellers those with $x_\theta^i(s^*) > 0$ (the set $\mathcal{L}$); countries with a zero position are excluded. The demand schedule at valuation m is $D(m) = \sum_{i\in\mathcal{B}:\, m^i \geq m} q^i$, the total quantity that buyers with reservation valuation at least m would take, and the supply schedule is $S(m) = \sum_{i\in\mathcal{L}:\, m^i \leq m} q^i$. Let $D^{-1}(q)$ and $S^{-1}(q)$ denote the corresponding marginal-valuation curves: $D^{-1}(q)$ is the reservation

valuation of the buyer who supplies the q-th unit of demand when buyers are ordered from the highest valuation down, and $S^{-1}(q)$ that of the seller who supplies the q-th unit of supply when sellers are ordered from the lowest valuation up. Both are step functions, D^{-1} nonincreasing and S^{-1} nondecreasing.

The predicted volume, the clearing valuation, and the surplus are

$$\mathcal{V}^* = \sup\left\{q \leq \mathcal{V} : D^{-1}(q) \geq S^{-1}(q)\right\}, \qquad m^* \in \left[S^{-1}(\mathcal{V}^*), D^{-1}(\mathcal{V}^*)\right],$$
$$\mathcal{S} = \int_0^{\mathcal{V}^*} \left(D^{-1}(q) - S^{-1}(q)\right) dq,$$

where $\mathcal{V} = \min\{\sum_{i\in\mathcal{B}} q^i, \sum_{i\in\mathcal{L}} q^i\}$ is the quantity of the short side of the market, and $\mathcal{V}^* = 0$ when the demand curve lies below the supply curve at $q = 0$. Three properties follow. First, $\mathcal{V}^* \leq \mathcal{V}$: no more can trade than the short side offers, and dispersion in valuations can only lower the matched quantity below that bound. Second, $\mathcal{S} \geq 0$, with equality if and only if $\mathcal{V}^* = 0$, since the integrand is nonnegative on $[0, \mathcal{V}^*]$ by construction. Third, the net position $\sum_i \chi^i \tilde{x}^i_\theta(s^*) = \sum_{i\in\mathcal{L}} q^i - \sum_{i\in\mathcal{B}} q^i$ cannot be matched at any valuation, because it is the excess of one side over the other. The entire short side clears, $\mathcal{V}^* = \mathcal{V}$, when the demand curve lies weakly above the supply curve on $[0, \mathcal{V}]$. Numerically, the code orders buyers by decreasing and sellers by increasing m^i, matches the two heads of the queues unit by unit as long as the buyer's valuation is at least the seller's, and accumulates the matched quantity and the valuation gap; this greedy procedure computes $\mathcal{V}^*$ and $\mathcal{S}$ exactly, since matching in that order is what the integral above describes.

Temperature Distributions. The CMIP6 series come from the IPCC Working Group I Atlas repository (Gutiérrez et al., 2021; Iturbide et al., 2022), which distributes, for every CMIP6 model, scenario, and ensemble member, annual regional averages of near-surface air temperature. We use the global average over land and sea. A model enters the status quo sample if it has a historical run covering 1995–2014, and it enters an SSP sample if it also has a run for that scenario covering 2041–2060; we keep one ensemble member per model and average the annual values over the twenty-year window. This yields 35 status quo realizations, ranging from 13.41 to 15.85°C with mean 14.52°C; 32 realizations under SSP1-2.6, ranging from 14.26 to 16.66°C; and 34 under SSP5-8.5, ranging from 14.95 to 17.29°C. The two future samples differ from the status quo sample in model composition only through data availability.

Each sample $\{s_k\}_{k=1}^n$ is smoothed with a Gaussian kernel of bandwidth b, truncated to

$[a, \bar{a}] = [13, 18]^\circ$ C, and renormalized:

$$f(s) = \frac{\frac{1}{nb}\sum_{k=1}^{n}\varphi\left(\frac{s-s_k}{b}\right)}{\frac{1}{n}\sum_{k=1}^{n}\left[\Phi\left(\frac{\bar{a}-s_k}{b}\right)-\Phi\left(\frac{a-s_k}{b}\right)\right]}, \qquad F(s) = \frac{\frac{1}{n}\sum_{k=1}^{n}\left[\Phi\left(\frac{s-s_k}{b}\right)-\Phi\left(\frac{a-s_k}{b}\right)\right]}{\frac{1}{n}\sum_{k=1}^{n}\left[\Phi\left(\frac{\bar{a}-s_k}{b}\right)-\Phi\left(\frac{a-s_k}{b}\right)\right]},$$

where φ and Φ are the standard normal density and distribution function, so that $F(a) = 0$ and $F(\bar{a}) = 1$ for every scenario and the cdf gap vanishes at both ends of the support. The bandwidth is the two-stage plug-in estimate of Sheather and Jones (1991): $b = 0.333$°C for the status quo, 0.383°C for SSP1-2.6, and 0.311°C for SSP5-8.5. Both functions are stored on a grid of 501 points with spacing 0.01°C; all integrals over s in the text are trapezoidal sums on this grid, and the density-weighted objects use the status quo density $\underline{f}$ on the same grid.

Growth–Climate Panel. We rebuild the country-year panel of Burke, Hsiang, and Miguel (2015) from its original sources and extend it through 2023. Monthly temperature and precipitation on a 0.5° grid come from CRU TS version 4.08 (Harris, Osborn, Jones, and Lister, 2020). For each grid cell and year we average the twelve monthly temperatures and sum the twelve monthly precipitation totals. Cells are assigned to countries with the Natural Earth 1:50m administrative boundaries, and each country's annual temperature and precipitation are population-weighted averages over its cells, with weights from the WorldPop year-2000 population raster (Tatem, 2017) summed from its native 1 km resolution to the CRU grid. The weights are the same in every year. Growth is the annual growth rate of real GDP per capita from the World Bank World Development Indicators (series NY.GDP.PCAP.KD.ZG), the series BHM label growthWDI. The estimation sample keeps the years 1961–2010, drops country-years with a missing growth or climate observation, and keeps countries with at least twenty remaining observations: 7,157 country-years for 168 countries. Country mean temperatures over this sample, $\bar{T}^i$, range from -3.5 to 29.1°C, with cross-country mean 18.9°C; Z^i is $\bar{T}^i$ minus this mean.

Regression Estimates. We estimate (17) by OLS with country and year fixed effects and country-specific quadratic time trends, with standard errors clustered by country; growth is measured as a fraction per year, temperature in degrees Celsius, and precipitation in meters per year. The curvature intercept ϕ_{20} and the interaction ϕ_{21} are both negative,

so every country's response is concave and warmer countries have more concave responses; the level coefficients ϕ_{10} and ϕ_{11} are both positive. The implied growth-maximizing temperature $T^{i*} = -\phi_1^i / (2\phi_2^i)$ has cross-country mean 16.5°C, standard deviation 2.8°C, and median 17.5°C; for the average country ($Z^i = 0$) the optimum is 16.9°C. The hedging analysis uses the point estimates; the cross-country ordering of exposures that drives the results in the text is determined by the sign pattern of the estimates and by each country's distance from its optimum.

Country Levels, Analysis Sample, and Baseline Growth. Population and GDP at purchasing power parity for 2023 come from the World Development Indicators (series SP.POP.TOTL and NY.GDP.MKTP.PP.CD, the latter in current international dollars), and the baseline growth rate $\bar{g}^i$ is the 1990–2019 average of the same growth series used in the regression. The analysis sample consists of the 162 of the 168 regression countries with all three inputs; the six excluded — Cuba, New Caledonia, French Polynesia, Syria, Venezuela, and Yemen — have no 2023 PPP GDP in the WDI. Population shares χ^i are renormalized to sum to one over the 162 countries, and their population-weighted average GDP per capita, 23,474 PPP dollars, is the unit of consumption. Countries with $g^{i\prime}(\bar{s}) < 0$, that is, hotter than their optimum at the status quo mean, hold 61% of the sample's population, which is the source of the negative net position at short horizons in Figure 4.

Country-Level Hedging Integrands. Figure OA.1 plots the hedging integrand $\tilde{x}_\theta^i(s)$ of (18), without population weights, for nine countries at the six horizons $t \in \{1, 5, 10, 20, 50, 100\}$ years, and Table OA.1 lists the inputs that determine each curve. Germany and Russia have s_0^i above 18°C and gain at every state; India, Indonesia, Brazil, and Nigeria have s_0^i below 13°C and lose at every state; the United States, Japan, and China have s_0^i inside the support and switch from gainers to losers at 16.1, 15.7, and 15.4°C, respectively. The scale of each curve is set by $c_t^i(s)\, t$: the United States and Germany, with high c_0^i, dominate at short horizons, while China and India, with high baseline growth, dominate at $t = 50$ and $t = 100$. Multiplying by χ^i produces the contributions stacked in Figure 4.

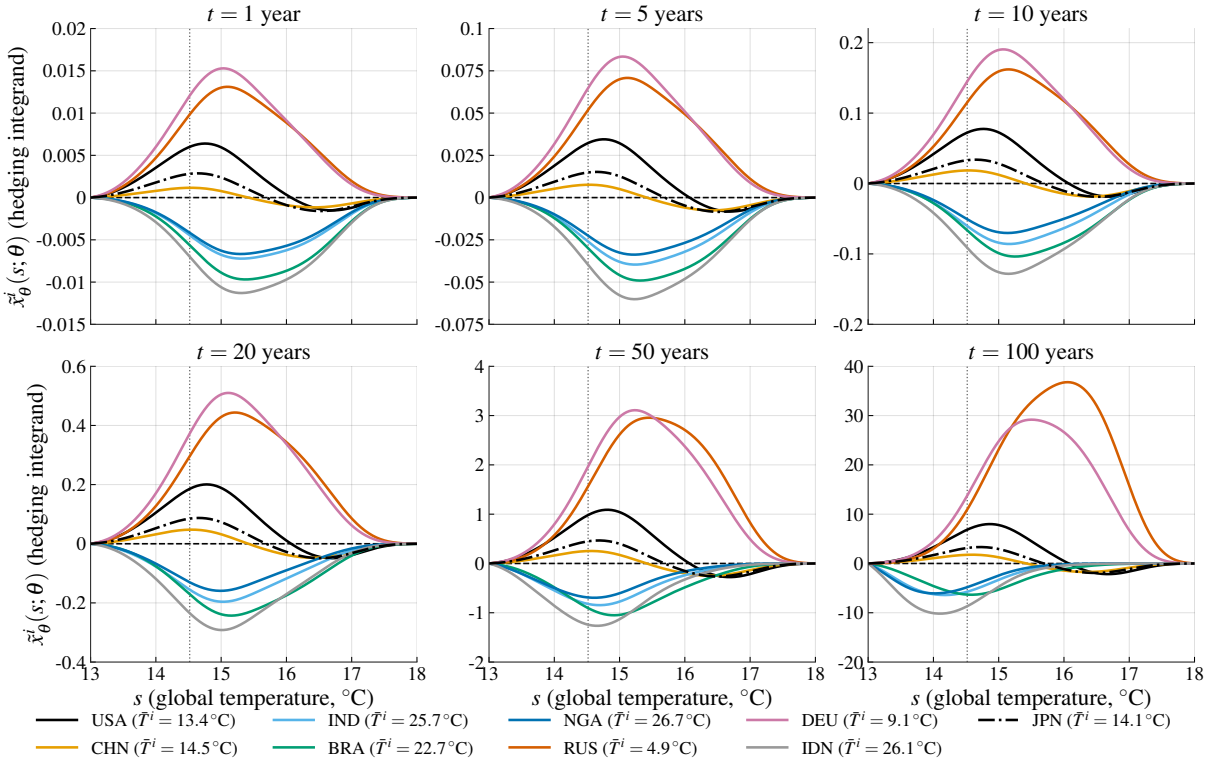


Figure OA.1: Country-Level Hedging Integrands

Note: Hedging integrand $\tilde{x}^i_\theta(s) = \left(\underline{F}(s) - \overline{F}(s)\right) dc^i_t(s)/ds$ for nine countries, under the aggressive scenario (SSP5-8.5), at horizons $t \in \{1, 5, 10, 20, 50, 100\}$ years (one panel per horizon). Positions are per capita and unweighted by population. The dotted vertical line marks the status quo mean of 14.52°C; each legend entry reports the country's mean temperature $\bar{T}^i$.

Country	χ^i	c_0^i	$\bar{g}^i$ (%)	$\bar{T}^i$ (°C)	T^{i*} (°C)	s_0^i (°C)
United States	0.043	3.51	1.52	13.4	15.0	16.1
China	0.180	1.07	8.53	14.5	15.4	15.4
India	0.184	0.44	4.49	25.7	18.7	7.6
Brazil	0.027	0.90	1.02	22.7	18.0	9.8
Nigeria	0.029	0.37	1.81	26.7	19.0	6.8
Russia	0.018	1.92	0.87	4.9	10.8	20.4
Germany	0.011	3.05	1.49	9.1	13.1	18.5
Indonesia	0.036	0.66	3.51	26.1	18.8	7.2
Japan	0.016	2.16	0.96	14.1	15.2	15.7

Table OA.1: Featured Countries

Note: Inputs for the nine countries in Figure OA.1: population share χ^i (2023, renormalized over the 162-country sample), normalized initial consumption c_0^i (2023 PPP GDP per capita relative to the population-weighted world average of $23,474$ PPP dollars), mean 1990–2019 growth of real GDP per capita $\bar{g}^i$, mean 1961–2010 temperature $\bar{T}^i$, growth-maximizing temperature $T^{i*} = -\phi_1^i / (2\phi_2^i)$, and the global state $s_0^i = \bar{s} + T^{i*} - \bar{T}^i$ at which the country's hedging integrand changes sign. Values of s_0^i outside $[13, 18]^\circ$ C mean the integrand has one sign on the whole support.

C Section 4: Supplementary Material

This appendix collects the derivations and proofs for Section 4. We first characterize the Gaussian public signal and the equilibrium of the endowment economy, derive the decomposition of welfare changes into the reallocation and probability channels, prove Propositions 2 and 3, characterize the threshold $\rho^*(\gamma, \delta)$, and show that the Hirshleifer result holds for any information structure. We then turn to the production economy and characterize the equilibrium investment rule, the split of each channel into production-efficiency and risk-sharing components, and the proof of Proposition 4.

Throughout, $\theta \in \{\nu, \sigma\}$ denotes the perturbed signal parameter, and superscripts ν and σ on Ξ^E, Ξ_c, and Ξ_ξ indicate the margin $d\theta = d\nu$ or $d\theta = d\sigma$. Since a reduction in noise raises informativeness, we write $\Xi^{-\sigma} \equiv -\Xi^\sigma$ for the noise margin measured in the direction of more informative signals, so that Propositions 2 through 4 sign Ξ^ν and $\Xi^{-\sigma}$.

C.1 Endowment Economy

Gaussian Public Signal. The conditional densities and distribution functions of the signal are $f(\xi \mid a) = \frac{1}{\sigma}\phi\left(\frac{\xi-\nu}{\sigma}\right)$, $f(\xi \mid b) = \frac{1}{\sigma}\phi\left(\frac{\xi+\nu}{\sigma}\right)$, $F(\xi \mid a) = \Phi\left(\frac{\xi-\nu}{\sigma}\right)$, and $F(\xi \mid b) = \Phi\left(\frac{\xi+\nu}{\sigma}\right)$, where $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal density and distribution function. With equal priors, the marginal density and distribution function of

the signal are

$$f(\xi;\theta) \equiv \sum_s \pi(s) f(\xi \mid s;\theta) = \frac{1}{2\sigma}\left[\phi\left(\frac{\xi-\nu}{\sigma}\right) + \phi\left(\frac{\xi+\nu}{\sigma}\right)\right],$$
$$F(\xi;\theta) \equiv \sum_s \pi(s) F(\xi \mid s;\theta) = \frac{1}{2}\left[\Phi\left(\frac{\xi-\nu}{\sigma}\right) + \Phi\left(\frac{\xi+\nu}{\sigma}\right)\right],$$

so $f(\xi;\theta)$ is the marginal density of the signal. It is even in ξ, with mean zero and variance $\nu^2+\sigma^2$. The log-likelihood ratio is

$$\log\frac{f(\xi \mid a)}{f(\xi \mid b)} = \frac{(\xi+\nu)^2-(\xi-\nu)^2}{2\sigma^2} = \frac{2\nu\xi}{\sigma^2} = 2\rho z, \qquad z \equiv \frac{\xi}{\sigma}, \quad \rho \equiv \frac{\nu}{\sigma},$$

which is the expression in the footnote of Section 4.1. Bayes' rule then gives the posterior probability of state a,

$$p(\xi) = \Pr(s=a \mid \xi) = \frac{\pi(a) f(\xi \mid a)}{\pi(a) f(\xi \mid a) + \pi(b) f(\xi \mid b)} = \left(1+\exp\left\{-\frac{2\nu\xi}{\sigma^2}\right\}\right)^{-1},$$

as stated in Section 4.1. For $\nu > 0$, $p(\xi)$ is strictly increasing in ξ, and $2p(\xi)-1 = \tanh\left(\frac{\nu\xi}{\sigma^2}\right) = \tanh(\rho z)$.

Differentiating the conditional distribution functions with respect to the signal parameters yields $\frac{\partial(1-F(\xi|a))}{\partial\nu} = f(\xi \mid a)$, $\frac{\partial(1-F(\xi|b))}{\partial\nu} = -f(\xi \mid b)$, $\frac{\partial(1-F(\xi|a))}{\partial\sigma} = \frac{\xi-\nu}{\sigma} f(\xi \mid a)$, and $\frac{\partial(1-F(\xi|b))}{\partial\sigma} = \frac{\xi+\nu}{\sigma} f(\xi \mid b)$. Hence the normalized survival changes of the conditional signal distributions, $g_s(\xi;\theta) \equiv \frac{\frac{\partial(1-F(\xi|s;\theta))}{\partial\theta}}{f(\xi|s;\theta)}$, are

$$g_a(\xi;\nu) = 1, \qquad g_b(\xi;\nu) = -1, \qquad g_a(\xi;\sigma) = \frac{\xi-\nu}{\sigma}, \qquad g_b(\xi;\sigma) = \frac{\xi+\nu}{\sigma}, \tag{OA.1}$$

and the survival changes of the marginal signal distribution are

$$\frac{\partial(1-F(\xi;\theta))}{\partial\nu} = \frac{1}{2\sigma}\left[\phi\left(\frac{\xi-\nu}{\sigma}\right) - \phi\left(\frac{\xi+\nu}{\sigma}\right)\right],$$
$$\frac{\partial(1-F(\xi;\theta))}{\partial\sigma} = \frac{h(\xi)}{2\sigma^2}, \qquad h(\xi) \equiv (\xi-\nu)\phi\left(\frac{\xi-\nu}{\sigma}\right) + (\xi+\nu)\phi\left(\frac{\xi+\nu}{\sigma}\right), \tag{OA.2}$$

and both survival changes are odd in ξ. The strength-margin survival change has the sign of ξ, since $|\xi-\nu| < |\xi+\nu|$ for $\xi > 0$: a stronger signal moves probability mass toward the tails. The sign of h is characterized in the proof of Proposition 3.

Blackwell ordering. The footnote in Section 4.1 states that, for $\rho_1 < \rho_2$, the signal with strength-to-noise ratio ρ_1 is a garbling of the signal with ratio ρ_2 in the sense of Blackwell (1953), so that the informativeness of the signal depends on (ν, σ) only through $\rho = \nu/\sigma$ and is increasing in ρ. To see this, note that the map $\xi \mapsto z = \xi/\sigma$ is invertible, so the signal with parameters (ν, σ) is Blackwell equivalent to the standardized signal $z \mid a \sim \mathcal{N}(\rho, 1)$, $z \mid b \sim \mathcal{N}(-\rho, 1)$, which depends on (ν, σ) only through ρ. Let z_2 be drawn from the standardized signal with ratio ρ_2 and define

$$z_1 \equiv \frac{\rho_1}{\rho_2} z_2 + \tau\eta, \qquad \tau^2 \equiv 1 - \left(\frac{\rho_1}{\rho_2}\right)^2 > 0,$$

where $\eta \sim \mathcal{N}(0, 1)$ is independent of (s, z_2). Conditional on $s = a$, z_1 is normal with mean ρ_1 and variance $(\rho_1/\rho_2)^2 + \tau^2 = 1$; conditional on $s = b$, it is normal with mean $-\rho_1$ and variance 1. Hence z_1 has the distribution of the standardized signal with ratio ρ_1, and it is obtained from z_2 through a transformation that does not depend on s, which is the definition of a garbling.

Equilibrium of the Endowment Economy. Given the signal ξ, individual i chooses Arrow–Debreu holdings $x_1^i(s \mid \xi)$ to maximize $V_1^i(\xi) = \sum_s u\left(c_2^i(s, \xi)\right) \Pr(s \mid \xi)$ subject to $\sum_s q_1(s \mid \xi)\, x_1^i(s \mid \xi) = 0$ and $c_2^i(s, \xi) = \chi^i(s)\, n_2(s) + x_1^i(s \mid \xi)$. The first-order conditions are

$$u'\left(c_2^i(s, \xi)\right) \Pr(s \mid \xi) = \eta^i(\xi)\, q_1(s \mid \xi), \qquad s \in \{a, b\}, \tag{OA.3}$$

where $\eta^i(\xi) > 0$ is the multiplier on the date-1 budget constraint. Taking ratios across states,

$$\frac{u'\left(c_2^A(a, \xi)\right)}{u'\left(c_2^A(b, \xi)\right)} = \frac{u'\left(c_2^B(a, \xi)\right)}{u'\left(c_2^B(b, \xi)\right)} = \frac{q_1(a \mid \xi) \Pr(b \mid \xi)}{q_1(b \mid \xi) \Pr(a \mid \xi)}.$$

Market clearing with no aggregate risk, $c_2^A(s, \xi) + c_2^B(s, \xi) = n_2$ for both s, rules out any state dependence of consumption: if $c_2^A(a, \xi) > c_2^A(b, \xi)$, then $c_2^B(a, \xi) < c_2^B(b, \xi)$, so the first ratio is below one while the second is above one, a contradiction, and the reverse inequality is ruled out symmetrically. Hence $c_2^i(s, \xi) = c_2^i(\xi)$, that is, full insurance conditional on the signal, and Equation (OA.3) implies that state prices are proportional to posterior probabilities, $q_1(s \mid \xi) \propto \Pr(s \mid \xi)$. Substituting into the budget constraint, $\sum_s \Pr(s \mid \xi)\left[c_2^i(\xi) - \chi^i(s)\, n_2\right] = 0$, so each individual consumes the conditional expectation of their endowment, $c_2^i(\xi) = \mathbb{E}\left[\chi^i(s) \mid \xi\right] n_2$. With $n_2 = 1$, $\chi^A(a) = \frac{1+\delta}{2}$, and

$\chi^A(b) = \frac{1-\delta}{2}$,

$$c_2^A(\xi) = p(\xi)\frac{1+\delta}{2} + (1-p(\xi))\frac{1-\delta}{2} = \frac{1}{2} + \frac{\delta}{2}(2p(\xi)-1),$$

which is Equation (19), and $c_2^B(\xi) = 1 - c_2^A(\xi)$. Using $2p(\xi) - 1 = \tanh(\nu\xi/\sigma^2)$,

$$c_2^A(\xi) = \frac{1}{2} + \frac{\delta}{2}\tanh\left(\frac{\nu\xi}{\sigma^2}\right) = \frac{1}{2} + \frac{\delta}{2}\tanh(\rho z), \tag{OA.4}$$

which is the representation used in Section 4.2: the equilibrium allocation depends on the signal only through ρ and the standardized signal $z = \xi/\sigma$.

Three properties of Equation (OA.4) are used repeatedly. First, for $\nu > 0$, $c_2^A(\xi)$ is strictly increasing in ξ, with $c_2^A(0) = \frac{1}{2}$ and $c_2^A(-\xi) = c_2^B(\xi)$. Second, its derivatives are

$$\begin{aligned}
\frac{\partial c_2^A(\xi)}{\partial \xi} &= \frac{\delta\nu}{2\sigma^2}\operatorname{sech}^2\left(\frac{\nu\xi}{\sigma^2}\right), \qquad \frac{\partial c_2^A(\xi)}{\partial \nu} = \frac{\delta\xi}{2\sigma^2}\operatorname{sech}^2\left(\frac{\nu\xi}{\sigma^2}\right), \\
\frac{\partial c_2^A(\xi)}{\partial \sigma} &= -\frac{\delta\nu\xi}{\sigma^3}\operatorname{sech}^2\left(\frac{\nu\xi}{\sigma^2}\right),
\end{aligned} \tag{OA.5}$$

so $\frac{\partial c_2^A}{\partial \xi}$ is even and strictly positive, largest at $\xi = 0$ and vanishing as $|\xi| \to \infty$, while $\frac{\partial c_2^A}{\partial \nu}$ and $\frac{\partial c_2^A}{\partial \sigma}$ are odd, with $\operatorname{sgn}\left(\frac{\partial c_2^A}{\partial \nu}\right) = \operatorname{sgn}(\xi)$, $\operatorname{sgn}\left(\frac{\partial c_2^A}{\partial \sigma}\right) = -\operatorname{sgn}(\xi)$, and $\frac{\partial c_2^A}{\partial \sigma} = -2\rho\frac{\partial c_2^A}{\partial \nu}$. Third, since consumption is independent of s given ξ, ex-ante welfare and the welfare numeraire are $V^i = \int u(c_2^i(\xi))\, f(\xi;\theta)\, d\xi$ and $\lambda^i = \int u'(c_2^i(\xi))\, f(\xi;\theta)\, d\xi$; because $c_2^B(\xi) = c_2^A(-\xi)$ and $f(\xi;\theta)$ is even, $V^A = V^B$ and $\lambda^A = \lambda^B \equiv \lambda$.

Reallocation and Probability Channels: Derivation of Equations (20) to (24). Individual i's expected utility is $V^i = \int \sum_s u(c_2^i(s,\xi;\theta))\, f(\xi \mid s;\theta)\, \pi(s)\, d\xi$, where $c_2^i(s,\xi;\theta)$ is the equilibrium mapping from histories to consumption, which depends on θ, and we allow for state-dependent consumption so that the derivation also covers the production economy. Differentiating,

$$\begin{aligned}
\frac{dV^i}{d\theta} = &\int \sum_s u'\left(c_2^i(s,\xi;\theta)\right)\frac{\partial c_2^i(s,\xi;\theta)}{\partial \theta} f(\xi \mid s;\theta)\, \pi(s)\, d\xi \\
&+ \sum_s \pi(s) \int u\left(c_2^i(s,\xi;\theta)\right)\frac{\partial f(\xi \mid s;\theta)}{\partial \theta} d\xi.
\end{aligned}$$

The second term is the probability pricing term of Proposition 1, applied state by state to the conditional distribution $F(\xi \mid s;\theta)$. Integrating by parts in ξ, and using that $F(\xi \mid s;\theta)$ converges to 0 and 1 as $\xi \to -\infty$ and $\xi \to \infty$ for every θ, so that the boundary terms vanish as in the proof of Proposition 1,

$$\begin{aligned}\int u\left(c_2^i(s,\xi;\theta)\right)\frac{\partial f(\xi \mid s;\theta)}{\partial\theta}d\xi &= \int u'\left(c_2^i(s,\xi;\theta)\right)\frac{\partial c_2^i(s,\xi;\theta)}{\partial\xi}\frac{\partial\left(1-F(\xi \mid s;\theta)\right)}{\partial\theta}d\xi \\ &= \int u'\left(c_2^i(s,\xi;\theta)\right)g_s(\xi;\theta)\frac{\partial c_2^i(s,\xi;\theta)}{\partial\xi}f(\xi \mid s;\theta)\,d\xi.\end{aligned}$$

Dividing by $\lambda^i = \int \sum_s u'\left(c_2^i(s,\xi;\theta)\right) f(\xi \mid s;\theta)\,\pi(s)\,d\xi$ and using the definition of the state price, $\omega_2^i(s,\xi) = \frac{\pi(s)f(\xi|s;\theta)u'\left(c_2^i(s,\xi;\theta)\right)}{\lambda^i}$,

$$\frac{\frac{dV^i}{d\theta}}{\lambda^i} = \int \sum_s \omega_2^i(s,\xi)\left(\frac{\partial c_2^i(s,\xi;\theta)}{\partial\theta} + g_s(\xi;\theta)\frac{\partial c_2^i(s,\xi;\theta)}{\partial\xi}\right)d\xi, \tag{OA.6}$$

which is the counterpart of Equation (7) and reduces to Equation (20) under full insurance, $c_2^i(s,\xi) = c_2^i(\xi)$. Normalizing by λ^i expresses gains in units of uncontingent date-2 consumption: a uniform increase dc in c_2^i at every history raises V^i by $\lambda^i dc$. Summing Equation (OA.6) over the two individuals gives $\Xi^E = \Xi_c + \Xi_\xi$, with the reallocation and probability channels defined in Equations (21) and (22).

Under full insurance, $u'\left(c_2^i(\xi)\right)$ does not depend on s, so summing the state prices over states gives

$$\sum_s \omega_2^i(s,\xi) = \frac{u'\left(c_2^i(\xi)\right)}{\lambda^i}\sum_s \pi(s)\,f(\xi \mid s;\theta) = \frac{f(\xi;\theta)\,u'\left(c_2^i(\xi)\right)}{\lambda^i} \equiv \omega_2^i(\xi).$$

Since the aggregate endowment is fixed, $\frac{\partial c_2^B}{\partial\theta} = -\frac{\partial c_2^A}{\partial\theta}$, and the reallocation channel becomes $\Xi_c = \int\left[\omega_2^A(\xi) - \omega_2^B(\xi)\right]\frac{\partial c_2^A(\xi;\theta)}{\partial\theta}d\xi$, which is Equation (23). For the probability channel, the normalized survival change enters conditional on s; multiplying by the state price and summing over states,

$$\begin{aligned}\sum_s \omega_2^i(s,\xi)\,g_s(\xi;\theta) &= \frac{u'\left(c_2^i(\xi)\right)}{\lambda^i}\sum_s \pi(s)\frac{\partial\left(1-F(\xi \mid s;\theta)\right)}{\partial\theta} = \frac{u'\left(c_2^i(\xi)\right)}{\lambda^i}\frac{\partial\left(1-F(\xi;\theta)\right)}{\partial\theta} \\ &= \omega_2^i(\xi)\frac{\frac{\partial(1-F(\xi;\theta))}{\partial\theta}}{f(\xi;\theta)},\end{aligned}$$

so the probability channel depends on the signal distribution only through the

normalized survival change of its marginal distribution. Using $\frac{\partial c_2^B}{\partial \xi} = -\frac{\partial c_2^A}{\partial \xi}$, $\Xi_\xi = \int \left[\omega_2^A(\xi) - \omega_2^B(\xi)\right] \frac{\partial c_2^A(\xi;\theta)}{\partial \xi} \frac{\partial (1-F(\xi;\theta))}{\partial \theta} \frac{}{f(\xi;\theta)} d\xi$, which is Equation (24). Finally, define the state-price gap

$$\Delta\omega(\xi) \equiv \omega_2^A(\xi) - \omega_2^B(\xi) = \frac{f(\xi;\theta)}{\lambda} \Delta u'(\xi), \qquad \Delta u'(\xi) \equiv u'\left(c_2^A(\xi)\right) - u'\left(c_2^B(\xi)\right). \tag{OA.7}$$

For $\xi > 0$, $c_2^A(\xi) > c_2^B(\xi)$ and u' is strictly decreasing, so $\Delta u'(\xi) < 0$; since $c_2^A(-\xi) = c_2^B(\xi)$, $\Delta u'$ is odd. Hence $\operatorname{sgn}(\Delta\omega(\xi)) = \operatorname{sgn}(\Delta u'(\xi)) = -\operatorname{sgn}(\xi)$ for $\xi \neq 0$, as claimed in Section 4.2.

Proof of Proposition 2 (Sign of the Reallocation Channel).

Proof. By Equation (23), $\Xi_c = \int \Delta\omega(\xi) \frac{\partial c_2^A(\xi;\theta)}{\partial \theta} d\xi$. By Equation (OA.7), $\Delta\omega(\xi)$ has the sign of $-\xi$, while by Equation (OA.5) both $\frac{\partial c_2^A}{\partial \nu}$ and $\frac{\partial c_2^A}{\partial(-\sigma)} = -\frac{\partial c_2^A}{\partial \sigma}$ have the sign of ξ. The integrand is therefore strictly negative for every $\xi \neq 0$ on both margins. Since $f(\xi;\theta) > 0$ and $\xi = 0$ has measure zero, integration yields $\Xi_c^\nu < 0$ and $\Xi_c^{-\sigma} < 0$ whenever $\nu > 0$. Moreover, since $\frac{\partial c_2^A}{\partial \sigma} = -2\rho \frac{\partial c_2^A}{\partial \nu}$ pointwise,

$$\Xi_c^\sigma = -2\rho\, \Xi_c^\nu, \tag{OA.8}$$

an identity used below. □

Proof of Proposition 3 (Sign of the Probability Channel) and Characterization of ρ^*. *Strength and noise margins.* We first derive how welfare and the two channels depend on (ν, σ), which links the noise margin to the strength margin. With $z = \xi/\sigma$, the standardized signal has density $\frac{1}{2}[\phi(z-\rho) + \phi(z+\rho)]$ and, by Equation (OA.4), the allocation is $c_2^A = \frac{1}{2} + \frac{\delta}{2}\tanh(\rho z)$. Hence $V^i = \int u(c_2^i(z)) \frac{1}{2}[\phi(z-\rho) + \phi(z+\rho)]\, dz$ and $\lambda = \int u'\left(c_2^A(z)\right) \frac{1}{2}[\phi(z-\rho) + \phi(z+\rho)]\, dz$ depend on (ν, σ) only through ρ. Write $V(\rho) \equiv V^i$. By the chain rule, $\frac{\partial V}{\partial \nu} = \frac{V'(\rho)}{\sigma}$ and $\frac{\partial V}{\partial \sigma} = -\frac{\nu}{\sigma^2} V'(\rho) = -\rho \frac{V'(\rho)}{\sigma}$. Since $\Xi^{E,\theta} = \frac{2}{\lambda} \frac{\partial V}{\partial \theta}$, the strength and noise margins satisfy

$$\Xi^{E,\sigma} = -\rho\, \Xi^{E,\nu}, \qquad \Xi_c^\sigma = -2\rho\, \Xi_c^\nu, \qquad \Xi_\xi^\sigma = \rho\left(\Xi_c^\nu - \Xi_\xi^\nu\right), \tag{OA.9}$$

where the first identity is the chain rule, the second is Equation (OA.8), and the third follows from $\Xi_\xi = \Xi^E - \Xi_c$ on each margin: $\Xi_\xi^\sigma = -\rho\, \Xi^{E,\nu} + 2\rho\, \Xi_c^\nu = -\rho\left(\Xi_c^\nu + \Xi_\xi^\nu\right) + 2\rho\, \Xi_c^\nu =$

$\rho\left(\Xi_c^\nu - \Xi_\xi^\nu\right)$. Moreover, $\Xi^{E,\nu} = \frac{2V'(\rho)}{\sigma\lambda(\rho)}$ equals $\frac{1}{\sigma}$ times a function of (ρ, γ, δ), the same change of variables applied to Equation (23) shows that $\sigma\,\Xi_c^\nu$ depends only on (ρ, γ, δ), and by Equation (OA.9) the same scaling holds for every other margin and channel. This justifies the normalization $\sigma = 1$ in Section 4.2: with $\sigma = 1$, every welfare object is a function of (ρ, γ, δ) alone.

Proof. We prove Proposition 3. Substituting Equation (OA.7) into Equation (24), the marginal density cancels and

$$\Xi_\xi = \frac{1}{\lambda}\int \Delta u'(\xi)\,\frac{\partial c_2^A(\xi)}{\partial \xi}\,\frac{\partial\left(1 - F(\xi;\theta)\right)}{\partial \theta}d\xi, \tag{OA.10}$$

where $\Delta u'$ is odd with $\operatorname{sgn}\left(\Delta u'(\xi)\right) = -\operatorname{sgn}(\xi)$, and $\frac{\partial c_2^A}{\partial \xi} > 0$ is even by Equation (OA.5). The sign of the integrand at each ξ is therefore determined by the survival change of the marginal signal distribution in Equation (OA.2). □

Strength margin. The survival change $\frac{\partial(1-F(\xi;\theta))}{\partial \nu}$ is odd and has the sign of ξ. The integrand in Equation (OA.10) is the product of $\Delta u'$, which has the sign of $-\xi$, a strictly positive even factor, and a factor with the sign of ξ, so it is strictly negative for every $\xi \neq 0$. Hence $\Xi_\xi^\nu < 0$ whenever $\nu > 0$, which proves part (i). (At $\nu = 0$, $\Delta u'$ vanishes identically and $\Xi_\xi^\nu = 0$.)

Noise margin. For $\theta = -\sigma$, the survival change is $-\frac{h(\xi)}{2\sigma^2}$ by Equation (OA.2). Since $\Delta u'$ and h are odd while $\frac{\partial c_2^A}{\partial \xi}$ is even, the integrand is even, and

$$\Xi_\xi^{-\sigma} = \frac{1}{\sigma^2\lambda}\int_0^\infty w(\xi)\,h(\xi)\,d\xi, \qquad w(\xi) \equiv \left[u'\left(c_2^B(\xi)\right) - u'\left(c_2^A(\xi)\right)\right]\frac{\partial c_2^A(\xi)}{\partial \xi} > 0 \quad \text{for } \xi > 0, \tag{OA.11}$$

so the sign of $\Xi_\xi^{-\sigma}$ is the sign of a w-weighted average of h over $\xi > 0$. For $\xi \geq \nu$, both terms of h are nonnegative and $h(\xi) > 0$. For $0 < \xi < \nu$, $h(\xi) > 0$ if and only if $(\xi + \nu)\,\phi\left(\frac{\xi+\nu}{\sigma}\right) > (\nu - \xi)\,\phi\left(\frac{\nu-\xi}{\sigma}\right)$, which after taking logarithms reads $r(\xi) > 0$, where

$$r(\xi) \equiv \log\left(\frac{\nu+\xi}{\nu-\xi}\right) - \frac{2\nu\xi}{\sigma^2}, \qquad r(0) = 0, \qquad r'(\xi) = \frac{2\nu}{\nu^2 - \xi^2} - \frac{2\nu}{\sigma^2}.$$

Unimodal regime, $\nu \leq \sigma$. Then $\nu^2 - \xi^2 < \sigma^2$ for every $\xi \in (0, \nu)$, so $r' > 0$ and $r > 0$ on $(0, \nu)$: $h(\xi) > 0$ for every $\xi > 0$, the integrand in Equation (OA.11) is positive pointwise, and $\Xi_\xi^{-\sigma} > 0$ whenever $\nu > 0$. Because this argument uses only the concavity of u and

the symmetry and monotonicity of the equilibrium allocation, the condition $\nu \leq \sigma$ yields a positive probability channel for any strictly concave utility function, isoelastic or not. *Bimodal regime*, $\nu > \sigma$. Then $r'(\xi) < 0$ for $\xi < \sqrt{\nu^2 - \sigma^2}$ and $r'(\xi) > 0$ beyond, with $r(\xi) \to \infty$ as $\xi \uparrow \nu$. Hence r is strictly negative on $\left(0, \sqrt{\nu^2 - \sigma^2}\right]$ and crosses zero exactly once, at $\xi^* \in \left(\sqrt{\nu^2 - \sigma^2}, \nu\right)$: $h < 0$ on $(0, \xi^*)$ and $h > 0$ on (ξ^*, ∞). The central region, where noise reduction hollows out the marginal distribution, and the tail region, where it concentrates mass around $\pm\nu$, contribute to $\Xi_\xi^{-\sigma}$ with opposite signs, and the net sign depends on the weight w, that is, on γ and δ. Two limits pin down the sign at both ends of the bimodal regime. First, at $\nu = \sigma$ the integrand is strictly positive almost everywhere, so $\Xi_\xi^{-\sigma} > 0$; since $\Xi_\xi^{-\sigma}$ is continuous in (ν, σ), the probability channel remains strictly positive for ρ in a neighborhood above one. Second, for ρ sufficiently large, $\Xi_\xi^{-\sigma} < 0$. To see this, normalize $\sigma = 1$ (without loss, by the scaling argument above), so that $\nu = \rho$. The weight $w(\xi)$ is proportional to $\rho \tanh(\rho\xi) \operatorname{sech}^2(\rho\xi)$ up to a bounded factor, so it concentrates on signals of order $1/\rho$, where $h(\xi) = h'(0)\xi + O(\xi^2)$ with $h'(0) = 2\phi(\rho)(1 - \rho^2) < 0$. The contribution of $(0, \xi^*)$ is therefore negative and at least of order $\rho\,\phi(\rho)$ in absolute value. The positive contribution begins at $\xi^* > \sqrt{\rho^2 - 1}$, where w is bounded by a multiple of $e^{-2\rho\xi^*}$ and h grows at most linearly, so it is bounded by a multiple of $e^{-2\rho\sqrt{\rho^2-1}}$, which is exponentially smaller than $\phi(\rho)$ for large ρ. The negative central contribution dominates and $\Xi_\xi^{-\sigma} < 0$.

Characterization of the threshold. By continuity, $\Xi_\xi^{-\sigma}$ changes sign at least once as ρ increases beyond one. The third identity in Equation (OA.9) characterizes the crossing in terms of the strength margin:

$$\Xi_\xi^{-\sigma} = \rho\left(\Xi_\xi^\nu - \Xi_c^\nu\right) = \rho\left(|\Xi_c^\nu| - \left|\Xi_\xi^\nu\right|\right), \tag{OA.12}$$

where the second equality uses $\Xi_c^\nu < 0$ and $\Xi_\xi^\nu < 0$, from Proposition 2 and part (i). The probability channel on the noise margin is thus positive exactly when, on the strength margin, the reallocation channel dominates the probability channel in absolute value, and it vanishes when the two strength-margin channels coincide. Since every channel scales with $\frac{1}{\sigma}$, the ratio Ξ_ξ^ν / Ξ_c^ν depends only on (ρ, γ, δ), so the threshold

$$\rho^*(\gamma, \delta) \equiv \min\left\{\rho > 1 : \Xi_c^\nu(\rho; \gamma, \delta) = \Xi_\xi^\nu(\rho; \gamma, \delta)\right\}$$

is well defined and depends only on preferences and endowment risk, which is the statement

in part (ii) of Proposition 3. Numerically, the crossing is unique: $\Xi_{\xi}^{-\sigma} > 0$ for $1 < \rho < \rho^*$ and $\Xi_{\xi}^{-\sigma} < 0$ for $\rho > \rho^*$, with $\rho^*(2, 0.4) \approx 1.22$ at the baseline calibration and $\rho^* \in [1.20, 1.56]$ for $\gamma \in \{0.5, 1, 2, 5, 10\}$ and $\delta \in \{0.1, 0.4, 0.8\}$. At the baseline, the noise-margin switch in Figure 9 occurs at $\sigma = \nu/\rho^* \approx 0.62$ for $\nu = 0.75$.

The Hirshleifer Effect: Total Welfare and Signal Discrimination. Two features of the equilibrium allocation drive the Hirshleifer effect. First, insured consumption is a *martingale* in the signal: by the law of iterated expectations, $\mathbb{E}[c_2^i(\xi)] = \frac{n_2}{2}$ for every information structure. Second, ex-ante welfare is the expectation of a concave function of that martingale, $V^i = \mathbb{E}[u(c_2^i(\xi))]$. An uninformative signal leaves both individuals fully insured at $c_2^i = \frac{n_2}{2}$ (the first best), while a fully revealing signal eliminates all insurance and leaves each individual consuming their endowment (autarky). Hirshleifer's conclusion does not depend on any parametric structure. For a public signal T, let $V(T)$ denote the ex-ante welfare of either individual. If T and T' are two public signals about s, with T more informative than T' in the sense of Blackwell (1953), then $V(T) \leq V(T')$, strictly whenever the two signals induce different distributions of insured consumption. To see this, let $c_T = \mathbb{E}\left[\chi^A(s) n_2(s) \mid T\right]$ and $c_{T'} = \mathbb{E}\left[\chi^A(s) n_2(s) \mid T'\right]$ denote individual A's insured consumption under the two signals. Since T' is a garbling of T, the law of iterated expectations implies $\mathbb{E}[c_T \mid T'] = c_{T'}$, so c_T is a mean-preserving spread of $c_{T'}$. Because $V = \mathbb{E}[u(c)]$ with u concave, a mean-preserving spread lowers welfare (Rothschild and Stiglitz, 1970), strictly when the spread is nondegenerate; the same argument applies to individual B. Information that arrives before insurance can be purchased disperses the consumption that the two individuals could otherwise have pooled, and risk-averse individuals dislike the spread. Combined with the garbling argument above, this shows that in the Gaussian family $V(\rho)$ is nonincreasing in ρ.

Total welfare effect on both margins. Section 4.2 states that in the endowment economy the reallocation channel always dominates, so that the total effect Ξ^E in Figure 9 is positive at every ν and σ. In terms of the informativeness margins, $\Xi^{E,\nu} < 0$ and $\Xi^{E,-\sigma} < 0$ for every $\nu > 0$ and $\sigma > 0$. On the strength margin, $\Xi^{E,\nu} = \Xi_c^{\nu} + \Xi_{\xi}^{\nu} < 0$ by Proposition 2 and part (i) of Proposition 3. On the noise margin, the first identity in Equation (OA.9) gives $\Xi^{E,-\sigma} = -\Xi^{E,\sigma} = \rho\,\Xi^{E,\nu} < 0$. In particular, whenever the probability channel is positive on the noise margin, $\Xi_{\xi}^{-\sigma} > 0$, the reallocation channel dominates it, $|\Xi_c^{-\sigma}| > \Xi_{\xi}^{-\sigma}$, since $\Xi^{E,-\sigma} = \Xi_c^{-\sigma} + \Xi_{\xi}^{-\sigma} < 0$ with $\Xi_c^{-\sigma} < 0$.

Signal discrimination. Section 4.2 calibrates ρ to the area under the receiver operating

characteristic curve of the signal, AUC, defined as the probability that a signal drawn in state a is larger than an independent signal drawn in state b. With $\xi_a \sim \mathcal{N}(\nu, \sigma^2)$ and $\xi_b \sim \mathcal{N}(-\nu, \sigma^2)$ independent, $\xi_a - \xi_b \sim \mathcal{N}(2\nu, 2\sigma^2)$, so

$$\text{AUC} = \Pr(\xi_a > \xi_b) = \Phi\left(\frac{2\nu}{\sqrt{2}\sigma}\right) = \Phi\left(\sqrt{2}\rho\right),$$

which is strictly increasing in ρ, from $\frac{1}{2}$ at $\rho = 0$ toward 1 as $\rho \to \infty$. Hence ρ maps one-to-one into AUC, and $\rho = 0.75$ corresponds to $\text{AUC} = \Phi(1.06) \approx 0.86$.

C.2 Production Economy

Equilibrium. After observing ξ, individual i chooses investment k^i and Arrow–Debreu holdings $x_1^i(s \mid \xi)$ to maximize $\sum_s u(c_2^i(s,\xi)) \Pr(s \mid \xi)$ subject to $\sum_s q_1(s \mid \xi) x_1^i(s \mid \xi) = 0$ and $c_2^i(s,\xi) = \chi^i(s) n_2 + R(s) k^i - \frac{\kappa}{2}(k^i)^2 + x_1^i(s \mid \xi)$. Because markets against s are complete at date 1, the investment decision separates from the consumption decision: k^i maximizes the market value of production, $\sum_s q_1(s \mid \xi)\left[R(s) k^i - \frac{\kappa}{2}(k^i)^2\right]$, which is strictly concave in k^i because $\kappa > 0$, while the first-order condition for consumption is again Equation (OA.3). Combining the first-order condition for investment, $\sum_s q_1(s \mid \xi)(R(s) - \kappa k^i) = 0$, with Equation (OA.3) and $\Pr(s \mid \xi) = \frac{\pi(s) f(\xi \mid s)}{f(\xi;\theta)}$ yields

$$\sum_s u'\left(c_2^i(s,\xi)\right)\left(R(s) - \kappa k^i\right) f(\xi \mid s)\, \pi(s) = 0,$$

which is Equation (25). Since both individuals face the same state prices, they choose the same investment, $k^A = k^B = k(\xi) = \frac{\sum_s q_1(s|\xi) R(s)}{\kappa \sum_s q_1(s|\xi)}$, so $\kappa k(\xi)$ is a weighted average of $R(a)$ and $R(b)$ with strictly positive weights, and aggregate consumption is $c_2(s,\xi) = \sum_i c_2^i(s,\xi) = n_2 + 2\left[R(s) k(\xi) - \frac{\kappa}{2} k(\xi)^2\right]$. Taking ratios of Equation (OA.3) across states, as in the endowment economy, $\frac{u'\left(c_2^A(a,\xi)\right)}{u'\left(c_2^A(b,\xi)\right)} = \frac{u'\left(c_2^B(a,\xi)\right)}{u'\left(c_2^B(b,\xi)\right)}$; with isoelastic utility this means $\frac{c_2^A(a,\xi)}{c_2^A(b,\xi)} = \frac{c_2^B(a,\xi)}{c_2^B(b,\xi)}$, so both individuals consume constant shares of aggregate consumption across states, $c_2^i(s,\xi) = \alpha^i(\xi) c_2(s,\xi)$ with $\sum_i \alpha^i(\xi) = 1$, which is the full-insurance property stated in Section 4.3. Because $u'(\alpha c) = \alpha^{-\gamma} u'(c)$, the shares cancel from Equation (OA.3), and state prices satisfy $q_1(s \mid \xi) \propto u'(c_2(s,\xi)) \Pr(s \mid \xi)$. The welfare numeraire and state prices are as in the main text, $\lambda^i = \int \sum_s u'(c_2^i(s,\xi)) f(\xi \mid s;\theta)\, \pi(s)\, d\xi$ and $\omega_2^i(s,\xi) = \frac{\pi(s) f(\xi|s;\theta) u'\left(c_2^i(s,\xi)\right)}{\lambda^i}$, and Equation (OA.6) applies with state-dependent consumption.

Investment increases in the signal. For $\mu > 0$, the equilibrium investment satisfies $R(b) < \kappa k(\xi) < R(a)$ and is strictly increasing in the signal, $k'(\xi) > 0$; for $\mu = 0$, $\kappa k(\xi) = 1$ and $k'(\xi) = 0$. To see this, use $q_1(s \mid \xi) \propto u'(c_2(s,\xi)) \Pr(s \mid \xi)$ to write the investment first-order condition as $\Omega_k(k(\xi), p(\xi)) = 0$, where

$$\Omega(k,p) \equiv p\, u(c_2(a;k)) + (1-p)\, u(c_2(b;k)), \qquad c_2(s;k) \equiv n_2 + 2\left[R(s)\, k - \frac{\kappa}{2}k^2\right],$$

is the expected utility of aggregate consumption under the posterior p. Since $\Omega_{kk} = \sum_s \Pr(s \mid \xi)\left[4(R(s) - \kappa k)^2 u''(c_2(s;k)) - 2\kappa u'(c_2(s;k))\right] < 0$, $k(\xi)$ is the unique maximizer of $\Omega(\cdot, p(\xi))$, and the first-order condition $\Omega_k = 2\sum_s \Pr(s \mid \xi)\, u'(c_2(s;k))(R(s) - \kappa k) = 0$ expresses $\kappa k(\xi)$ as a weighted average of $R(a) = e^{\mu}$ and $R(b) = e^{-\mu}$ with strictly positive weights, so $R(b) < \kappa k(\xi) < R(a)$ when $\mu > 0$. The cross-partial $\Omega_{kp} = 2\left[u'(c_2(a;k))(R(a) - \kappa k) - u'(c_2(b;k))(R(b) - \kappa k)\right]$ is strictly positive at the optimum, since $R(a) - \kappa k > 0 > R(b) - \kappa k$. By the implicit function theorem, $k'(\xi) = -\frac{\Omega_{kp}}{\Omega_{kk}} p'(\xi) > 0$, because $p'(\xi) > 0$ for $\nu > 0$. When $\mu = 0$, $R(a) = R(b) = 1$, so $\kappa k(\xi) = 1$ for every ξ.

Production-efficiency and risk-sharing components. Summing Equation (OA.6) over individuals, $\Xi^E = \sum_i \int \sum_s \omega_2^i(s,\xi)\left[\frac{\partial c_2^i(s,\xi;\theta)}{\partial \theta} + g_s(\xi;\theta)\frac{\partial c_2^i(s,\xi;\theta)}{\partial \xi}\right] d\xi$. Define the aggregate state price $\omega_2(s,\xi) \equiv \sum_i \omega_2^i(s,\xi)$. For any pair of individual responses $y^i(s,\xi)$ with aggregate $y(s,\xi) \equiv \sum_i y^i(s,\xi)$, the identity

$$\sum_i \omega_2^i(s,\xi)\, y^i(s,\xi) = \frac{\omega_2(s,\xi)}{2} y(s,\xi) + \sum_i \left(\omega_2^i(s,\xi) - \frac{\omega_2(s,\xi)}{2}\right)\left(y^i(s,\xi) - \frac{y(s,\xi)}{2}\right) \tag{OA.13}$$

is the two-individual case of the decomposition in Dávila and Schaab (2025): the first term values the aggregate response at the average state price, and the second term is twice the cross-sectional covariance between state prices and individual responses, which is nonzero only when the response is reallocated across individuals with different state prices. Applying Equation (OA.13) at every history to $y^i = \frac{\partial c_2^i}{\partial \theta}$ and to $y^i = g_s \frac{\partial c_2^i}{\partial \xi}$, whose aggregates are $\frac{\partial c_2}{\partial \theta}$ and $g_s \frac{\partial c_2}{\partial \xi}$, and integrating, the production-efficiency components of the two channels are

$$\Xi_c^{PE} = \int \sum_s \frac{\omega_2(s,\xi)}{2} \frac{\partial c_2(s,\xi;\theta)}{\partial \theta} d\xi \qquad \text{and} \qquad \Xi_\xi^{PE} = \int \sum_s \frac{\omega_2(s,\xi)}{2} g_s(\xi;\theta) \frac{\partial c_2(s,\xi)}{\partial \xi} d\xi, \tag{OA.14}$$

and the risk-sharing components, $\Xi_c^{RS} = \Xi_c - \Xi_c^{PE}$ and $\Xi_\xi^{RS} = \Xi_\xi - \Xi_\xi^{PE}$, collect the covariance terms. This is the split $\Xi^E = \Xi_c^{PE} + \Xi_c^{RS} + \Xi_\xi^{PE} + \Xi_\xi^{RS}$ stated in Section 4.3.

Because individuals consume at a single date, the intertemporal-sharing component of Dávila and Schaab (2025) is absent. In the endowment economy, aggregate consumption does not respond to θ or to ξ, so $\Xi_c^{PE} = \Xi_\xi^{PE} = 0$ and $\Xi^E = \Xi^{RS}$: the two channels of Section 4.2 are entirely risk-sharing components.

Proof of Proposition 4 (Production Efficiency).

Proof. For part (i), since the endowment does not depend on θ, differentiating aggregate consumption at a given signal yields $\frac{\partial c_2(s,\xi;\theta)}{\partial \theta} = 2\left(R\left(s\right) - \kappa k\left(\xi\right)\right)\frac{\partial k(\xi;\theta)}{\partial \theta}$. Substituting into Equation (OA.14) and using the definition of ω_2,

$$\Xi_c^{PE} = \int \frac{\partial k\left(\xi;\theta\right)}{\partial \theta} \sum_i \frac{1}{\lambda^i} \left[\sum_s u'\left(c_2^i\left(s,\xi\right)\right)\left(R\left(s\right) - \kappa k\left(\xi\right)\right) f\left(\xi \mid s;\theta\right)\pi\left(s\right)\right] d\xi = 0,$$

because the bracket is individual i's first-order condition for investment, Equation (25). The investment response $\frac{\partial k}{\partial \theta}$ therefore carries no first-order welfare value at given probabilities, so $\Xi^{PE} = \Xi_\xi^{PE}$ and the entire reallocation channel contributes to risk sharing, $\Xi_c^{RS} = \Xi_c$. For part (ii), differentiating aggregate consumption with respect to the signal at a given θ yields $\frac{\partial c_2(s,\xi)}{\partial \xi} = 2\left(R\left(s\right) - \kappa k\left(\xi\right)\right)k'\left(\xi\right)$, so

$$\Xi_\xi^{PE} = \int k'\left(\xi\right)\sum_s \omega_2\left(s,\xi\right)\left(R\left(s\right) - \kappa k\left(\xi\right)\right) g_s\left(\xi;\theta\right) d\xi.$$

Summing the two first-order conditions, Equation (25), each divided by λ^i, gives $\sum_s \omega_2\left(s,\xi\right)\left(R\left(s\right) - \kappa k\left(\xi\right)\right) = 0$ for every ξ, so

$$A\left(\xi\right) \equiv \omega_2\left(a,\xi\right)\left(R\left(a\right) - \kappa k\left(\xi\right)\right) = -\omega_2\left(b,\xi\right)\left(R\left(b\right) - \kappa k\left(\xi\right)\right),$$

and $A\left(\xi\right) > 0$ for $\mu > 0$, since $R\left(b\right) < \kappa k\left(\xi\right) < R\left(a\right)$ as shown above. Hence $\Xi_\xi^{PE} = \int k'\left(\xi\right) A\left(\xi\right)\left[g_a\left(\xi;\theta\right) - g_b\left(\xi;\theta\right)\right] d\xi$. By Equation (OA.1), $g_a - g_b = 2$ on the strength margin and $g_a - g_b = -\frac{2\nu}{\sigma}$ for $\theta = \sigma$, that is, $\frac{2\nu}{\sigma}$ for $\theta = -\sigma$. Therefore,

$$\Xi_\xi^{PE,\nu} = 2\int k'\left(\xi\right) A\left(\xi\right) d\xi > 0 \qquad \text{and} \qquad \Xi_\xi^{PE,-\sigma} = \frac{\nu}{\sigma}\Xi_\xi^{PE,\nu} > 0,$$

since $k'\left(\xi\right) > 0$ for $\mu > 0$ as shown above. Both informativeness margins therefore carry a strictly positive production-efficiency component when $\mu > 0$, and a negative one for decreases in informativeness, with the noise-margin component proportional to the

strength-margin component with factor $\rho = \frac{\nu}{\sigma}$. When $\mu = 0$, $A(\xi) = 0$ and $k'(\xi) = 0$, so $\Xi_{\xi}^{PE} = 0$ and the endowment economy is recovered. □

D Extensions and Comparisons

D.1 Alternative Derivations

Quantile Representation. With a predetermined consumption profile $c_1(s)$, change variables from states to quantiles, $q = F(s;\theta)$, and let $s_\theta(q) = F^{-1}(q;\theta)$, so that $\int_{\underline{s}}^{\bar{s}} u(c_1(s)) f(s;\theta)\, ds = \int_0^1 u(c_1(s_\theta(q)))\, dq$: all dependence on θ sits in the integrand, and the uniform measure dq does not depend on θ. Differentiating $F(s_\theta(q);\theta) = q$ at fixed q gives $\frac{ds_\theta(q)}{d\theta} = \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f(s;\theta)}$, so consumption at a fixed quantile moves at the rate $\frac{dc_1(s_\theta(q))}{d\theta} = \frac{ds_\theta(q)}{d\theta}\frac{dc_1(s)}{ds} = x_\theta(s)$.

Differentiating under the integral and converting back to states with $dq = f(s;\theta)\, ds$,

$$\begin{aligned} p_\theta &= \frac{\beta}{u'(c_0)} \frac{d}{d\theta} \int_0^1 u(c_1(s_\theta(q)))\, dq \\ &= \frac{\beta}{u'(c_0)} \int_0^1 u'(c_1(s_\theta(q)))\, x_\theta(s_\theta(q))\, dq = \int_{\underline{s}}^{\bar{s}} \omega(s)\, x_\theta(s)\, ds, \end{aligned} \tag{OA.15}$$

with $\omega(s) = \frac{\beta u'(c_1(s))}{u'(c_0)} f(s;\theta)$: the probability pricing formula, where the chain rule produces the marginal utility without integration by parts. The consumption-equivalent cash flow $x_\theta(s)$ is the derivative of consumption with respect to θ holding the quantile fixed. When consumption responds to θ, $\frac{dc_1(s_\theta(q);\theta)}{d\theta} = \frac{\partial c_1(s;\theta)}{\partial \theta} + x_\theta(s)$, which gives the consumption/probability decomposition.

Integral Representation. With a predetermined consumption profile $c_1(s)$, write

$$u(c_1(s)) = u(c_1(\underline{s})) + \int_{\underline{s}}^{s} u'(c_1(t)) \frac{dc_1(t)}{dt} dt,$$

by the fundamental theorem of calculus, substitute into date-1 expected utility, and exchange the order of integration, using $\int_t^{\bar{s}} f(s;\theta)\, ds = 1 - F(t;\theta)$:

$$\int_{\underline{s}}^{\bar{s}} u(c_1(s)) f(s;\theta)\, ds = u(c_1(\underline{s})) + \int_{\underline{s}}^{\bar{s}} (1 - F(s;\theta)) u'(c_1(s)) \frac{dc_1(s)}{ds} ds. \tag{OA.16}$$

This integral representation (equivalently, one integration by parts) puts all dependence on θ in the survival function; the baseline $u(c_1(\underline{s}))$ and the marginal utility $u'(c_1(s))\frac{dc_1(s)}{ds}$ do not depend on θ.

Differentiating Equation (OA.16) with respect to θ and scaling by $\beta/u'(c_0)$, the baseline drops out and $p_\theta = \frac{\beta}{u'(c_0)}\int_{\underline{s}}^{\bar{s}} \frac{d(1-F(s;\theta))}{d\theta} u'(c_1(s))\frac{dc_1(s)}{ds}ds = \int_{\underline{s}}^{\bar{s}} \omega(s)\, x_\theta(s)\, ds$, with $\omega(s) = \frac{\beta u'(c_1(s))}{u'(c_0)} f(s;\theta)$ and $x_\theta(s) = \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f(s;\theta)}\frac{dc_1(s)}{ds}$: the probability pricing formula. Relative to the proof of Proposition 1, the order of the steps is reversed, rewrite first and differentiate second, so the boundary term is $u(c_1(\underline{s}))$, nonzero but independent of θ, rather than $u(c_1(s))\frac{d(1-F(s;\theta))}{d\theta}$, which vanishes at both ends. When consumption responds to θ, the product rule separates $\int_{\underline{s}}^{\bar{s}} u'(c_1(s;\theta))\frac{\partial c_1(s;\theta)}{\partial\theta} f(s;\theta)\, ds$, which needs no rewriting, from $\int_{\underline{s}}^{\bar{s}} u(c_1(s;\theta))\frac{\partial f(s;\theta)}{\partial\theta}ds$, to which Equation (OA.16) applies with $c_1(s;\theta)$ frozen; together they give the consumption/probability decomposition.

D.2 Discrete States

The counterpart to Equation (1) in the discrete case is given by

$$V = u(c_0) + \beta\sum_s \pi(s;\theta)\, u(c_1(s)),$$

where $\pi(s)$ denotes the probability of one of the finitely many states $s \in \{1,\ldots,S\}$. The agent's valuation for a marginal change in probabilities satisfies the formula

$$p_\theta = \sum_s \frac{d\pi(s;\theta)}{d\theta}\frac{\beta u(c_1(s))}{u'(c_0)},$$

where it must be that $\sum_s \frac{d\pi(s;\theta)}{d\theta} = 0$. Summation by parts (Courant and John, 1989) implies that

$$\sum_s u(c_1(s))\frac{d\pi(s;\theta)}{d\theta} = \sum_{s=1}^{S-1}\left(u(c_1(s)) - u(c_1(s+1))\right)\frac{d\Pi(s;\theta)}{d\theta},$$

where $\frac{d\Pi(s;\theta)}{d\theta} = \sum_{u=1}^{s}\frac{d\pi(u;\theta)}{d\theta}$ denotes the change in cdf, with $\Pi(0;\theta) = 0$, and where we use the fact that $\frac{d\Pi(S)}{d\theta} = 0$. Therefore, we can express p_θ as

$$p_\theta = \beta\sum_{s=1}^{S-1}\pi(s;\theta)\left(\frac{u(c_1(s+1)) - u(c_1(s))}{u'(c_0)}\right)\frac{\frac{d(1-\Pi(s;\theta))}{d\theta}}{\pi(s;\theta)},$$

which is the exact counterpart of Equation (4). Formally, as the difference between $s+1$ and s becomes small,

$$u\left(c_1\left(s+1\right)\right)-u\left(c_1\left(s\right)\right)\rightarrow u'\left(c_1\left(s\right)\right)\frac{dc_1\left(s\right)}{ds}.$$

Binomial Case. In the binomial case, with states denoted by $s\in\{L,H\}$, the probability price can be written as

$$p_\theta=\beta\pi\left(L;\theta\right)\left(\frac{u\left(c_1\left(H\right)\right)-u\left(c_1\left(L\right)\right)}{u'\left(c_0\right)}\right)\frac{\frac{d(1-\Pi(L))}{d\theta}}{\pi\left(L;\theta\right)}=-\beta\left(\frac{u\left(c_1\left(H\right)\right)-u\left(c_1\left(L\right)\right)}{u'\left(c_0\right)}\right)\frac{d\pi\left(L\right)}{d\theta},$$

where the second equality uses the fact that $\Pi\left(L\right)=\pi\left(L\right)$ in the binomial case. Even though the binomial case is the easiest to compute, it does not clearly identify that the shift in the cdf is the relevant object in general, since a change in the probability of one state necessarily implies an opposite change in the other. For this reason, we do not use it as a benchmark in the paper.

D.3 Leisure

When preferences include both consumption and leisure, then $U\left(s\right)=\beta u\left(c_1\left(s\right),n_1\left(s\right)\right)$, so

$$\frac{dU\left(s\right)}{ds}=\beta\frac{\partial u\left(c_1\left(s\right),n_1\left(s\right)\right)}{\partial c}\left(\frac{dc_1\left(s\right)}{ds}+\frac{\frac{\partial u(c_1(s),n_1(s))}{\partial n}}{\frac{\partial u(c_1(s),n_1(s))}{\partial c}}\frac{dn_1\left(s\right)}{ds}\right). \tag{OA.17}$$

The term in parentheses in Equation (OA.17) corresponds to a consumption-equivalent, expressing changes in $n_1\left(s\right)$ in consumption units. Hence, all results in the paper apply using this leisure-augmented consumption-equivalent. The same logic applies to environments with multiple consumption goods.

D.4 Multidimensional States

The results in the paper are stated for a scalar state s. Here we show that Proposition 1 extends to a vector of states $s=(s_1,\ldots,s_n)$ with support $\Omega\subseteq\mathbb{R}^n$, pdf $f\left(s;\theta\right)>0$, and

$\int_\Omega \frac{df(s;\theta)}{d\theta} ds = 0$.[20] In this case, the counterpart of Equation (3) is

$$p_\theta = \int_\Omega \frac{\beta u\left(c_1\left(s\right)\right)}{u'\left(c_0\right)} \frac{df\left(s;\theta\right)}{d\theta} ds. \tag{OA.18}$$

With a scalar state, integration by parts replaces the change in the pdf for the change in the cdf. The divergence theorem plays the same role with a vector of states. Define a vector field $J\left(s;\theta\right) : \Omega \to \mathbb{R}^n$ by

$$\frac{df\left(s;\theta\right)}{d\theta} = -\nabla_s \cdot J\left(s;\theta\right), \tag{OA.19}$$

with $J\left(s;\theta\right) \cdot \mathbf{n}\left(s\right) = 0$ on the boundary $\partial\Omega$, where $\mathbf{n}\left(s\right)$ denotes the outward normal (when Ω is unbounded, J must vanish sufficiently fast at infinity). When $n = 1$, Equation (OA.19) and the boundary condition imply $J\left(s;\theta\right) = \frac{d(1-F(s;\theta))}{d\theta}$.

Let $U\left(s\right) = \beta u\left(c_1\left(s\right)\right)$. Since $\nabla_s \cdot \left(UJ\right) = \nabla_s U \cdot J + U \nabla_s \cdot J$, the divergence theorem and the boundary condition on J imply that

$$\int_\Omega U\left(s\right) \frac{df\left(s;\theta\right)}{d\theta} ds = -\underbrace{\oint_{\partial\Omega} U\left(s\right) J\left(s;\theta\right) \cdot \mathbf{n}\left(s\right) dS}_{=0} + \int_\Omega \nabla_s U\left(s\right) \cdot J\left(s;\theta\right) ds.$$

Using the chain rule, $\nabla_s U\left(s\right) = \beta u'\left(c_1\left(s\right)\right) \nabla_s c_1\left(s\right)$, so Equation (OA.18) becomes

$$p_\theta = \int_\Omega \omega\left(s\right) x_\theta\left(s\right) ds, \quad \text{where} \quad x_\theta\left(s\right) = \frac{J\left(s;\theta\right)}{f\left(s;\theta\right)} \cdot \nabla_s c_1\left(s\right) = \sum_{k=1}^{n} \frac{J_k\left(s;\theta\right)}{f\left(s;\theta\right)} \frac{\partial c_1\left(s\right)}{\partial s_k}, \tag{OA.20}$$

and where $\omega\left(s\right) = \frac{\beta u'(c_1(s))}{u'(c_0)} f\left(s;\theta\right)$ is the state price, as in Equation (4). Equation (OA.20) is the exact counterpart of Equations (4) and (5).

As an example, for a Gaussian state, $s \sim \mathcal{N}\left(\mu, \Sigma\left(\theta\right)\right)$, the density satisfies $\frac{df(s;\theta)}{d\theta} = \frac{1}{2}\mathrm{tr}\left(\Sigma'\left(\theta\right) \nabla_s^2 f\left(s;\theta\right)\right) = -\nabla_s \cdot \left(-\frac{1}{2}\Sigma'\left(\theta\right) \nabla_s f\left(s;\theta\right)\right)$, so

$$\frac{J\left(s;\theta\right)}{f\left(s;\theta\right)} = \frac{1}{2}\Sigma'\left(\theta\right) \Sigma\left(\theta\right)^{-1} \left(s - \mu\right),$$

which reduces to $\frac{s-\mu}{\sigma}$ when $n = 1$ and $\theta = \sigma$, as in Equation (10).

[20] We would like to thank our discussant Sergei Glebkin for pointing us towards this generalization.

D.5 Redistributional Concerns

In our applications, we have focused on characterizing (Kaldor-Hicks) efficiency. It is however straightforward to compute welfare gains for any welfarist social welfare function, given by $W = \mathcal{W}\left(V^1, \ldots, V^I\right)$, using the efficiency/redistribution decomposition in Dávila and Schaab (2025). Formally, the normalized marginal social welfare effects of any perturbation for any welfarist social welfare function can be expressed as

$$\frac{dW^\lambda}{d\theta} = \sum_i \omega^i \frac{dV^{i|\lambda}}{d\theta} = \underbrace{\sum_i \frac{dV^{i|\lambda}}{d\theta}}_{\Xi^E \text{ (Efficiency)}} + \underbrace{I \cdot \mathbb{C}ov_i\left[\omega^i, \frac{dV^{i|\lambda}}{d\theta}\right]}_{\Xi^{RD} \text{(Redistribution)}}, \tag{OA.21}$$

where $\frac{dV^{i|\lambda}}{d\theta} = \frac{\frac{dV^i}{d\theta}}{\lambda^i}$ and $\omega^i = \frac{\frac{\partial \mathcal{W}}{\partial V^i}\lambda^i}{\frac{1}{I}\sum_i \frac{\partial \mathcal{W}}{\partial V^i}\lambda^i}$. This is the unique decomposition in which a normalized welfare assessment can be expressed as Kaldor-Hicks efficiency, Ξ^E, and its complement, Ξ^{RD}. The choice of λ^i simply accounts for the chosen unit to make interpersonal comparisons (welfare numeraire). Hence, perturbations in which $\Xi^E > 0$ can be turned into Pareto improvements if transfers are feasible and costless. Given (OA.21), it is thus straightforward to augment our analysis to also draw insights over Ξ^{RD}.

D.6 Probability Pricing vs. Comparative Statics of Cash Flow Pricing

The valuation for a marginal change in probabilities (probability pricing) differs from the change in the valuation for an asset induced by a change in probabilities (comparative statics of cash-flow pricing). This distinction arises even in the canonical representative-agent, consumption-based asset-pricing model à la Lucas (1978), when the asset's payoff coincides with aggregate consumption.

Formally, interpreting the environment in Section 2 as a representative-agent economy à la Lucas (1978), the price of a claim to aggregate consumption satisfies

$$\frac{p_x}{c_0} = \int_{\underline{s}}^{\overline{s}} \frac{\beta u'\left(c_1\left(s\right)\right)}{u'\left(c_0\right)} \frac{c_1\left(s\right)}{c_0} f\left(s\right) ds,$$

since $x\left(s\right) = c_1\left(s\right)$ in equilibrium. Parametrizing the pdf by θ, the sensitivity of the price

of this claim to a marginal change in probabilities is

$$\frac{d\left(\frac{p_x}{c_0}\right)}{d\theta} = \int_{\underline{s}}^{\bar{s}} \frac{\beta u'\left(c_1\left(s\right)\right)}{u'\left(c_0\right)} \frac{c_1\left(s\right)}{c_0} \frac{\partial \ln f\left(s;\theta\right)}{\partial \theta} f\left(s;\theta\right) ds. \tag{OA.22}$$

This expression contrasts with the probability pricing exercise — computing the valuation for a marginal change in probabilities — which yields

$$\frac{p_\theta}{c_0} = \int_{\underline{s}}^{\bar{s}} \frac{\beta u'\left(c_1\left(s\right)\right)}{u'\left(c_0\right)} \frac{c_1\left(s\right)}{c_0} \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f\left(s;\theta\right)} f\left(s;\theta\right) ds, \tag{OA.23}$$

after assuming that $c_1\left(s\right) = e^s c_0$, so $\frac{\partial c_1(s)}{\partial s} = c_1\left(s\right)$ — this is similar to the exponential consumption mapping in our climate application (Section 3).

The difference between (OA.22) and (OA.23) is easiest to see under log utility. In that case, $\frac{p_x}{c_0} = \beta$ and $\frac{d\left(\frac{p_x}{c_0}\right)}{d\theta} = 0$ for any change in the distribution $f\left(s\right)$. For example, increasing the mean of aggregate consumption growth leaves the willingness-to-pay for the consumption claimleaves the valuation for the consumption claim, p_x/c_0, unchanged, as cash-flow and discount-rate effects exactly offset. By contrast, such a shift in the mean implies a positive probability price, $p_\theta/c_0 > 0$, since $\frac{d(1-F(s;\theta))}{d\theta} > 0$ by Equation (9) in the text. This reasoning, which highlights the differences between computing an individual's valuation for a change in probabilities (probability pricing) and the comparative statics of cash-flow pricing with respect to changes in probabilities, also applies beyond the log-utility case.

D.7 Varying Consumption vs. Probabilities

A central implication of our results is that changes in consumption allocations (holding probabilities fixed) are equivalent to changes in probabilities (holding the consumption-state mapping fixed). Formally, as in Equation (7) in the text,

$$\frac{dV}{d\theta} = \beta \int_{\underline{s}}^{\bar{s}} u'\left(c_1\left(s;\theta\right)\right) \left(\frac{\partial c_1\left(s;\theta\right)}{\partial \theta} + \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f\left(s;\theta\right)} \frac{\partial c_1\left(s;\theta\right)}{\partial s} \right) f\left(s;\theta\right) ds.$$

This expression makes clear that the valuation for changes in either the consumption mapping $c_1(s;\theta)$ or the probability distribution $F(s;\theta)$ can be identical. This equivalence is useful to contrast, for example, two different ways of formulating the valuation for

changes in the business cycle. First, consider the approach of Alvarez and Jermann (2004), which holds probabilities fixed — so that $\frac{\partial(1-F(s;\theta))}{\partial\theta} = 0$ — and instead parameterizes changes in the consumption mapping:

$$c_1(s;\theta) = (1-\theta)\,\underline{c}_1(s) + \theta\bar{c}_1(s),$$

where $\underline{c}_1(s)$ denotes the status quo consumption allocation and $\bar{c}_1(s)$ a counterfactual allocation without fluctuations. In this case,

$$\frac{\partial c_1(s;\theta)}{\partial\theta} = \bar{c}_1(s) - \underline{c}_1(s) \neq 0.$$

Alternatively, one can consider perturbations that shift the probabilities of states while keeping the consumption mapping fixed, that is, $\frac{\partial c_1(s;\theta)}{\partial\theta} = 0$ but $\frac{\partial(1-F(s;\theta))}{\partial\theta} \neq 0$. The language in Lucas (1987) appears to describe such a change in the *distribution* of consumption, although the distinction between a change in the consumption mapping and a change in probabilities is not explicit — since Lucas's results are derived for a fully parameterized model in which expectations are explicitly computed. More generally, for marginal valuation computations, our results account exactly for changes in both consumption allocations and probabilities, so the conclusions do not depend on the approach taken.

D.8 Global Probability Prices

The main text prices marginal perturbations. This section shows that the marginal probability price p_θ is the derivative, evaluated at the status quo, of a global price defined as an equivalent variation, and that defining the global price as a compensating variation instead yields the same derivative.

Index the environment by θ and write lifetime utility as $V(\theta) \equiv u(c_0(\theta)) + \beta\int_{\underline{s}}^{\bar{s}} u(c_1(s;\theta))\,f(s;\theta)\,ds$. The marginal probability price is the date-0 consumption that compensates the agent for a change $d\theta$, $p_\theta = \frac{1}{u'(c_0(\theta))}\frac{dV(\theta)}{d\theta}$. When consumption is predetermined, $\frac{dV(\theta)}{d\theta} = \beta\int_{\underline{s}}^{\bar{s}} u(c_1(s))\frac{\partial f(s;\theta)}{\partial\theta}ds$, and integrating by parts gives the probability pricing formula $p_\theta = \int_{\underline{s}}^{\bar{s}} \omega(s)\,x_\theta(s)\,ds$.

The global price of moving from the status quo $\underline{\theta}$ to θ, $\mathcal{P}(\theta|\underline{\theta})$, is the date-0 consumption

that, added at the status quo, makes the agent as well off as under θ:

$$u\left(c_0\left(\underline{\theta}\right)+\mathcal{P}\left(\theta|\underline{\theta}\right)\right)+\beta\int_{\underline{s}}^{\bar{s}}u\left(c_1\left(s;\underline{\theta}\right)\right)f\left(s;\underline{\theta}\right)ds=V\left(\theta\right). \tag{OA.24}$$

This is an equivalent variation, and by construction $\mathcal{P}\left(\underline{\theta}|\underline{\theta}\right)=0$. The marginal probability price is the slope of the global price at the status quo:

$$\left.\frac{d\mathcal{P}\left(\theta|\underline{\theta}\right)}{d\theta}\right|_{\theta=\underline{\theta}}=p_{\underline{\theta}}. \tag{OA.25}$$

To see this, note that on the left of (OA.24) only $\mathcal{P}\left(\theta|\underline{\theta}\right)$ depends on θ, since the status-quo continuation term is a constant. Differentiating both sides in θ gives

$$u'\left(c_0\left(\underline{\theta}\right)+\mathcal{P}\left(\theta|\underline{\theta}\right)\right)\frac{d\mathcal{P}\left(\theta|\underline{\theta}\right)}{d\theta}=\frac{dV\left(\theta\right)}{d\theta}=u'\left(c_0\left(\theta\right)\right)p_{\theta}. \tag{OA.26}$$

At $\theta=\underline{\theta}$, $\mathcal{P}\left(\underline{\theta}|\underline{\theta}\right)=0$ implies that the marginal utility on the left is $u'\left(c_0\left(\underline{\theta}\right)\right)$, the same factor that multiplies $p_{\underline{\theta}}$ on the right, and dividing through yields (OA.25). Away from the status quo, (OA.26) shows that the slope of the global price is p_{θ} rescaled by $u'\left(c_0\left(\theta\right)\right)/u'\left(c_0\left(\underline{\theta}\right)+\mathcal{P}\left(\theta|\underline{\theta}\right)\right)$, which converts date-0 consumption at θ into date-0 consumption at the status quo. This factor equals one at $\underline{\theta}$, so $\mathcal{P}\left(\theta|\underline{\theta}\right)=p_{\underline{\theta}}\left(\theta-\underline{\theta}\right)$ up to second order.

Compensating Variation. The same conclusion holds if the global price is defined as the date-0 consumption taken away under θ that returns the agent to status-quo utility:

$$V\left(\underline{\theta}\right)=u\left(c_0\left(\theta\right)-\mathcal{P}^{CV}\left(\theta|\underline{\theta}\right)\right)+\beta\int_{\underline{s}}^{\bar{s}}u\left(c_1\left(s;\theta\right)\right)f\left(s;\theta\right)ds. \tag{OA.27}$$

Differentiating in θ and rearranging,

$$u'\left(c_0\left(\theta\right)-\mathcal{P}^{CV}\left(\theta|\underline{\theta}\right)\right)\frac{d\mathcal{P}^{CV}\left(\theta|\underline{\theta}\right)}{d\theta}=\frac{dV\left(\theta\right)}{d\theta}-\left[u'\left(c_0\left(\theta\right)\right)-u'\left(c_0\left(\theta\right)-\mathcal{P}^{CV}\left(\theta|\underline{\theta}\right)\right)\right]\frac{dc_0\left(\theta\right)}{d\theta}.$$

At $\theta=\underline{\theta}$, where $\mathcal{P}^{CV}\left(\underline{\theta}|\underline{\theta}\right)=0$, the bracket vanishes and the marginal utility on the left is $u'\left(c_0\left(\underline{\theta}\right)\right)$, so $\left.\frac{d\mathcal{P}^{CV}\left(\theta|\underline{\theta}\right)}{d\theta}\right|_{\theta=\underline{\theta}}=\left.\frac{d\mathcal{P}\left(\theta|\underline{\theta}\right)}{d\theta}\right|_{\theta=\underline{\theta}}=p_{\underline{\theta}}$. The equivalent and compensating variations share the same slope at the status quo and differ only at second order; they are mirror images of each other, $\mathcal{P}^{CV}\left(\theta|\underline{\theta}\right)=-\mathcal{P}\left(\underline{\theta}|\theta\right)$.

D.9 More General Preferences

It is straightforward to extend the results to more general preferences. Here we illustrate the Epstein–Zin case. It is equally straightforward to allow for robust control preferences.

Epstein–Zin. In the environment of Section 2, replace expected utility with Epstein–Zin preferences $V = u(c_0) + \beta u(\mu)$, where $\mu = h^{-1}\left(\int_{\underline{s}}^{\bar{s}} h(c_1(s)) f(s;\theta)\, ds\right)$ is the certainty equivalent of date-1 consumption, $u(\cdot)$ is the time aggregator, and $h(\cdot)$ is the risk aggregator; expected utility is the case $h = u$, and the isoelastic case is $u(c) = \frac{c^{1-1/\psi}}{1-1/\psi}$, $h(c) = \frac{c^{1-\gamma}}{1-\gamma}$. Since $\frac{d\mu}{dq} = \frac{1}{h'(\mu)} \int_{\underline{s}}^{\bar{s}} h'(c_1(s))\, x(s)\, f(s;\theta)\, ds$, cash-flow pricing gives $p_x = \int_{\underline{s}}^{\bar{s}} \omega(s)\, x(s)\, ds$ with state-price

$$\omega(s) = \beta \frac{u'(\mu)}{u'(c_0)} \frac{h'(c_1(s))}{h'(\mu)} f(s;\theta) = \beta \left(\frac{c_1(s)}{c_0}\right)^{-\frac{1}{\psi}} \left(\frac{c_1(s)}{\mu}\right)^{\frac{1}{\psi}-\gamma} f(s;\theta), \tag{OA.28}$$

where the second equality holds in the isoelastic case: the expected-utility state-price scaled by $(c_1(s)/\mu)^{\frac{1}{\psi}-\gamma}$, which raises the prices of states with consumption below μ when $\gamma > 1/\psi$.

For a perturbation $d\theta$ with predetermined consumption, $\frac{d\mu}{d\theta} = \frac{1}{h'(\mu)} \int_{\underline{s}}^{\bar{s}} h(c_1(s)) \frac{df(s;\theta)}{d\theta} ds$, so $p_\theta = \frac{\beta u'(\mu)}{u'(c_0) h'(\mu)} \int_{\underline{s}}^{\bar{s}} h(c_1(s)) \frac{df(s;\theta)}{d\theta} ds$. Integrating by parts, with boundary terms vanishing because $F(\underline{s};\theta) = 0$ and $F(\bar{s};\theta) = 1$ for all θ, $\int_{\underline{s}}^{\bar{s}} h(c_1(s)) \frac{df(s;\theta)}{d\theta} ds = \int_{\underline{s}}^{\bar{s}} h'(c_1(s)) \frac{dc_1(s)}{ds} \frac{d(1-F(s;\theta))}{d\theta} ds$, so that

$$p_\theta = \int_{\underline{s}}^{\bar{s}} \omega(s)\, x_\theta(s)\, ds, \quad \text{where} \quad x_\theta(s) = \frac{\frac{d(1-F(s;\theta))}{d\theta}}{f(s;\theta)} \frac{dc_1(s)}{ds} \tag{OA.29}$$

is the consumption-equivalent cash flow of the main text and $\omega(s)$ is the state-price in Equation (OA.28). Epstein–Zin preferences change only the state-price, so every use of probability pricing carries over with $\omega(s)$ as in Equation (OA.28).