

Behavioral Bias, Relative to What?

Robustness and Discounting in a Bewley Economy*

Lars Peter Hansen[†], Francesco Lippi[‡], and Thomas J. Sargent[§]

October 7, 2026

Abstract

We construct a linear-quadratic Bewley economy in which heterogeneity in discounting and in concerns about model misspecification leaves equilibrium quantities unchanged. A consumer's multiplier-preference problem reduces to a scalar Riccati equation whose solution yields an explicit locus of robustness parameters and discount factors along which consumption-saving rules coincide with a trusting benchmark's. Holding the initial allocation and the endowment law fixed, any distribution of types on the locus supports the same interest rate and joint distribution of consumption, debt, and income. Quantity data do not separately identify patience and robustness, even when the endowment law is known. Variation in income risk identifies the two separately. Types differ in worst-case forecasts and welfare evaluations. We bound admissible preference heterogeneity and compute detection error probabilities. Under an explicit reporting assumption, elicited discount factors and forecasts satisfy an Euler-equation restriction. Interpretations depend on assumptions that observed choices do not identify.

Keywords— Bewley model, linear-quadratic economies, multiplier preferences, model uncertainty, observational equivalence, identification, behavioral biases.

JEL codes: D81, D84, E21, D52, E71, C61

*Claude (Anthropic) assisted with verifications of some guesses, some derivations, numerical checks of the formulas, and Python code that generates the quantitative illustrations. The authors are responsible for the proofs, including the admissibility and boundary arguments.

[†]University of Chicago; lhansen@uchicago.edu

[‡]Luiss University and EIEF; flippi@luiss.it

[§]University of Pennsylvania and Hoover Institution. Email: sargentt@upenn.edu

Distrust and Caution are the parents of security.

Benjamin Franklin, 1733

1 Introduction

A behavioral bias is usually defined as a gap between what someone does and what a model says he should do. Measuring one requires a reference model. We ask how far quantity data constrain that choice in a class of heterogeneous-preference, linear-quadratic incomplete-markets economies in which an equilibrium can be verified explicitly without solving a distributional fixed-point problem. In standard Bewley models the equilibrium interest rate depends on a stationary distribution of wealth that itself depends on the interest rate (Bewley, 1977, Huggett, 1993, Aiyagari, 1994), and preference heterogeneity makes the fixed point a joint one across types (Krusell and Smith, 1998). Our construction does not solve that problem. Its equilibrium is nonstationary: the cross-section dispersion of consumption grows without bound,¹ utility has a bliss point, and admissible types are defined relative to a candidate interest rate. An imposed compensation between robustness and impatience makes the construction work.

We extend the “borrowers and lenders” closed economy of Ljungqvist and Sargent (2026, section 6.5),² in which *ex ante* identical consumers with quadratic preferences share a discount factor β , trade one risk-free bond, and face the equilibrium gross interest rate $R = \beta^{-1}$. Our consumers differ *ex ante* in their discount factors and in how much they distrust a common baseline model of their endowments. Each orders consumption streams by the multiplier preferences of Hansen and Sargent (2001, 2008). With distortions to both conditional means and conditional variances, these coincide with the risk-sensitive preferences of Jacobson (1973) and Hansen and Sargent (1995). A consumer who distrusts the baseline model saves more;³ a less patient consumer saves less. Hansen et al. (1999) show that the two forces can offset in quantities, and they feature how a concern about robustness alters saving in

¹A constant death rate with fair annuities maps the economy into a stationary version of itself without disturbing any identification result of Section 4: see Remark 5 and Appendix B.

²In the tradition of Aiyagari (1994) and Huggett (1993), the model combines components of the permanent income models of Friedman (1957), Hall (1978), and Bewley (1977, 1980). Wang (2003) obtains the permanent income hypothesis as a Bewley equilibrium with constant absolute risk aversion. See also Sargent and Stachurski (2026a). Sargent and Stachurski (2026b) set out the equilibrium of this paper and verify it numerically.

³This precautionary motive operates with quadratic preferences, which have linear marginal utility. It differs from the precautionary motive of Leland (1968), Kimball (1990), and Zeldes (1989), which requires convex marginal utility. It comes from guarding against misspecified conditional means of endowment shocks; the worst-case variance distortion leaves decisions unchanged (Section 3).

a risk-free asset. In a portfolio problem a concern about robustness looks like risk aversion (Maenhout, 2004).

A robust consumer's problem reduces to a scalar Riccati equation, which we solve in closed form. Its solution traces a one-dimensional locus of robustness parameters σ_i and discount factors $\hat{\beta}_i$ along which the marginal propensity to consume equals $1 - \beta$ (Proposition 3). Populate the economy with any distribution Φ of types on that locus. Hold fixed the initial distribution of assets and endowment states and the law that generates endowments. Then $R = \beta^{-1}$ clears markets, and the joint distribution of consumption, debt, and income does not depend on Φ (Proposition 4). Hansen et al. (1999) and Hansen and Sargent (2008, ch. 10) use observational equivalence to restrict one decision maker's robustness parameter. Nie et al. (2019) give every consumer in a Huggett economy with constant absolute risk aversion the same concern about robustness, which lowers the equilibrium interest rate. Here a whole distribution of types occupies the locus at $R = \beta^{-1}$, and the equilibrium is invariant to it.

Even when the law generating realized income is known, consumption and saving choices do not separately identify patience and concern about model misspecification (Corollary 2). The equivalence holds within an environment. Observing one consumer at two levels of income risk separates robustness from patience (Corollary 3). Behavioral and welfare interpretations therefore require assumptions beyond the observed decision rule. They require keeping apart the law that generates the data, each consumer's approximating model, his worst-case model, and the analyst's reference criterion (Section 6.1). Consumers on the locus have different worst-case forecasts (Proposition 5). A worst-case forecast is not by itself a belief. Our consumers fully trust neither the baseline model nor the worst-case model, and the worst case is a device for constructing a robust decision rule. Treating a worst-case forecast as a belief imposes subjective expected utility.

Along the locus a bias is a gap between two models, the reference model and the model attributed to the consumer. Within the maintained model, quantity data fix the decision rule, the locus, and the admissible range of types. They select no point on the locus, no distribution of types, and no reference model. Measured from the trusting benchmark, robust consumers' worst-case forecasts look pessimistic and their discount factors look impatient. Measured from a point near the robustness boundary, other consumers look optimistic and patient. Section 6.1 states the observational-equivalence class precisely and then describes it as a gauge freedom.

The model bounds that freedom. Admissible discount factors lie in $(\beta^2, \beta]$: no consumer on the locus discounts one year more heavily than the benchmark discounts two (Corollary 1). Every robust consumer, on or off the locus, satisfies an Euler-equation restriction that links his discount factor to the persistence of his worst-case forecasts; the locus adds the

martingale restriction on consumption. Under an explicit assumption about what survey respondents report, the Euler-equation restriction can be confronted with elicited discount factors and forecasts.

Several regularities in survey expectations have been read as departures from rational expectations. Households overestimate the persistence of shocks to their own income (Rozsypal and Schlafmann, 2023). Household forecasts of unemployment and inflation are pessimistic (Bhandari et al., 2025). Forecasts underreact to news in some data (Coibion and Gorodnichenko, 2015) and overreact in others (Bordalo et al., 2020). Our consumers’ worst-case forecasts display these patterns (Section 5). They overstate the persistence of shocks. Below bliss they are pessimistic. They overreact to news in samples in which distances from bliss are concentrated and underreact in samples in which those distances are dispersed. Our consumers optimize and doubt their model.⁴ Indarte et al. (2026) link elicited expectations to transactions data and measure departures from full information rational expectations; that exercise names its reference model. We calibrate the economy to none of these measurements. We use it to state what such readings presuppose.

Section 2 describes the benchmark economy and isolates the scalar α^2 , the innovation variance of consumption. Section 3 states the standing assumptions, reduces a robust consumer’s problem to a scalar Riccati equation, and computes the observational equivalence locus. Section 4 describes the equilibrium and the non-identification result. Section 5 describes worst-case forecasts. Section 6 describes what a bias is measured relative to, the freedom quantity data leave in choosing that reference, and what bounds and restricts it. Section 7 describes statistical discrimination between models. Section 8 describes welfare comparisons. Section 9 reports a numerical illustration. Section 10 concludes. Appendix A describes admissible robustness, Appendix B shows that a constant death rate restores a stationary cross-section without disturbing any identification result, and Appendix C contains the proofs.

2 The benchmark economy

We restate the “borrowers and lenders” closed economy of Ljungqvist and Sargent (2026, section 6.5), following their treatment and their notation closely, because everything that we add in later sections attaches to it.

⁴These studies concern forecasts of income and of macroeconomic aggregates; ours concern own consumption.

2.1 An individual consumer's problem

A consumer orders consumption streams by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t), \quad u(c_t) = -\frac{1}{2} (c_t - \gamma)^2, \quad (1)$$

where $\beta \in (0, 1)$ and $\gamma > 0$ is a bliss level of consumption. Consumption is allowed to be negative, which Ljungqvist and Sargent (2026, section 6.4) interpret as production of the good at a disutility cost.

The consumer chooses $\{c_t, b_{t+1}\}_{t=0}^{\infty}$ subject to the sequence of budget constraints

$$c_t + b_t = R^{-1}b_{t+1} + y_t, \quad t \geq 0, \quad (2)$$

where b_t is one-period risk-free debt falling due at the beginning of period t , b_0 is given, and R is a constant gross risk-free interest rate. We assume $R^{-1} = \beta$. We rule out Ponzi schemes by imposing

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t b_t^2 < +\infty. \quad (3)$$

The endowment has the state-space representation

$$z_{t+1} = Az_t + Cw_{t+1}, \quad y_t = Gz_t. \quad (4)$$

The eigenvalues of A determine the persistence of the endowment state z ; the matrix C describes how the shocks w_{t+1} impinge on z at $t + 1$; the row vector G maps z_t into the endowment y_t . The vectors z_t and w_{t+1} are column vectors; w_{t+1} is independently and identically distributed Gaussian with mean zero and identity contemporaneous covariance matrix; and the spectral radius $\rho(A) < \beta^{-1}$, so $(I - \beta A)^{-1}$ exists. That condition does not make the plans below satisfy (3). A deterministic income component that grows at rate $\lambda \in (\beta^{-1/2}, \beta^{-1})$ has finite human wealth but makes discounted squared debt diverge. We therefore impose the moment condition

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t H_t^2 < +\infty, \quad H_t \equiv G(I - \beta A)^{-1} z_t, \quad (5)$$

and assume that initial conditions have finite second moments. A spectral radius below $\beta^{-1/2}$ on the components of z_t that affect human wealth is sufficient for (5). The state confronting the consumer at t is $\begin{pmatrix} b_t & z_t^\top \end{pmatrix}^\top$.

Because marginal utility is linear, the Euler equation $\mathbb{E}_t u'(c_{t+1}) = u'(c_t)$ becomes

$$\mathbb{E}_t c_{t+1} = c_t. \quad (6)$$

Linearity does the work here (Ljungqvist and Sargent, 2026, section 6.4). Under the more standard assumptions $u' > 0$, $u'' < 0$, $u''' > 0$, Jensen's inequality applied to the Euler equation gives $\mathbb{E}_t c_{t+1} \geq c_t$, strictly when future consumption is conditionally nondegenerate, so consumption is a submartingale. That is the usual precautionary saving mechanism. Robustness concerns replace it below.

Define net wealth as

$$a_t \equiv G(I - \beta A)^{-1} z_t - b_t, \quad (7)$$

which is human wealth valued at the market discount factor $\beta = R^{-1}$ less debt. Combining the martingale (6) with the intertemporal budget constraint that (2) and (3) imply gives the consumption function

$$c_t = (1 - \beta) a_t, \quad (8)$$

and the law of motion

$$c_{t+1} = c_t + h w_{t+1}, \quad h = (1 - \beta) G(I - \beta A)^{-1} C. \quad (9)$$

Consumption is a random walk with innovation $h w_{t+1}$. Marginal utility is $\mu_t = \gamma - c_t$, so

$$\mu_{t+1} = \mu_t - h w_{t+1}. \quad (10)$$

2.2 A closed economy

Up to this point $R = \beta^{-1}$ has been imposed, an assumption usually excused by saying that one is studying a small open economy. In the spirit of Bewley (1977, 1980), Ljungqvist and Sargent (2026, section 6.5) fill out the environment so that $R = \beta^{-1}$ becomes an equilibrium outcome.

Let the consumer be one of a continuum of unit measure indexed by $i \in [0, 1]$, all of whom trade a risk-free one-period bond whose price is β . Consumers are *ex ante* identical: they share the preferences (1) and the endowment process (4). They are *ex post* heterogeneous in the realizations of their non-financial incomes. Initial vectors z_0^i are independent draws from a cross-section distribution that satisfies

$$\int_0^1 \mathbb{E}_0 y_j^i di = Y \quad \text{for every } j \geq 0, \quad (11)$$

where Y is the cross-section mean of non-financial income, and subsequent realizations are driven by idiosyncratic draws w_{t+1}^i that average to zero across the continuum at every date.⁵ Initial asset positions satisfy⁶

$$\int_0^1 b_0^i di = 0. \quad (12)$$

With (11), equation (8) gives

$$c_0^i = (1 - \beta) \sum_{j=0}^{\infty} \beta^j \mathbb{E}_0 y_j^i - (1 - \beta) b_0^i = (1 - \beta) a_0^i, \quad (13)$$

and integrating over consumers and using (12),

$$\int_0^1 c_0^i di = (1 - \beta) \sum_{j=0}^{\infty} \beta^j \int_0^1 \mathbb{E}_0 y_j^i di = (1 - \beta) \frac{Y}{1 - \beta} = Y. \quad (14)$$

The goods market clears at $t = 0$. Integrating (9) over consumers and using $\int_0^1 w_{t+1}^i di = 0$ gives $\int_0^1 c_{t+1}^i di = \int_0^1 c_t^i di$, so applying (14) recursively yields $\int_0^1 c_t^i di = Y$ for every t . The same condition makes the cross-section average endowment state evolve deterministically. Writing $\bar{z}_t \equiv \int_0^1 z_t^i di$ and $\bar{b}_t \equiv \int_0^1 b_t^i di$,

$$\bar{z}_{t+1} = A \bar{z}_t, \quad \bar{y}_t = G \bar{z}_t = Y, \quad \bar{c}_t = Y, \quad \bar{b}_{t+1} = R (\bar{b}_t + \bar{c}_t - \bar{y}_t) = 0, \quad (15)$$

where $\bar{y}_t = Y$ at every date follows from (11) and $\bar{b}_0 = 0$ from (12). Goods markets clear at every date, and risk-free loans are in zero net supply. Hence $R = \beta^{-1}$ is an equilibrium gross interest rate. Loans are made from one consumer to another within a closed set of agents.

Higher moments of the cross-section distribution change with the passage of time. From (9),

$$\mathbb{E}_0 c_t^i = c_0^i, \quad \mathbb{E}_0 (c_t^i - c_0^i)^2 = t h h^\top. \quad (16)$$

Each consumer's consumption follows a random walk. By (8) and (7),

$$c_t + (1 - \beta) b_t - y_t = \beta G (A - I) (I - \beta A)^{-1} z_t. \quad (17)$$

⁵We take the cross-section identities (15) and the cross-section distributions of Proposition 4 as primitives. A continuum of independent draws does not by itself deliver jointly measurable sample paths and exact cross-section distributions. Sun (2006) constructs a probability space on which it does.

⁶Condition (11) is stronger than requiring the cross-section mean of date-zero income to equal Y . It requires the cross-section mean of *forecast* income to equal Y at every horizon, which (14) below uses. For the two-factor endowment (18) it says that $\int_0^1 z_{1,0}^i di = Y$ and $\int_0^1 z_{2,0}^i di = 0$; a cross-section carrying a nonzero average transitory component would not clear the goods market at $t = 0$. Condition (12) replaces the assumption $b_0^i = 0$ for every i that Ljungqvist and Sargent (2026, section 6.5) make. Only the average matters, and the weaker condition leaves room for a nondegenerate initial distribution of financial wealth.

The right side is stationary when $A - I$ annihilates the nonstationary components of z_t , as it does when those components are semisimple unit roots and the rest are stable. The spectral bound of Section 2 alone does not deliver that. It admits mildly explosive roots, and two distinct nonstationary modes can leave no stationary combination of c_t , b_t , and y_t . For the two-factor endowment (18) below, $c_t + (1 - \beta)b_t - y_t = -\beta z_{2t}$, so the three series are jointly cointegrated. Consumption and debt alone are not cointegrated there, because consumption carries the permanent-income trend and debt does not. Debt then follows $b_{t+1} = b_t - z_{2t}$, in which net new borrowing stays transitory while the cross-section dispersion of outstanding positions widens, and average debt stays at zero. Without the transitory shock, $\eta_2 = 0$, and with $z_{2,0}^i = 0$, debt never moves from its initial value of zero. Cross-section consumption distributions spread out as a cohort ages; Deaton and Paxson (1994) and Storesletten et al. (2004) report evidence of such spreading.

2.3 A two-factor endowment process

For the numerical work in Section 9 we specialize (4) to

$$\begin{pmatrix} z_{1,t+1} \\ z_{2,t+1} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} z_{1t} \\ z_{2t} \end{pmatrix} + \begin{pmatrix} \eta_1 & 0 \\ 0 & \eta_2 \end{pmatrix} \begin{pmatrix} w_{1,t+1} \\ w_{2,t+1} \end{pmatrix}, \quad y_t = z_{1t} + z_{2t}. \quad (18)$$

The component z_{1t} is permanent and z_{2t} is purely transitory. A consumer hedges the permanent component least easily and smooths the transitory one most easily. We reserve σ for the robustness parameter and denote the standard deviations of the two shocks by η_1 and η_2 .

Remark 1. *The matrix $A = \text{diag}(1, 0)$ in (18) has $\rho(A) = 1$, which satisfies the assumption $\rho(A) < \beta^{-1}$ of Section 2 but leaves (4) without a stationary distribution. We assign an initial cross-section distribution for z_0^i with $\int_0^1 z_{1,0}^i di = Y$ and $\int_0^1 z_{2,0}^i di = 0$, as Sargent and Stachurski (2026a) do. This is condition (11) for the two-factor endowment, because the permanent component has a constant cross-section mean and the transitory component is forecast to be zero at every horizon $j \geq 1$. The Neumann series for $(I - \beta A)^{-1}$ converges because $\beta < 1$, and the market clearing argument of Section 2.2 uses only (11), (12), and the fact that idiosyncratic shocks average to zero across the continuum. The matrix $\text{diag}(1, 0)$ also satisfies the stronger sufficient condition of Section 2, that discounted second moments of the endowment be finite.*

2.4 The consumption innovation variance

Define

$$\alpha^2 \equiv hh^\top = (1 - \beta)^2 G(I - \beta A)^{-1} C C^\top (I - \beta A^\top)^{-1} G^\top. \quad (19)$$

For the two-factor endowment (18) we have $(I - \beta A)^{-1} = \text{diag}((1 - \beta)^{-1}, 1)$, so that $h = \begin{pmatrix} \eta_1 & (1 - \beta)\eta_2 \end{pmatrix}$ and

$$\alpha^2 = \eta_1^2 + (1 - \beta)^2 \eta_2^2. \quad (20)$$

A unit innovation to the permanent component z_1 passes into consumption one for one, because it is fully capitalized; a unit innovation to the transitory component z_2 passes into consumption with coefficient $1 - \beta$, because only its annuity value is consumed. Hence the entries η_1 and $(1 - \beta)\eta_2$ of h , and the two coefficients in (20).

The scalar α^2 is the one-step-ahead variance of the consumption innovation in (9). It is also the rate at which the cross-section variance of consumption grows with age; (16) reads $\mathbb{E}_0(c_t^i - c_0^i)^2 = t \alpha^2$. The product $\sigma \alpha^2$ determines the worst-case persistence ζ derived below.

3 A robust consumer's problem

A consumer doubts whether (4) is correctly specified. We call (4) the baseline or *approximating model*. He wants a consumption rule that performs well under it and under a neighborhood of statistically close alternatives. Section 3.1 describes his preferences and states the standing assumptions. Section 3.2 reduces his problem to a single scalar state. Section 3.3 solves the resulting Riccati equation in closed form. Section 3.4 gives the locus of $(\sigma, \hat{\beta})$ pairs along which his decision rule coincides with a trusting consumer's. Throughout this section β is the benchmark discount factor and $\hat{\beta}$ is this consumer's.

3.1 Multiplier preferences

To express these concerns we attribute to consumers what Hansen and Sargent (2001) and Hansen and Sargent (2008) call “multiplier preferences”, to distinguish them from the “constraint preferences” in which the σ parameter below is itself a function of a deeper entropy budget.⁷ To represent alternative models, we let a minimizing agent replace the standard normal conditional distribution of the shock w_{t+1} by another conditional distribution. Let

⁷See https://python-advanced.quantecon.org/five_preferences.html for an elementary user guide to multiplier, constraint, and risk-sensitive preferences. Hansen and Sargent (2001) and Hansen and Sargent (2008) discuss the connections between these preferences, those of Jacobson (1973) and Whittle (1981), and the maxmin expected utility of Gilboa and Schmeidler (1989).

m_{t+1} be the ratio of the alternative conditional density of w_{t+1} to the standard normal density, so that $m_{t+1} \geq 0$ and $\mathbb{E}_t m_{t+1} = 1$. The ratio m_{t+1} can be a history-dependent, nonlinear function of $\{z_s\}_{s=-\infty}^t$ and w_{t+1} . One class of alternatives shifts the conditional mean of the shock by v_{t+1} ,

$$z_{t+1} = Az_t + C(w_{t+1} + v_{t+1}), \quad (21)$$

and leaves its conditional covariance matrix at I . Other alternatives also change the covariance matrix.

Let $x_t = \begin{pmatrix} 1 & b_t & z_t^\top \end{pmatrix}^\top$ and

$$V(x) \equiv -\frac{1}{2}x^\top \mathcal{P}x - \mathbf{p}_{\text{vec}}. \quad (22)$$

The state includes a constant because the bliss point puts terms linear in (b_t, z_t) into the value function. We normalize the constant entry of $\mathcal{P}$ to zero, so that the scalar $\mathbf{p}_{\text{vec}}$ collects the terms that do not depend on the state. The scalar reduction of Section 3.2 centers wealth at bliss and writes its value function with a constant $\mathbf{p}_s$; the two differ by a bliss term, as (83) records. Following Hansen and Sargent (1995), Hansen et al. (1999), and Hansen and Sargent (2008, ch. 3),⁸ the consumer's value function solves the Bellman equation

$$V(x) = \sup_c \inf_{m' \geq 0, \mathbb{E}m'=1} \left\{ -\frac{1}{2}(c - \gamma)^2 + \hat{\beta} \mathbb{E} [m' V(x') + \theta m' \log m'] \right\}. \quad (23)$$

A prime denotes a next-period value. The density ratio m' is a function of next period's shock w' , and it multiplies the continuation value $V(x')$. The expectation $\mathbb{E}$ is taken under the approximating model (4), and $\hat{\beta}$ is the consumer's discount factor.

The quantity $\mathbb{E} [m' \log m']$ is the conditional relative entropy of the alternative shock distribution with respect to the baseline one, and $\theta \in (0, \infty]$ prices it. We work with the parameter

$$\sigma \equiv -\theta^{-1} \leq 0. \quad (24)$$

Remark 2 (Entropy minimization and the log-exp form). *For a continuation value V' ,*

$$\min_{m' \geq 0, \mathbb{E}m'=1} \mathbb{E} [m' V' + \theta m' \log m'] = -\theta \log \mathbb{E} \exp \left(-\frac{V'}{\theta} \right) = \frac{1}{\sigma} \log \mathbb{E} \exp (\sigma V'), \quad (25)$$

attained at the minimizer $m^ = \exp(\sigma V') / \mathbb{E} \exp(\sigma V')$. Substituting (25) into (23) gives*

⁸Hansen and Sargent (1995) constructed an infinite-horizon recursion under discounting that extends the risk-sensitivity framework of Jacobson (1973), but did not feature the robustness connection that Jacobson (1973) noted. Hansen et al. (1999) implemented that recursion and featured robustness and what later came to be called multiplier preferences. With minimization over the entire conditional distribution of the shock, as in (23), the two are one preference ordering (Remark 2).

the risk-sensitive recursion $V = \sup_c \left\{ u(c) + \hat{\beta} \sigma^{-1} \log \mathbb{E} \exp(\sigma V') \right\}$ of Hansen and Sargent (1995). Multiplier preferences and risk-sensitive preferences are one preference ordering. We adopt the robustness interpretation. Section 6.6 discusses the choice.

An asterisk marks a minimizer. The minimizer m^* tilts the standard normal density of w' exponentially by the continuation value. Because the value function is quadratic, the worst-case conditional distribution of w' is normal. It has a mean v^* and a covariance matrix Σ^* weakly larger, in the positive-semidefinite order, than I , and its conditional relative entropy is

$$\frac{1}{2} \left[(v^*)^\top v^* + \text{trace}(\Sigma^* - I) - \log \det \Sigma^* \right] \quad (26)$$

(Hansen and Sargent, 2008, ch. 3). We let $\hat{\mathbb{E}}$ denote an expectation taken under the worst-case law. Restricting the minimizing agent to the mean shifts (21), as in Hansen et al. (1999, eq. 9), gives the same matrix $\mathcal{P}$, the same decision rule, and the same worst-case mean v^* ; only the constant $\mathbf{p}_s$ differs (Appendix C.3).⁹

The worst case also inflates the variance of the shock. That changes neither decisions nor the interest rate. Consumption and bond holdings do not alter the loading of next period's state on the shock w_{t+1} , so the worst-case covariance affects only the constant $\mathbf{p}_s$. The consumer's Euler equation $\hat{\beta} R \hat{\mathbb{E}}_t \mu_{t+1} = \mu_t$ involves the worst-case conditional mean of marginal utility and not its variance, so the covariance distortion leaves the equilibrium gross interest rate R unchanged. It matters for welfare.¹⁰

Setting $\sigma = 0$ recovers (1). Decreasing σ below zero expresses concerns about model misspecification. For $\sigma < 0$ admissible plans have finite discounted second moments under the worst case, at the consumer's own discount factor.¹¹

Standing assumptions. (i) Consumption risk is nondegenerate: $\alpha > 0$. (ii) The moment condition (5) holds and initial conditions (b_0^i, z_0^i) have finite second moments. (iii) Two admissibility restrictions play different roles. Financial solvency (3), at the market discount factor, comes from the trading environment. Integrability $\hat{\mathbb{E}}_0 \sum_{t=0}^{\infty} \hat{\beta}^t \tilde{a}_t^2 < +\infty$, where $\tilde{a}_t$ is net wealth relative to bliss (27), under the consumer's worst-case law and at his own discount

⁹Suppressing the random innovations while keeping the adversarial disturbance channel leaves a deterministic problem with the same $\mathcal{P}$, decision rule, and v^* , and with $\mathbf{p}_s = 0$. That is not the same as setting $C = 0$, which removes the channel and changes the robust rule. Appendix C.3 computes the three constants.

¹⁰A consumer who chose a portfolio of risky securities would control that loading, and the worst-case covariance would then affect his decisions. Hansen and Sargent (2001) work in continuous time, where the covariance distortion vanishes because an absolutely continuous change of measure alters only drifts (Hansen et al., 1999, fn. 23).

¹¹The robust linear regulator of Hansen and Sargent (2008) restricts sequences in the same way. The restriction keeps the value of the game between the consumer and the minimizing agent finite. It changes the admissible set, not notation.

factor, defines his preference criterion. A plan is admissible for a consumer with parameters $(\sigma, \hat{\beta})$ if it depends only on the history of endowment shocks and satisfies both restrictions. A preference parameter does not remove a restriction that the trading environment imposes. We solve an auxiliary problem on the integrability domain alone; when its solution also satisfies (3), as every rule on the locus does by (5), that solution also solves the problem on the intersection. Auxiliary solutions that violate (3), among them the bliss plan of Proposition 9 and some plans of the impatient twin of Section 8.1, are not feasible equilibrium choices. (iv) Across the economies compared in Section 4, the joint distribution of initial conditions and the law that generates realized endowments are held fixed; idiosyncratic shocks are independent across consumers and independent of types.

The parameter σ carries units of consumption⁻², because θ is measured in units of utility per unit of relative entropy and utility here is measured in units of consumption squared. The product $\sigma\alpha^2$ is dimensionless. It is the only way in which σ enters anything below, and Section 7 exploits this property.

3.2 A scalar reduction

As above, $\hat{\mathbb{E}}_t$ denotes a conditional expectation taken under the worst-case law and $\mathbb{E}_t$ one taken under the approximating model.

Proposition 1 (Scalar reduction). *The vector problem collapses to a scalar one. Let a_t be the net wealth defined in (7). Define deviations from bliss as*

$$\tilde{a}_t \equiv a_t - \frac{\gamma}{1 - \beta}, \quad \tilde{c}_t \equiv c_t - \gamma = -\mu_t. \quad (27)$$

Then the budget constraints (2) and the endowment process (4) imply the single scalar law of motion

$$\tilde{a}_{t+1} = R(\tilde{a}_t - \tilde{c}_t) + s\epsilon_{t+1}, \quad s \equiv \frac{\alpha}{1 - \beta}, \quad \epsilon_{t+1} \equiv \frac{hw_{t+1}}{\alpha} = \tilde{\epsilon}_{t+1} + \tilde{v}_{t+1}, \quad (28)$$

where ϵ_{t+1} is a standardized scalar shock.¹² Under the approximating model $\tilde{v}_{t+1} = 0$ and ϵ_{t+1} is standard normal. Under an alternative normal law $\tilde{\epsilon}_{t+1}$ has conditional mean zero and conditional variance $\tilde{\Sigma}_{t+1}$, so ϵ_{t+1} has conditional mean $\tilde{v}_{t+1}$ and conditional variance $\tilde{\Sigma}_{t+1}$. For a quadratic continuation value the unrestricted entropy minimizer is normal, so normal alternatives suffice here. The conditional relative entropy of the alternative with respect to

¹²The distortions $\tilde{v}_{t+1}$ and $\tilde{\Sigma}_{t+1}$ are kept standardized, like ϵ_{t+1} , so that entropy has the form (29); $s = \alpha/(1 - \beta)$ carries the units.

the baseline is

$$\mathcal{E}(\tilde{v}_{t+1}, \tilde{\Sigma}_{t+1}) \equiv \frac{1}{2} \left[\tilde{v}_{t+1}^2 + \tilde{\Sigma}_{t+1} - 1 - \log \tilde{\Sigma}_{t+1} \right]. \quad (29)$$

The distortion of the law of the vector w_{t+1} enters only through the conditional variance of hw_{t+1}/α and the scalar

$$\tilde{v}_{t+1} = \frac{hv_{t+1}}{\alpha}, \quad (30)$$

which parallels the standardized shock $\epsilon_{t+1} \equiv hw_{t+1}/\alpha$. Consumer i chooses $\tilde{c}$ to maximize and the minimizing agent chooses $(\tilde{v}, \tilde{\Sigma})$ to minimize

$$\mathbb{E}_0^{\tilde{v}, \tilde{\Sigma}} \sum_{t=0}^{\infty} \beta^t \left[-\frac{1}{2} \tilde{c}_t^2 + \hat{\beta} \theta \mathcal{E}(\tilde{v}_{t+1}, \tilde{\Sigma}_{t+1}) \right] \quad (31)$$

subject to (28), where $\mathbb{E}_0^{\tilde{v}, \tilde{\Sigma}}$ is taken under the alternative law indexed by the distortions $(\tilde{v}_{t+1}, \tilde{\Sigma}_{t+1})$.

Proof. See Appendix C, sections C.1 and C.2. □

Net wealth $\tilde{a}_t$ is the only state variable in (28).

The value function takes the form $-\frac{1}{2}\Omega\tilde{a}^2 - \mathbf{p}_s$ for a scalar Ω (Appendix C.3 computes the constant $\mathbf{p}_s$), and the decision rule is $\tilde{c}_t = \Omega\tilde{a}_t$, so that Ω is the marginal propensity to consume out of net wealth. In levels the rule reads

$$c_t = \Omega a_t + \gamma \left(1 - \frac{\Omega}{1 - \beta} \right). \quad (32)$$

3.3 A Riccati equation

Proposition 2. For $\sigma > \underline{\sigma}$ and $\hat{\beta}R^2 > 1$, the admissible interior region of Appendix A, the marginal propensity to consume Ω solves the scalar Riccati equation

$$1 + \sigma\Omega s^2 = \hat{\beta}R^2 (1 - \Omega). \quad (33)$$

Its solution is

$$\Omega(\sigma, \hat{\beta}) = \frac{\hat{\beta}R^2 - 1}{\sigma\alpha^2(1 - \beta)^{-2} + \hat{\beta}R^2}. \quad (34)$$

The worst-case and approximating laws of motion for $\tilde{a}$ are

$$\hat{\mathbb{E}}_t \tilde{a}_{t+1} = \zeta \tilde{a}_t, \quad \mathbb{E}_t \tilde{a}_{t+1} = R(1 - \Omega) \tilde{a}_t, \quad \zeta = \frac{R(1 - \Omega)}{1 + \sigma\Omega s^2} = \frac{1}{\hat{\beta}R}, \quad (35)$$

and the same two persistences govern $\mu_t = -\Omega\tilde{a}_t$. The minimizing agent's distortion to the conditional mean of ϵ_{t+1} is

$$s\tilde{v}_{t+1} = [\zeta - R(1 - \Omega)]\tilde{a}_t, \quad (36)$$

which on the locus (39) is $(\zeta - 1)\tilde{a}_t$. Under the worst-case law the conditional variance of ϵ_{t+1} is

$$\tilde{\Sigma} = \frac{1}{1 + \sigma\Omega s^2} = \frac{1}{\hat{\beta}R^2(1 - \Omega)}. \quad (37)$$

Proof. See Appendix C. □

Equation (33) is the reduced form, linear in Ω after dividing out the trivial root $\Omega = 0$, the bliss plan; the unreduced form is quadratic. Appendix C solves the adversary's problem for $(\tilde{v}, \tilde{\Sigma})$, then the consumer's now-ordinary problem for $\tilde{c}$, then matches coefficients of the value function. The first step replaces the robust-control problem by an ordinary linear-quadratic one with curvature operator $\mathcal{D}(\Omega) \equiv \Omega/(1 + \sigma\Omega s^2)$. For $\sigma < 0$ we have $\mathcal{D}(\Omega) > \Omega$: the minimizing agent raises the effective curvature of the continuation value. For $\sigma < 0$ the worst-case variance (37) exceeds one. The minimizing agent also makes the effect of a shock on net wealth more persistent than it is under the approximating model: by (35), ζ exceeds $R(1 - \Omega)$ by the factor $(1 + \sigma\Omega s^2)^{-1} > 1$. A permanent income consumer insures himself against permanent shocks least easily.

Setting $\sigma = 0$ and $\hat{\beta} = \beta$ in (34) gives $\Omega = 1 - \beta$, which is (8), and $\zeta = 1$, which is the martingale (6). The last equality in (35) is the consumer's Euler equation evaluated under his own worst-case model, $\hat{\beta}R\hat{\mathbb{E}}_t\mu_{t+1} = \mu_t$. It holds for every $(\sigma, \hat{\beta})$, not only on the locus constructed next, and it makes the worst-case persistence of marginal utility a ratio of two discount factors. On the locus $R(1 - \Omega) = 1$, so marginal utility has a unit root under the approximating model and is mildly explosive under the worst-case model.

3.4 Observational equivalence

The next proposition is due to Hansen et al. (1999). We apply it at the equilibrium interest rate $R = \beta^{-1}$ and compute α^2 in closed form for the endowment (18).

Proposition 3 (Observational equivalence). *Fix all parameters other than $(\sigma, \hat{\beta})$, and let $R = \beta^{-1}$. Let*

$$\underline{\sigma} \equiv -\frac{(1 - \beta)^2}{\alpha^2} < 0. \quad (38)$$

For every $\sigma \in (\underline{\sigma}, 0]$ the consumption-saving rule (8) chosen by a consumer with parameters

$(0, \beta)$ is also chosen by a consumer with parameters $(\sigma, \hat{\beta}(\sigma))$, where the locus function is

$$\hat{\beta}(\sigma) = \beta \left(1 + \frac{\sigma \alpha^2}{1 - \beta} \right). \quad (39)$$

This is equivalent to saying that the marginal propensity to consume satisfies $\Omega = 1 - \beta$ at every point on the locus. Along that locus

$$\zeta(\sigma) = \frac{\beta}{\hat{\beta}(\sigma)} = \left[1 + \frac{\sigma \alpha^2}{1 - \beta} \right]^{-1}, \quad (40)$$

and the worst-case conditional variance (37) of ϵ_{t+1} equals $\zeta(\sigma)$. For $\sigma < 0$ we have $\hat{\beta}(\sigma) < \beta$ and $1 < \zeta(\sigma) < R$, with $\zeta(\sigma) \rightarrow R$ as $\sigma \downarrow \underline{\sigma}$. For $\sigma \leq \underline{\sigma}$ the admissible interior construction does not support rule (8) (Proposition 9).

Proof. See Appendix C. □

Proposition 9 shows that an integrability-admissible interior solution requires $\sigma > \underline{\sigma}$ and $\hat{\beta}R^2 > 1$, and that it then satisfies $\hat{\beta}\zeta^2 < 1$. Under standing assumption (ii), every plan on the locus also satisfies the financial solvency restriction (3) under the approximating model. On the locus net wealth relative to bliss is a random walk, so $\mathbb{E}_0 \tilde{a}_t^2 = \tilde{a}_0^2 + ts^2$ and $\sum_t \beta^t \mathbb{E}_0 \tilde{a}_t^2 < +\infty$. Since $b_t = H_t - a_t$ and $(H_t - a_t)^2 \leq 2H_t^2 + 2a_t^2$, the moment condition (5) delivers (3).

Two things distinguish the economy below from earlier uses of the same apparatus. Hansen et al. (1999), Anderson et al. (2003), and Barillas et al. (2009) apply observational equivalence to one decision maker or to a representative agent, for asset pricing or for robust policy design, and use it to restrict a single robustness parameter. Here an entire distribution of types occupies the locus (39) at once, and the equilibrium rather than any one consumer's decision rule is invariant to the population (Proposition 4 below).¹³ The closed-form locus admits such a population and lets the equilibrium be computed by hand.

For a consumer below bliss, lowering σ raises saving, because the minimizing agent in (23) makes future endowments look worse than the approximating model says they are. Lowering $\hat{\beta}$ lowers saving. Along (39) the two movements cancel exactly, and the consumption plan is unchanged.¹⁴

¹³The locus (39) fixes the trusting benchmark and solves for the compensating $(\sigma, \hat{\beta})$; a different construction, fixing the robust consumer instead, yields a different, non-linear locus that does not support the aggregation of Proposition 4. Sargent and Stachurski (2026b) compare the two constructions directly.

¹⁴The script `check_decision_rules.py` in the replication files checks this algebra against the vector problem for the two-factor endowment (18). The closed-form value function is a fixed point of the robust Riccati map built from the $\mathcal{D}$ operator of Hansen and Sargent (2008), with residual below 10^{-12} , at five

Corollary 1 (How much heterogeneity the economy admits). *As σ ranges over the admissible interval $(\underline{\sigma}, 0]$, the discount factor $\hat{\beta}$ ranges over $(\beta^2, \beta]$. No consumer on the locus discounts one year more heavily than the benchmark discounts two. The lower end is the admissibility condition $\hat{\beta}R^2 > 1$ at $R = \beta^{-1}$, and the square is the one in that condition: admissibility restricts discounted squared net wealth. The implied rate of time preference $\hat{\beta}^{-1} - 1$ can therefore be inflated by a factor whose supremum, not attained, is*

$$\frac{\beta^{-2} - 1}{\beta^{-1} - 1} = \frac{1 + \beta}{\beta}, \quad (41)$$

slightly more than two for β near one. Neither bound depends on α^2 , and hence neither depends on the endowment process.

Proof. See Appendix C. □

Figure 1 plots the locus at the parameter values of Section 9.

4 Equilibria with heterogeneous concerns about misspecification

We describe a Bewley model with *ex ante* heterogeneous agents, all of whom share a common baseline approximating model (4) for each consumer's endowment process and all of whom face a common gross interest rate $R = \beta^{-1}$ on one-period risk-free loans. They do not share a discount factor or a degree of confidence in the baseline model. We call β the *benchmark* discount factor. It is the preference parameter of a consumer who trusts the approximating model of the endowment process, the reciprocal of the equilibrium interest rate, and the point $\sigma = 0$ on the observational equivalence locus (39). The distribution of consumer types across the locus is denoted Φ , supported on $\sigma_i \in (\underline{\sigma}, 0]$, which induces a distribution of $\hat{\beta}_i$ values through the bijection in equation (39) (see Corollary 1).

Definition 1. *An equilibrium consists of a gross interest rate R and, for each consumer i , a consumption-saving rule that solves consumer i 's multiplier problem (23) at that R , such that the goods market clears and bonds are in zero net supply at every date along the realized endowment paths.*

points on the locus and three off it. It delivers (8) on the locus and (34) off it, and the worst-case closed-loop dynamics are stable under discounting at $\hat{\beta}$. Iterating on the same map from zero, with a penalty $10^{-9}b^2/2$ on debt that selects the non-bliss root, reproduces (8) to eight digits at every tested point on the locus, including a zero coefficient on the bliss point. Off the locus that iteration converges at two of the three tested points.




Figure 1: The observational equivalence locus (39) at $\beta = 0.95$, $\eta_1 = 0.15$, $\eta_2 = 0.30$. Each point pairs a robustness parameter σ with the discount factor $\hat{\beta}(\sigma)$ that leaves the marginal propensity to consume at $1 - \beta$. The admissible discount factors are $(\beta^2, \beta]$, shaded. The right axis reports the worst-case persistence $\zeta = \beta/\hat{\beta}$, which runs opposite to $\hat{\beta}$ because persistence rises as discounting falls. Filled circles mark the four admissible types of Table 1. The open circle at $\underline{\sigma}$ is the breakdown point, which the admissible interval excludes; it is the boundary limit reported in the table's last row.

Market clearing requires a law that generates those realizations. Given the initial cross-section conditions (11) and (12), it requires realized shocks to average to zero across consumers at every date, $\int_0^1 w_{t+1}^i di = 0$, so that aggregate income stays at Y .¹⁵ The approximating model (4) has this property. Definition 1 therefore does not require any consumer's model to generate the data. Consumers with different σ_i hold different worst-case models. Taken as data-generating laws, those models lack that property whenever types are independent of consumption, a positive mass of consumers is robust, and $\gamma \neq Y$ (Section 6.1).

Those worst-case models differ across consumers, and they need not be jointly compatible with the aggregate resource constraint. They cause no trouble here because endowments are idiosyncratic and the only price is the constant R : no consumer's worst-case model of another consumer's endowment enters any decision rule or any market-clearing condition.

¹⁵The weaker condition that the consumption innovations hw_{t+1}^i average to zero does not suffice. It holds aggregate consumption at Y and lets aggregate income move. In the two-factor endowment (18), cross-section average innovations $\Delta \bar{z}_1 = -(1 - \beta)d$ and $\Delta \bar{z}_2 = d$ give $\int_0^1 hw_{t+1}^i di = \Delta \bar{z}_1 + (1 - \beta)\Delta \bar{z}_2 = 0$ and $\Delta \bar{y} = \Delta \bar{z}_1 + \Delta \bar{z}_2 = \beta d$. Goods do not clear at that date.

Proposition 4 (Equilibrium invariant to Φ). *Let consumer $i \in [0, 1]$ be indexed by a robustness parameter $\sigma_i \in (\underline{\sigma}, 0]$ drawn from a distribution Φ , and let consumer i have discount factor $\hat{\beta}_i \equiv \hat{\beta}(\sigma_i)$, so that every pair $(\sigma_i, \hat{\beta}_i)$ lies on the locus (39). Let initial conditions (b_0^i, z_0^i) have a fixed joint distribution with finite second moments satisfying (11) and (12), and let realized endowments follow (4) with shocks that satisfy $\int_0^1 w_{t+1}^i di = 0$ at every date. Then*

- (i) *every consumer chooses the consumption-saving rule (8) of a baseline $(0, \beta)$ consumer. Each $(\sigma_i, \hat{\beta}_i)$ lies on the locus, so $\Omega = 1 - \beta$ for every consumer and the consumption process c_t^i is a martingale under the baseline approximating model;*
- (ii) *$R = \beta^{-1}$ is an equilibrium gross interest rate;*
- (iii) *aggregate consumption is constant at Y . Under the approximating model (4) and standing assumption (iv), the cross-section distribution of consumption at date t is $F_t = F_0 * N(0, t\alpha^2)$, where F_0 is the distribution of c_0^i and $*$ denotes convolution, so its variance is $\text{Var}(c_0^i) + t\alpha^2$. Both are independent of Φ .*

Proof. See Appendix C. □

Parts (i) and (ii) hold path by path. For any common realization of initial conditions and endowment innovations, reassigning preferences along the locus leaves (c_t^i, b_t^i, y_t^i) unchanged, because every type uses the same rule. Applying the same probability law to these pathwise maps gives identical joint distributions. Part (iii) adds the approximating model and standing assumption (iv) to obtain the Gaussian cross-section.

The scalar α^2 in (39) is evaluated at the benchmark discount factor $\beta = R^{-1}$ rather than at any individual's $\hat{\beta}_i$, because the consumption innovation h is built from the market interest rate.

Proposition 4 ties our two forms of *ex ante* heterogeneity together. Consumer i with lower σ_i has a stronger precautionary motive; equation (39) assigns $\hat{\beta}_i$ just small enough to offset it.

Remark 3 (Scope of the result). *Three facts remove the fixed point. The linear-quadratic structure makes each consumer's decision rule linear in his net wealth (Section 5.1). The locus (39) makes the coefficient $\Omega = 1 - \beta$ the same for every type; off the locus the coefficient depends on the volatility of the endowment as well. Market clearing (15) then uses only the cross-section conditions (11) and (12) and the fact that the shocks themselves average to zero across the continuum. None of the three refers to a distribution of wealth. The candidate $R = \beta^{-1}$ is therefore verified rather than solved for, and the aggregation is exact: we appeal to no limit, no approximate aggregation theorem, and no numerical fixed point.*

Two things are assumed rather than derived. Proposition 4 places every consumer on the locus, which is an assumption about how nature drew $(\sigma_i, \hat{\beta}_i)$. We impose it deliberately. Off the locus decision rules differ across types and quantities depend on Φ , and $R = \beta^{-1}$ need not remain an equilibrium. At $R = \beta^{-1}$, robust consumers below bliss who discount at β save more than the trusting benchmark. When bliss exceeds mean income, $\gamma > Y$, that extra saving means that $R = \beta^{-1}$ does not clear the bond market. On the locus lower discount factors offset the extra saving exactly. When $\gamma = Y$ and types are independent of initial conditions, mean net wealth relative to bliss is zero within every type and stays zero, so $R = \beta^{-1}$ still clears markets off the locus. On the locus, quantities are invariant to Φ exactly rather than approximately, so Section 6 states an identification problem as an equality. We establish existence of an equilibrium at $R = \beta^{-1}$, not uniqueness. The locus is defined relative to a given interest rate, so the type space is indexed by R .

Remark 4 (Impatience). For every $\sigma_i < 0$ we have $\hat{\beta}_i < \beta = R^{-1}$, so $\hat{\beta}_i R < 1$: every consumer other than the benchmark type is strictly impatient relative to the market interest rate. Equivalently $R \leq \hat{\beta}_i^{-1}$ for all i , with equality only for the benchmark type. The equilibrium reproduces the inequality between the interest rate and the rate of time preference that is familiar from Bewley (1977), Huggett (1993), and Aiyagari (1994), but the wedge is produced by worst-case model distortion rather than by convex marginal utility, and the wedge is exactly large enough to leave the martingale (6) intact under the approximating model.

In the deterministic economy of Becker (1980), with heterogeneous discount factors and borrowing constraints, the most patient household eventually holds all the capital. In the tradition of Krusell and Smith (1998), dispersion in discount factors generates dispersion in wealth. Here the compensating variation in robustness eliminates the differences in consumption-saving rules that discount-factor heterogeneity would otherwise generate: every type on the locus has $\Omega = 1 - \beta$.

The risk-free bond is the only traded asset, so the only equilibrium price is R . Consumers' shadow valuations of risky payoffs differ across the locus, but no traded security prices the difference.

Proposition 4 has a non-identification corollary.

Corollary 2. Suppose the approximating model (4) generates the data, and hold fixed the joint distribution of initial conditions under the hypotheses of Proposition 4. Let Φ and Φ' be any two distributions supported in $(\underline{\sigma}, 0]$. The two economies they index generate identical joint distributions of $\{c_t^i, b_t^i, y_t^i\}$ at every date and identical equilibrium prices. No function of quantity data distinguishes them.

The equilibrium interest rate identifies $\beta = R^{-1}$, and consumption innovations identify their variance α^2 . Those two numbers pin down the locus (39).¹⁶ They say nothing about where along it any consumer sits.¹⁷ Under the approximating model every consumer's consumption is a martingale, so Hall's (1978) test of the permanent income hypothesis, that no variable known at t helps forecast $c_{t+1} - c_t$, is passed by every consumer in this economy. Consumers with $\sigma_i < 0$ reach the benchmark rule by evaluating continuation prospects under distorted laws (Section 5). Corollary 2 concerns preference identification under one probability law. It says nothing about how well data discriminate between that law and a worst-case law, which is the question of Section 7. The two questions are complementary, not the same question at different sample sizes. Econometrician A of Section 6.4 faces the situation of Corollary 2.

Remark 5 (Mortality and a stationary cross-section). *The cross-section dispersion of Proposition 4(iii) grows without bound because consumption is a martingale under the approximating model, Hall's (1978) random walk, and not because of robustness. A constant death rate $\chi \in (0, 1)$ with fair annuities maps the economy into one of the same form: survivors face $\hat{\beta}_i(1 - \chi)$ and $R/(1 - \chi)$ in place of $\hat{\beta}_i$ and R , and newborns enter every date with zero financial wealth and the average income state. $R = \beta^{-1}$ remains an equilibrium rate, the locus and observational equivalence survive with α^2 and $1 - \beta$ replaced by their survival-adjusted counterparts, and the equilibrium remains invariant to Φ for the same reasons as before. What changes is that the cross-section of consumption becomes stationary rather than drifting without bound (as does the cross-section of debt, absent an explosive endowment root), and the bound of Corollary 1 becomes $\tilde{\beta}_i \in (\tilde{\beta}^2, \tilde{\beta}]$ in the survival-adjusted factors $\tilde{\beta}_i = \hat{\beta}_i(1 - \chi)$ and $\tilde{\beta} = \beta(1 - \chi)$, equivalently $\hat{\beta}_i > \beta^2(1 - \chi)$. Proposition 10 in Appendix B states this precisely and gives the proof; Section B.4 of that appendix draws out a further consequence, that the share of the population above bliss converges to a value set by χ and $(\gamma - Y)/\alpha$ rather than to one half. Since no identification result above uses stationarity, this is a robustness check on the construction, not a repair of it.*

¹⁶Quantities alone need not identify β . If the endowment has no transitory component, $\eta_2 = 0$, and $b_0^i = z_{2,0}^i = 0$, then $c_t^i = y_t^i$ and $b_t^i = 0$ for every β , so the whole quantity process is invariant to β even though consumption risk is nondegenerate. Reading $1 - \beta$ off the decision rule requires enough variation in the state to separate the response to permanent and transitory shocks, here $\eta_2 > 0$. Within the model $R = \beta^{-1}$, so the observed interest rate settles the question without that condition; an analyst who does not maintain the model learns nothing about β from R .

¹⁷In the language of Manski (2003), quantity data set-identify $(\sigma_i, \hat{\beta}_i)$ conditional on the remaining primitives, and the identified set is the locus (39). The bliss point drops out of the rule (8), so quantity data do not identify γ either. Over a maintained set Γ of bliss points the identified set is $\{(\sigma, \hat{\beta}(\sigma), \gamma) : \sigma \in (\underline{\sigma}, 0], \gamma \in \Gamma\}$, whose projection on $(\sigma, \hat{\beta})$ is again the locus. Forecast levels, marginal utility, and the detection error probabilities of Section 7 all depend on γ . Corollary 1 bounds its projection on $\hat{\beta}$. Corollary 4 describes what elicited forecasts and discount factors add.

5 Worst-case forecasts

Consumers on the observational equivalence locus choose the same consumption rule and have different worst-case forecasts of their own future consumption. A worst-case forecast becomes a belief only under an added interpretation (Section 6.6). We use “belief” only where a reporting or subjective-expected-utility interpretation has been imposed.

5.1 Certainty equivalence

Simon (1956) and Theil (1957) established that an optimal decision rule for a linear-quadratic problem can be computed in two separate steps. Step one solves a perfect foresight problem: treat the future path of the forcing process as known and compute the best response to it. Step two replaces the unknown future values by their conditional means. The rule that emerges is the rule for the stochastic problem. The variances of the shocks affect the constant term in the value function but not the decision rule.

We follow the notation of the introduction to Lucas and Sargent (1981) but write f for their h , because h denotes the row vector in (9). Write a decision rule as a composition

$$u_t = f(z_t, q) = f_1[f_2(z_t), q], \quad (42)$$

in which q collects preference and technology parameters, f_2 maps the current state of the environment into the sequence of conditional means of its current and future values, and f_1 is the perfect foresight best response to that sequence. Without concerns about misspecification, the forecasting function f_2 depends on the law of motion of the environment and not on preferences or technology. The response function f_1 depends on preferences and technology and not on the law of motion.

The benchmark economy of Section 2 has this structure. Step two is the optimal linear prediction of the endowment process (4), which delivers $\sum_{j \geq 0} \beta^j \mathbb{E}_t y_{t+j} = G(I - \beta A)^{-1} z_t$. Step one is the annuity formula: consume the fraction $1 - \beta$ of the present value of the forecast path net of debt, which is (8).

Hansen and Sargent (2005) extended the Simon and Theil “certainty equivalence” principle, also called the “separation principle”, to a decision maker who distrusts his model.

Step one is unchanged as a function of the forecast sequence. The sequence fed into it changes: the conditional means are computed under the worst-case model rather than under the approximating model. Call the resulting forecasting function $\hat{f}_2$. The worst-case model depends on the value function, so $\hat{f}_2$ depends on preferences as well as on the law of motion, and the separation of f_2 from f_1 in (42) fails.

Distorted forecasts also depend on the volatility of the shocks, so the marginal propensity to consume (34) depends on $\sigma\alpha^2$. A consumer off the locus therefore has a decision rule that reflects the volatility of his endowment as well as his concern about misspecification. On the locus (39) that dependence disappears: $\Omega = 1 - \beta$ for every admissible σ and every α^2 . An observer who sees only the decision rules of Proposition 4 concludes that certainty equivalence holds, even though consumers with $\sigma_i < 0$ distrust the model that generates the data.

Composition (42) rephrases the identification problem of Section 4 in the language of Lucas and Sargent (1981). Quantity data reveal the composite function f . Consumers on the locus have different f_1 's, through $\hat{\beta}$, and different $\hat{f}_2$'s, through ζ , in offsetting ways. Rational expectations econometrics typically requires that the forecasting function coincide with the optimal predictor for the law of motion that generates the data; this requirement leads to the cross-equation restrictions between the parameters of the forecasting rule and the parameters of the environment described by Lucas (1976) and Sargent (1981). Section 6.4 returns to what such tests can and cannot detect here, and Section 6.2 to what would break the equivalence.

5.2 The worst-case model as a model of endowments

The minimizing agent's distortion (36) depends on net wealth $\tilde{a}_t$, which the consumer chooses. The recursive worst-case model is therefore not a model of the endowment process alone. An equivalent representation removes that feedback. Let the minimizing agent commit at date zero to a law for the endowment shocks, and let the consumer best respond to it.¹⁸

Let $\tilde{a}_t^*$ denote the net wealth of the equilibrium plan (8), with $\tilde{a}_0^* = \tilde{a}_0$, let $\mathcal{F}_t$ be the history of shocks through date t , and define the law $\mathbb{Q}^*$ by

$$\tilde{a}_{t+1}^* = \zeta \tilde{a}_t^* + s \nu_{t+1}, \quad \nu_{t+1} \mid \mathcal{F}_t \sim N(0, \zeta), \quad \epsilon_{t+1} = \frac{\zeta - 1}{s} \tilde{a}_t^* + \nu_{t+1}. \quad (43)$$

A consumer who departs from the equilibrium plan faces

$$\tilde{a}_{t+1} = R(\tilde{a}_t - \tilde{c}_t) + (\zeta - 1)\tilde{a}_t^* + s \nu_{t+1}. \quad (44)$$

The law depends on $\tilde{a}_t^*$, which the equilibrium plan determines, and not on the deviating consumer's $\tilde{a}_t$. A consumer who follows the equilibrium rule from the same initial condition

¹⁸With this timing the exogeneity of the law is immediate. We do not rely on the two orders of move delivering the same value; Hansen and Sargent (2008, ch. 7) give conditions under which they do.

has $\tilde{a}_t = \tilde{a}_t^*$ at every date; a consumer who departs from it does not. Against this fixed law the best response on the enlarged state is

$$\tilde{c}_t = k \tilde{a}_t + \beta(\zeta - 1)\tilde{a}_t^*, \quad k = 1 - \beta\zeta, \quad (45)$$

which is verified as follows. Substituting (45) into (44) gives $\tilde{a}_{t+1} = \zeta\tilde{a}_t + s\nu_{t+1}$, because $R\beta = 1$ cancels the terms in $\tilde{a}_t^*$ and $R(1 - k) = \zeta$. Both $\tilde{a}_t$ and $\tilde{a}_t^*$ then grow at rate ζ in conditional expectation, so $\hat{\mathbb{E}}_t \tilde{c}_{t+1} = \zeta \tilde{c}_t$, which with $\hat{\beta}R\zeta = 1$ is the Euler equation; and $\hat{\beta}\zeta^2 = \beta\zeta < 1$, so discounted second moments are finite. On the diagonal $\tilde{a}_t = \tilde{a}_t^*$ the rule (45) collapses to $\tilde{c}_t = (1 - \beta)\tilde{a}_t$. The equilibrium plan is therefore the best response to $\mathbb{Q}^*$ from the specified initial state, even though its off-diagonal feedback is not $(1 - \beta)\tilde{a}_t$. Appendix C.10 shows that it is the unique best response among plans with finite discounted second moments under $\mathbb{Q}^*$.

Once the law is fixed there is nothing left to minimize over, so the consumer faces an ordinary linear-quadratic problem with discount factor $\hat{\beta}$ and interest rate R . That is the problem of the impatient twin of Section 8.1, with the equilibrium plan's wealth entering as a forcing term: the own-wealth coefficient $k = 1 - \beta\zeta$ is exactly the twin's marginal propensity to consume (66). The consumer saves like the twin out of his own wealth and adds a term in $\tilde{a}_t^*$ that restores (8) on the diagonal. The twin is not an extra construct; he is this problem with the forcing term removed.

On the locus (39) the plan (8) gives $\tilde{a}_t = \tilde{a}_0 + s \sum_{j=1}^t \epsilon_j$, where $\epsilon_j \equiv hw_j/\alpha$ is the realized standardized endowment innovation, so the distortion (36) is

$$\tilde{v}_{t+1} = (\zeta - 1) \left[\frac{\tilde{a}_0}{s} + \sum_{j=1}^t \epsilon_j \right]. \quad (46)$$

Under the approximating model the ϵ_j are standard normal. In terms of the shocks ν_j of (43), which are the centered innovations under the worst-case law, the same distortion is

$$\tilde{v}_{t+1} = (\zeta - 1) \left[\zeta^t \frac{\tilde{a}_0}{s} + \sum_{j=1}^t \zeta^{t-j} \nu_j \right], \quad (47)$$

a geometric moving average whose initial-condition term carries the same factor ζ^t as the shocks farthest in the past. Along the equilibrium plan the distortion is a function of the initial condition and of past endowment shocks, and of nothing the consumer chooses. The worst-case model is then a probability model of the endowment process, an object that a survey could elicit. Forecasts of consumption are a different matter: they must incorporate

the plan being evaluated.

5.3 Forecasts across the observational equivalence locus

Although consumers with preferences on the observational equivalence locus behave identically, they do not forecast identically. In terms of the composition (42), consumers on the locus (39) have different f_1 's, each evaluated at that consumer's own $(\sigma_i, \hat{\beta}_i)$. They have different $\hat{f}_2$'s, each governed by $\zeta_i = (\hat{\beta}_i R)^{-1}$. The two differences offset in the composite. Every consumer's realized decision rule matches the benchmark $(0, \beta)$ consumer's, with coefficient $\Omega = 1 - \beta$. Neither f_1 nor $\hat{f}_2$ is the benchmark's; only their composition is.

Proposition 5. *Consumer i 's worst-case model implies the forecasts*

$$\gamma - \hat{\mathbb{E}}_t c_{t+j}^i = \zeta_i^j (\gamma - c_t^i), \quad j = 0, 1, 2, \dots, \quad (48)$$

where $\zeta_i = \zeta(\sigma_i) = \beta/\hat{\beta}_i$ is given by (40). For a consumer on the locus (39), $\mathbb{E}_t c_{t+j}^i = c_t^i$ for all j under the approximating model.

Proof. See Appendix C. □

A consumer below his bliss point faces a worst-case scenario in which consumption drifts away from bliss at the geometric rate ζ_i .¹⁹ Under the worst-case model the gap doubles in $\log 2 / \log \zeta_i$ periods. Section 9 reports doubling horizons for our parameter values. Section 5.2 makes the distortion a function of the endowment history alone, so the forecast does not respond to the consumer's choices. Its persistence ζ_i still comes from his problem.

Two consumers i and j with $\sigma_i < \sigma_j \leq 0$ make identical choices at every date and in every state, yet the ratio of their worst-case forecasts of $\gamma - c_{t+j'}$ at horizon j' is $(\zeta_i/\zeta_j)^{j'}$, which diverges as the horizon lengthens.

Worst-case forecasts over- or underreact to news. Coibion and Gorodnichenko (2015) regress the error of a forecast on its revision. A positive coefficient indicates underreaction and a negative one overreaction. Take consumer i 's one-period forecast error $c_{t+1}^i - \hat{\mathbb{E}}_t c_{t+1}^i$ and his revision $\hat{\mathbb{E}}_t c_{t+1}^i - \hat{\mathbb{E}}_{t-1} c_{t+1}^i$. By (48) and (9), under the approximating model

$$c_{t+1}^i - \hat{\mathbb{E}}_t c_{t+1}^i = (\zeta_i - 1) (\gamma - c_{t-1}^i - h w_t^i) + h w_{t+1}^i, \quad \hat{\mathbb{E}}_t c_{t+1}^i - \hat{\mathbb{E}}_{t-1} c_{t+1}^i = \zeta_i (\zeta_i - 1) (\gamma - c_{t-1}^i) + \zeta_i h w_t^i. \quad (49)$$

¹⁹Quadratic preferences make this worst-case drift symmetric about γ . The consumer fears that the *gap* $|\gamma - c|$ will grow, so a consumer who happens to be above the bliss point fears being showered with goods. This is an artifact of quadratic utility rather than of robustness, and it is the same artifact that makes bliss attainable in the first place.

In a sample of consumers of one type in which $\gamma - c_{t-1}^i$ has variance ς^2 , the covariance of the error with the revision is

$$\zeta_i(\zeta_i - 1) [(\zeta_i - 1)\varsigma^2 - \alpha^2]. \quad (50)$$

The revision has variance $\zeta_i^2 [(\zeta_i - 1)^2\varsigma^2 + \alpha^2]$, so the population regression coefficient of the error on the revision, fitted with an intercept, is

$$\psi_i = \frac{\zeta_i - 1}{\zeta_i} \frac{(\zeta_i - 1)\varsigma^2 - \alpha^2}{(\zeta_i - 1)^2\varsigma^2 + \alpha^2}. \quad (51)$$

It is zero for the trusting type and shares the sign of (50). Regressions of forecast errors on revisions report ψ_i . The calculation describes a cross-section of one type sampled so that the lagged consumption gap is independent of the new shocks. Pooling types, conditioning inclusion on realized shocks, or running a time-series regression on nonstationary levels changes it. Forecasts overreact when $\varsigma^2 < \alpha^2/(\zeta_i - 1)$ and underreact when $\varsigma^2 > \alpha^2/(\zeta_i - 1)$. The sign depends on how dispersed distances from bliss are in the sample, not on how far from bliss consumers are. In a cohort that starts from a common c_0 , $\varsigma^2 = (t - 1)\alpha^2$, so forecasts overreact for about $1/(\zeta_i - 1)$ periods and underreact afterwards. Halfway to the breakdown point at the parameter values of Section 9, $1/(\zeta_i - 1) = 39$ years. The studies cited in the introduction measure forecasts of macroeconomic variables; ours are forecasts of own consumption.

Why does a consumer not learn the correct specification? Robustness is a preference over decision rules, not a hypothesis under test. It values good performance against a whole neighborhood of models. A single realized path consistent with the approximating model no more refutes that preference than a lucky outcome refutes risk aversion. Section 7 supplies a complementary statistical answer. For a consumer near bliss the two models remain difficult to distinguish, under the criterion of Section 7, over samples of a few decades. Even a consumer inclined to test would find limited average discriminating power in the classification experiment of Section 7, which reports neither a rejection probability nor a posterior odds.²⁰

6 Biases relative to what?

A bias is a gap between two models. An analyst chooses both. Quantity data restrict the rationalizations to the locus and select no point on it.

²⁰Hansen and Sargent (2008, sec. 1.12) take up this question in more depth. Hansen and Sargent (2022) and Hansen and Sargent (2024) treat it more generally.

6.1 The observational-equivalence class as a gauge freedom

Fix the law that generates the data, the joint distribution of initial conditions, and the observables: the paths of c_t^i , b_t^i , and y_t^i , and the interest rate R . For those observables the locus (39) is an observational-equivalence class. Moving the whole population along it leaves their joint distribution unchanged (Corollary 2). It changes preferences, worst-case laws, welfare evaluations, and shadow valuations of risky payoffs. In that bounded sense the locus is a gauge freedom: a transformation of unobserved parameters that leaves the specified observables invariant.

To speak of a consumer's bias one must first choose a *reference type* r on the locus and then report departures from it. The benchmark type $\sigma = 0$ is one admissible choice of r and has no claim to priority. Let r have parameters $(\sigma_r, \hat{\beta}_r)$ and worst-case persistence ζ_r . Relative to r , consumer i displays

$$\underbrace{\frac{\zeta_i}{\zeta_r}}_{\text{excess pessimism}} = \underbrace{\frac{\hat{\beta}_r}{\hat{\beta}_i}}_{\text{excess impatience}}, \quad (52)$$

by (40). The two ratios are the same number, and their effects on saving cancel. Both are functions of r , and by Corollary 2 quantity data do not identify r . Setting r at $\sigma = 0$ makes every robust consumer's worst-case forecasts look pessimistic and his discount factor look impatient. Taking as reference a limiting point at the robustness boundary makes every consumer look optimistic and patient. Setting r in the middle splits the population into both kinds. The three descriptions are descriptions of the same economy. The biases in this economy are doubly subjective. One choice of r fixes both the model against which behavior is measured, σ_r , and the preferences supposed to have generated it, $\hat{\beta}_r$. Quantity data constrain neither. The Euler equation fixes the product $\hat{\beta}_i \zeta_i$ at R^{-1} for every robust consumer (Section 6.5), and quantity data identify R .

The reference type is not the only thing the analyst supplies. Four objects need to be kept apart: the law that generates the data, each consumer's approximating model, the worst-case model that he uses to evaluate a plan, and the analyst's reference criterion. Corollary 2 takes the approximating model to generate the data, and that assumption makes the $\sigma = 0$ consumer's forecasts the unbiased ones. Definition 1 does not need it. Suppose that consumer i 's endowment followed his worst-case model. His worst-case forecasts would then be correct, and the approximating model's forecast of his consumption would be the optimistic one. For one consumer near bliss, samples of a few decades make the two hypotheses hard to distinguish under the criterion of Section 7. The hypothesis cannot hold for all consumers

at once. On the locus $h v_{i,t+1} = \alpha \tilde{v}_{i,t+1} = (\zeta_i - 1)(1 - \beta) \tilde{a}_t^i = -(\zeta_i - 1)(\gamma - c_t^i)$ by (30), (36), and $\alpha/s = 1 - \beta$, so

$$\int_0^1 h v_{i,t+1} di = - \int_0^1 (\zeta_i - 1)(\gamma - c_t^i) di = -\mathbb{E}_i[\zeta_i - 1](\gamma - Y) + \text{Cov}_i(\zeta_i - 1, c_t^i), \quad (53)$$

where $\mathbb{E}_i$ and Cov_i are cross-section moments. When types are independent of consumption and a positive mass of consumers is robust, the distortions average to zero only if $\gamma = Y$. Otherwise aggregate income would not stay at Y , contrary to the construction of Section 2.2. An analyst cannot escape choosing a reference by promoting one consumer's worst-case model to the law that generates the data.

6.2 What bounds the freedom

Three restrictions limit how much bias can be attributed, and they work on different margins. Two of them take one economy as given. Section 6.3 takes up the third, which compares environments.

The first restriction is the bound of Corollary 1. The scale of endowment risk determines how large a $|\sigma|$ the economy tolerates, but it has no effect on the range of behavior that robustness can rationalize. A concern about misspecification can make a consumer look about twice as impatient as the benchmark, and no more, whatever his income process. The factor comes from the square in $\hat{\beta} R^2 > 1$, not from compounding. In the numerical example of Section 9, with $\beta = 0.95$, the admissible rates of time preference run from 5.3 percent up to, but not including, 10.8 percent per year. The bound limits how much impatience an analyst can attribute to a consumer to whom he also attributes a worst-case forecast. Evidence on that margin requires elicited discount factors and elicited forecasts from the same individuals (Section 6.5), with the discount factors elicited as Remark 7 describes.²¹

The second restriction is statistical: the detection error probability of Section 7, which Section 9 reports over a grid of pairs. Detection error probabilities are usually read as a restriction on how much robustness a decision maker may entertain. Read from the other side they measure how much behavioral bias may be attributed to a population before data

²¹The bound does not speak to dispersion in structurally estimated discount rates. Gourinchas and Parker (2002) estimate 3.94 to 5.71 percent across education groups. Lawrence (1991) estimates 12 to 19 percent across income, education and race in her preferred specification; her estimates are sensitive to the interest rate households are assumed to face. Frederick et al. (2002) survey a far wider range of estimates. The largest come from horizons of days or weeks, where annualizing magnifies small deviations. The estimators of Lawrence (1991) and Gourinchas and Parker (2002) assign households one probability model, the one that generates the data, and infer discounting and aversion to risk jointly under it. An econometrician who fits the homogeneous quadratic benchmark of this economy recovers β for every consumer. Those estimators maintain other models, so this is a statement about the benchmark and not about what they would report.

of a given length contradict the attribution under that criterion. Section 7.1 states the threshold that convention uses.

6.3 Variation in the environment

The third restriction compares economies. The two restrictions of Section 6.2 ask what may be attributed to a population observed once. The observational equivalence of Corollary 2 is itself a statement about one environment, and a consumer observed in two of them need not be unidentified.

Remark 6 (Equivalence within an environment, not across environments). *Consider two environments with different income risk, that is, with different values of α^2 . If every consumer sits on the locus of each environment, $\Omega = 1 - \beta$ for every consumer in both. The decision rule is then free of α^2 only because (39) assigns $\hat{\beta}$ as a function of $\sigma\alpha^2$; robustness alone would not make it so. Hold a consumer's $(\sigma, \hat{\beta})$ and the interest rate fixed instead, and change income risk from α_1^2 to α_2^2 . Staying on both loci requires*

$$\beta \left(1 + \frac{\sigma\alpha_1^2}{1-\beta} \right) = \beta \left(1 + \frac{\sigma\alpha_2^2}{1-\beta} \right), \quad (54)$$

so $\sigma = 0$ whenever $\alpha_1^2 \neq \alpha_2^2$. After the change a robust consumer's marginal propensity to consume (34) moves away from $1 - \beta$, while a trusting consumer's does not. Observational equivalence holds within an environment, not across environments. An exogenous change in income risk that leaves preferences and the interest rate unchanged would let quantity data separate σ_i from $\hat{\beta}_i$. An application would have to justify both of those assumptions, and in equilibrium the interest rate could itself change. Luo and Young (2010) add limited capacity to process information about permanent income. A change in that capacity breaks the observational equivalence in their model, and observationally equivalent types would pay different amounts for better information. In a filtering problem Kasa (2006) shows that a stronger concern about robustness and a larger capacity to process information can be observationally equivalent.

Corollary 3 inverts the map from marginal propensities to consume to $(\sigma, \hat{\beta})$ in closed form. The reciprocal of the marginal propensity to consume is affine in income risk; the slope identifies σ and the intercept identifies $\hat{\beta}$.

Corollary 3 (Identification from two environments). *Hold $(\sigma, \hat{\beta})$ and R fixed, and observe the marginal propensity to consume Ω_j in two environments whose human-wealth innovation variances s_j^2 differ, $s_1^2 \neq s_2^2$, both in the admissible interior region. Let $B = \hat{\beta}R^2$. By (34),*

$\Omega_j^{-1} = (B + \sigma s_j^2)/(B - 1) = A_0 + A_1 s_j^2$ with $A_0 = B/(B - 1)$ and $A_1 = \sigma/(B - 1)$. Since $A_0 - 1 = (B - 1)^{-1}$,

$$A_1 = \frac{\Omega_2^{-1} - \Omega_1^{-1}}{s_2^2 - s_1^2}, \quad A_0 = \Omega_1^{-1} - A_1 s_1^2, \quad \sigma = \frac{A_1}{A_0 - 1}, \quad \hat{\beta} = \frac{A_0}{(A_0 - 1) R^2}. \quad (55)$$

The empirical content depends on the hypotheses. The change in s_j^2 must be exogenous, preferences and the interest rate must be the same in both environments, and both Ω_j must be identified. Off the locus the variance of the consumption innovation is $\Omega_j^2 s_j^2$ rather than α_j^2 , so s_j^2 must be measured from income rather than from consumption.

6.4 Three econometricians

An empirical literature that tests for behavioral biases calls the following joint hypothesis FIRE, or “full information rational expectations”: f_2 is the optimal predictor for the law of motion that generates the data, and the decision maker conditions on the entire history of that law of motion. Every consumer with $\sigma_i < 0$ violates the first half, in the sense that the forecasting function embedded in his decision rule is not the optimal predictor for the law that generates the data, and none violates the second. What he would report is a separate matter (Section 6.6).

A test of whether $c_{t+1} - c_t$ is forecastable examines neither half directly. It examines the composite function f , which is the same for everybody.

A second kind of test attacks f_2 and bypasses the composite. Coibion and Gorodnichenko (2015) regress forecast errors on forecast revisions, which asks whether the forecasting function is the optimal predictor, whatever any decision maker does with it. Applied to survey data the approach finds violations in both directions: rigidity in Coibion and Gorodnichenko (2015), overreaction in Bordalo et al. (2020). Worst-case forecasts generate both (Section 5.3). A rejection of f_2 says nothing about f_1 . A decision maker who optimizes perfectly against a badly forecasting f_2 is rational in the sense this paper uses the word. The $\sigma_i < 0$ consumers in our model are rational in that sense. In the FIRE sense they do not “have rational expectations.”²²

Econometrician B below elicits discount factors. Remark 7 states what that requires.

²²Lucas (1987, p. 13): “The term ‘rational expectations,’ as Muth used it, refers to a consistency axiom for economic models, so it can be given precise meaning only in the context of specific models. I think this is why attempts to define rational expectations in a model-free way tend to come out either vacuous (‘People do the best they can with the information they have’) or silly (‘People know the true structure of the world they live in’).” FIRE seems to emerge from a model-free definition of the second kind. In our model economy every consumer’s approximating model is the model that generates the data, so Lucas’s consistency axiom holds for the baseline. Robust consumers doubt that model and guard against alternatives to it.

Remark 7 (Eliciting a discount factor). *A choice among payments that the consumer can trade in the bond market reveals R , not $\hat{\beta}_i$. Below bliss he takes the payment with the higher present value at R and uses the bond to reach the plan he prefers; with a bliss point and no free disposal the ranking is local, not global. A laboratory choice between a sum of money now and a larger sum later is such a choice whenever the subject can offset it with the bond (Cubitt and Read, 2007).²³ A marginal choice over consumption at his optimum also reveals R , because his Euler equation sets $\hat{\beta}_i \zeta_i = R^{-1}$. Eliciting $\hat{\beta}_i$ requires a protocol that imposes consumption the subject cannot undo, such as nontradable consumption or effort allocations; Augenblick et al. (2015) implement one with real effort. Such a choice reveals $\hat{\beta}_i$ times a ratio of worst-case expected marginal utilities. Recovering $\hat{\beta}_i$ requires the subject's distance from bliss and his background income risk. We call a discount factor elicited when it comes from such a protocol. Econometrician B and the test of Section 6.5 use elicited discount factors in that sense.*

Adopt the parameter values of Section 9, $\beta = 0.95$, $\eta_1 = 0.15$ and $\eta_2 = 0.30$, and let the population be distributed uniformly over the admissible range. Three econometricians study it. They do not bring the same instruments. B elicits discount factors in the sense of Remark 7, and C treats reported forecasts as worst-case conditional means (Section 6.6).

- A. A observes consumption and asset holdings for a panel of any length. A finds that a homogeneous expected-utility model with $\beta = 0.95$ rationalizes the data: consumption is a martingale and Hall's (1978) test passes. A reports no heterogeneity in preferences. Passing that test does not show that only the benchmark type is present; the heterogeneous population that generated the data passes it too.
- B. B also elicits time preferences over consumption, with a protocol of the kind Remark 7 describes. B finds discount factors dispersed over the admissible range and observes that the consumption-saving data show no trace of the dispersion: the impatient households do not draw down their wealth. B reports an anomaly and looks for a friction that suppresses the effect of impatience. B's finding rests on the link B's protocol draws between laboratory choices and $\hat{\beta}_i$.
- C. C also collects survey forecasts of households' own future consumption and reads them as the subjective beliefs that rationalize observed actions for an expected-utility maximizer (Section 6.6). C finds worst-case forecasts that drift away from bliss, which for

²³Warner and Pleeter (2001) infer discount rates from choices of more than 60,000 U.S. military personnel between a lump sum and an annuity during a downsizing program. Mean estimated nominal rates are 10.4 percent for officers and 35.4 percent for enlisted personnel. Giving enlisted personnel the officers' lump-sum payments and demographic characteristics brings their mean to 17.3 percent. A consumer in our economy who can trade the bond compares the lump sum and the annuity at R , whatever his $\hat{\beta}_i$.

households below bliss are forecasts of declines in consumption that do not materialize, and forecast errors predictable from current consumption. C finds biased reported forecasts. Calling them irrational depends on C's definition of rationality and on C's reporting model; it does not establish a failure to optimize.

The three econometricians are looking at the same economy with different information. A's homogeneous rationalization, B's anomaly, and C's biased forecasts record one economy on different instruments, each measured from a reference point its author did not know was being chosen. B's and C's elicited data distinguish rationalizations that A's quantities cannot. A measured bias is a joint statement about a population, a benchmark, and an instrument. Quantity data identify the decision rule and α^2 , and with them the locus. They identify neither the distribution of types along it nor the benchmark.

All three econometricians use subjective expected utility as their decision theory. Each attributes to a consumer one probability model, held with full confidence. A adds that the model is the approximating model and that everyone has one discount factor. B adds that discount factors elicited in the sense of Remark 7 govern saving in the field. Under subjective expected utility the forecast that rationalizes an action is a belief, so C reads a worst-case forecast as one. The consumers in this economy use a different decision theory. None fully trusts any specification, the baseline or the worst case. The worst-case model is a device for computing a robust decision rule, not a statement of a unique subjective belief (Section 6.6). An econometrician who adopted that decision theory would try to identify the approximating model, σ_i , and $\hat{\beta}_i$, not a belief. Under the maintained assumption that the approximating model generates the data, quantity data identify the law of income and the locus. Under an assumption about what respondents report, elicited discount factors and forecasts test the Euler-equation restriction of Corollary 4, and on the locus the elicited discount factor then determines σ_i .

6.5 What elicited data restrict

Two restrictions need to be kept apart. The first holds for every robust consumer, on or off the locus. By (35), $\zeta = (\hat{\beta}R)^{-1}$ at every interior $(\sigma, \hat{\beta})$, so a robust consumer's discount factor and the persistence of his worst-case forecasts satisfy an Euler-equation restriction. The second is the locus itself: $\Omega = 1 - \beta$, equivalently a consumption martingale under the approximating model, which quantity data can check.

Corollary 4 (An Euler-equation restriction). *By the last equality in (35), every consumer*

with an admissible interior solution satisfies

$$\hat{\beta}_i \zeta_i = R^{-1}. \quad (56)$$

Equivalently $\log \hat{\beta}_i + \log \zeta_i = -\log R$, so in the plane $(\log \zeta_i, \log \hat{\beta}_i)$ the population lies on a line of slope -1 whose intercept is pinned down by the observed interest rate.

Elicit an individual's discount factor, in the sense of Remark 7, and his forecasts of his own consumption one and two periods ahead. By (48),

$$\hat{\beta}_i \left(\widehat{\mathbb{E}}_t c_{t+2}^i - \widehat{\mathbb{E}}_t c_{t+1}^i \right) - R^{-1} \left(\widehat{\mathbb{E}}_t c_{t+1}^i - c_t^i \right) = 0, \quad (57)$$

a restriction whose forecast side does not involve the unobserved bliss point: only differences of reported forecasts enter, and γ cancels from them. The elicited $\hat{\beta}_i$ is a separate matter, and Remark 7 states what recovering it requires. Because $\widehat{\mathbb{E}}_t c_{t+1}^i - c_t^i = (1 - \zeta_i)(\gamma - c_t^i)$, the ratio

$$\frac{\widehat{\mathbb{E}}_t c_{t+2}^i - \widehat{\mathbb{E}}_t c_{t+1}^i}{\widehat{\mathbb{E}}_t c_{t+1}^i - c_t^i} = \zeta_i \quad (58)$$

is defined only when $\zeta_i \neq 1$ and $c_t^i \neq \gamma$. The undivided form (57) avoids dividing by nearly zero reported changes. Flat forecasts remain uninformative, and measurement error in survey responses needs econometric treatment that the undivided form does not supply. Both forms assume that survey answers report worst-case forecasts (Section 6.6).

Elicited uncertainty restricts the same parameters.

Corollary 5 (A variance restriction). *On the locus, consumer i 's worst-case conditional variance of consumption one period ahead is*

$$\widehat{\text{Var}}_t(c_{t+1}^i) = \zeta_i \alpha^2, \quad \text{so} \quad \hat{\beta}_i \frac{\widehat{\text{Var}}_t(c_{t+1}^i)}{\alpha^2} = R^{-1}, \quad (59)$$

and at horizon $j \geq 1$, $\widehat{\text{Var}}_t(c_{t+j}^i) = \alpha^2 \zeta_i \sum_{k=0}^{j-1} \zeta_i^{2k}$.

Proof. See Appendix C.12. □

Restriction (59) needs elicited uncertainty rather than elicited mean forecasts, and it is informative where (57) is not. At $c_t^i = \gamma$ reported changes in mean forecasts are flat and (58) is undefined, while the reported variance still identifies ζ_i . It also gives the covariance distortion an observable role. That distortion leaves decision rules unchanged (Section 3), so quantity data never reveal it, and elicited uncertainty does. Like (57), (59) assumes a reporting model, here for uncertainty rather than for means.

A rejection of (57) rejects the maintained optimization and reporting structure. Satisfying it does not place a consumer on the locus, because it holds off the locus too. The locus adds the martingale restriction on consumption. A consumer whose consumption is a martingale and whose elicited discount factor and forecasts satisfy (57) is on the locus, and his elicited $\hat{\beta}_i$ then locates him there, because (39) maps $\hat{\beta}$ to σ one for one. Testing the Euler-equation restriction uses no data on savings or asset holdings. Checking the locus uses consumption data as well.

Manski (2004) surveys the measurement of subjective expectations. Bhandari et al. (2025) use survey data to identify belief distortions of the kind that appear here. Bachmann et al. (2026) use survey data on firms to show that subjective uncertainty about future sales growth differs from statistical measures of uncertainty such as volatility. Fermand et al. (2024) find that households' uncertainty about inflation and home prices extrapolates from recent experience. In our economy the worst-case variance $\zeta_i \alpha^2$ does not respond to experience. Brunnermeier and Parker (2005) study a different theory in which beliefs are chosen and the same identification question arises.

6.6 Readings of σ_i

The minimization that implements multiplier preferences does not identify the consumer's beliefs. Under the recursive construction the worst-case model is a device for computing decision rules that perform well when the consumer does not fully trust the baseline model. Under the construction of Section 5.2 it is the subjective belief that would rationalize the same plan for an expected-utility maximizer. What such a consumer would report in a survey is a separate question: a baseline forecast, a worst-case forecast, or something else. The survey literature seldom measures how much confidence a respondent places in a baseline model. The test of Section 6.5 and econometrician C of Section 6.4 presume that respondents report worst-case forecasts.

The parameter σ_i admits four readings drawn from different literatures. It can be read as a behavioral bias, pessimism. It can be read as a response to ambiguity: when an entropy budget replaces the penalty θ , multiplier preferences are related to the constraint preferences of Hansen and Sargent (2001), an instance of the max-min expected utility of Gilboa and Schmeidler (1989), by a Lagrangian duality that holds state by state rather than globally.²⁴ It can be read as a preference for robustness (Hansen and Sargent, 2001, 2008). It can be

²⁴Along the equilibrium plan the worst-case conditional relative entropy is $\frac{1}{2}\tilde{v}_{t+1}^2 = \frac{1}{2}(\zeta - 1)^2\tilde{a}_t^2/s^2$ by (36), which grows with the consumer's distance from bliss. A consumer who carries the same entropy budget at every date therefore has an implied multiplier that moves with $\tilde{a}_t$, so his rule is not $\Omega = 1 - \beta$ at every state and he is not a point on the locus. For a given state, and for a date-zero budget evaluated at a given initial condition, a θ that implements that budget exists.

read as risk sensitivity (Jacobson, 1973, Whittle, 1981, Hansen and Sargent, 1995). The decision rules, the locus, and the equilibrium are the same under the first, the third, and the fourth, which name one preference ordering three ways.

The readings differ in what restricts σ_i . The equivalence of Remark 2 links two decision theories. It does not make their interpretations interchangeable. Under risk sensitivity σ is a taste parameter calibrated to attitudes toward risk. Under robustness σ measures how much a decision maker distrusts a specification, and the detection error probabilities of Section 7 restrict it. We adopt the robustness interpretation.

The robustness reading applies outside this economy. Published forecasts of a policy institution differ systematically from its technical staff's baseline, with higher inflation and lower unemployment. Ellison and Sargent (2012) read the published forecasts as coming from a worst-case model, and the institution's policy as the best response to them, a decision rule robust to misspecification of the baseline model.

7 Detection error probabilities

The second restriction of Section 6.2 is statistical. Corollary 2 shows that quantity data never identify a consumer's type. We measure how far apart two types' models are by how hard it is to tell them apart from a time series of length T on marginal utility. Section 7.1 compares every type to the trusting benchmark. Section 7.2 compares any two types directly.

7.1 Against the benchmark type

The measure calibrates σ in units that do not depend on how consumption is measured. Section 3 noted that the numerical value of σ changes when consumption is remeasured. Detection error probabilities, introduced by Anderson et al. (2003) and used by Barillas et al. (2009), do not.

The minimizing agent's distortion affects the state only through its effect on the innovation to marginal utility (Proposition 1), so we pose the detection problem in terms of marginal utility. An econometrician with a sample of length T must decide which of the two laws

$$(\text{approximating}) \quad \mu_{t+1} = \mu_t + \alpha \tilde{w}_{t+1}, \quad (\text{worst case}) \quad \mu_{t+1} = \zeta \mu_t + \zeta^{1/2} \alpha \tilde{w}_{t+1} \quad (60)$$

generated the data, using a likelihood ratio with equal prior weights. Here $\tilde{w}_{t+1}$ is the standardized innovation to marginal utility of whichever of the two laws is in force: it is

standard normal under that law, and under the approximating model it equals $-\epsilon_{t+1}$. Along the locus the worst-case law moves the conditional mean of μ_{t+1} from μ_t to $\zeta\mu_t$ and multiplies its conditional variance by $\tilde{\Sigma} = \zeta$ (Proposition 3). Along the equilibrium path μ_t is a function of the initial condition and past shocks, so (60) is the representation of Section 5.2: two laws of the exogenous endowment history. The detection error probability is

$$\text{DEP}(\sigma; T) = \frac{1}{2} \left[\Pr\{\text{choose approximating} \mid \text{worst case true}\} + \Pr\{\text{choose worst case} \mid \text{approximating true}\} \right]. \quad (61)$$

This experiment compares two fully specified laws under equal prior weights and equal costs of the two mistakes. A detection error probability is neither a significance level nor a posterior probability that either model is true. The econometrician forms the log-likelihood ratio

$$S = \frac{1}{2} \sum_{t=0}^{T-1} \left[\left(\frac{\mu_{t+1} - \mu_t}{\alpha} \right)^2 - \frac{1}{\zeta} \left(\frac{\mu_{t+1} - \zeta\mu_t}{\alpha} \right)^2 - \log \zeta \right] \quad (62)$$

and chooses the worst-case law when $S > 0$.²⁵ The term $-\log \zeta$ in (62) accounts for the variance distortion; a statistic without it would need a threshold larger by $\frac{T}{2} \log \zeta$.

The detection problem (60) is posed for the scalar μ , while the minimizing agent distorts the law of a vector endowment. An econometrician who observes the whole vector z_t can do no better, because both the mean and the covariance distortions lie in the direction $h^\top$ of shock space (Lemma 1 in Appendix C.5).

Both laws in (60) are non-stationary, so the distribution of the likelihood ratio depends on where μ starts.

Proposition 6. *DEP($\sigma; T$) depends on the parameters only through ζ , the sample length T , and the scale-free initial condition μ_0/α . In particular it is invariant to a rescaling of the units in which consumption is measured.*

Proof. See Appendix C. □

Since ζ is a monotone function of σ by (40), Proposition 6 lets us report a concern about misspecification in a way that does not depend on units. Following Hansen et al. (1999)

²⁵A Bayesian who puts positive prior probabilities on the two laws in (60), and assigns one loss to choosing the approximating law when the worst case is true and another to choosing the worst case when the approximating law is true, chooses the worst-case law when S exceeds a threshold set by the prior probabilities and the losses. With equal prior probabilities and equal losses the threshold is zero. The threshold matters less as T grows, but not because S/T converges. μ_t is not stationary, and by (63) the variance distortion contributes information of order T while the mean distortion contributes information of order T^2 .

and Barillas et al. (2009), we regard values of σ with $\text{DEP} \geq 0.2$ as ones against which a consumer might wish to guard, on the ground that the two models are then difficult to distinguish under this criterion. That threshold is a calibration convention, not a theorem of statistical plausibility: with equal prior weights it means an average probability of correct classification of 0.8.

The two laws in (60) differ in their conditional means by $(\zeta - 1)\mu_t$ and in their conditional variances by the factor ζ . When the approximating model generates a sample of length T , the expected log-likelihood ratio of the approximating law against the worst-case law is

$$\frac{T}{2} (\zeta^{-1} - 1 + \log \zeta) + \frac{(\zeta - 1)^2}{2\zeta} \frac{1}{\alpha^2} \sum_{t=0}^{T-1} \mathbb{E}\mu_t^2. \quad (63)$$

The first term comes from the variance distortion and does not depend on μ_0 . It is small: at the breakdown point and $T = 40$ it is 2.5 percent of the total. Under the approximating model $\mathbb{E}\mu_t^2 = \mu_0^2 + t\alpha^2$, so

$$\frac{1}{\alpha^2} \sum_{t=0}^{T-1} \mathbb{E}\mu_t^2 = \frac{T(T-1)}{2} \left[1 + \frac{2}{T-1} \left(\frac{\mu_0}{\alpha} \right)^2 \right]. \quad (64)$$

The bracket equals one at $\mu_0 = 0$ and grows without bound in $|\mu_0|/\alpha$. At $T = 40$ it doubles at $|\mu_0|/\alpha = 4.4$.

We set $\mu_0 = 0$ throughout, so that the two models are compared starting from a common consumption level equal to bliss. That choice minimizes the discriminating information (64), so the detection error probabilities we report are upper bounds. It also places half of a cohort above bliss, where the worst case is a fear of being showered with goods rather than of privation. Section 9 reports both magnitudes.

$\text{DEP}(\sigma; T)$ is computed from one consumer's time series of marginal utility, testing the approximating law against one candidate worst-case law ζ . Pooling N consumers' marginal utility under that same pair of hypotheses adds their log-likelihood ratios and lowers the detection error probability, with no need to know anyone's type: the two hypotheses are laws for the whole panel's data, not claims about who shares what type. $\text{DEP}(\sigma; T)$ is therefore an upper bound on the detection error probability such a panel would attain, not the panel statistic itself. Reading the reported figure this way is close to how Section 6.2 uses a detection error probability as a restriction on how much bias may be attributed to a population: the reported figure delivers a weaker restriction than the panel statistic would, and the gap widens as the panel grows.

That pooled test targets one candidate ζ for the whole panel, which need not be any

individual consumer's own worst-case model: each type on the locus has its own ζ_i . Testing each consumer's own worst-case model instead – the hypothesis that consumer i 's endowment follows his own worst-case law, for every i at once – requires knowing which consumers share which type, information quantity data do not reveal (Corollary 2). That hypothesis is special in a second way as well: Section 6.1 shows those heterogeneous own-type worst-case laws cannot jointly generate the data except when $\gamma = Y$, since when types are independent of consumption and a positive mass of consumers is robust, the aggregate distortion (53) averages to zero only in that case, and otherwise aggregate income would not stay at Y (Section 2.2). The single-consumer statistic of (61) is the one this paper reports: it measures what one consumer's own record could tell him, the question Section 7.1 poses, without presuming an econometrician who already knows the type distribution.

7.2 Between two types

Section 7.1 put the benchmark type $\sigma = 0$ on one side of every comparison. We compare two arbitrary types on the locus, privileging neither as the truth.

For i and j on the locus, let $\text{DEP}(\zeta_i, \zeta_j; T)$ be defined by (61) with $\mu_{t+1} = \zeta_i \mu_t + \zeta_i^{1/2} \alpha \tilde{w}_{t+1}$ and $\mu_{t+1} = \zeta_j \mu_t + \zeta_j^{1/2} \alpha \tilde{w}_{t+1}$ as the two hypotheses. Section 9 reports Monte Carlo estimates $\widehat{\text{DEP}}(\zeta_i, \zeta_j; T)$ over a grid of pairs. A companion measure is the conditional relative entropy of j 's model with respect to i 's, which for these two laws is

$$\frac{1}{2} \left[\frac{\zeta_j}{\zeta_i} - 1 - \log \frac{\zeta_j}{\zeta_i} + \frac{(\zeta_i - \zeta_j)^2}{\zeta_i \alpha^2} \mu_t^2 \right]. \quad (65)$$

Expression (65) involves only the two types compared and has no absolute origin. It is not symmetric in i and j , so it is a discrepancy and not a distance; the detection error probability, which is symmetric, is the measure we report. Section 6.1 found no privileged reference type for measuring biases. This economy likewise distinguishes no model against which to measure discrepancies. It supplies only discrepancies between models.

8 What a bias costs

Consumer i 's plan is rule (8), since $(\sigma_i, \hat{\beta}_i)$ lies on the locus (39) by construction. Measure the cost of departing from a reference criterion by comparing the value of this fixed plan with the value of the plan that the criterion would itself choose. Against the benchmark $(0, \beta)$ criterion the cost is zero, because rule (8) is the benchmark consumer's own optimal plan. Against consumer i 's own $(\sigma_i, \hat{\beta}_i)$ criterion the cost is zero as well, because $\Omega = 1 - \beta$

solves consumer i 's own problem (Proposition 3). Of these three criteria, only that of a trusting consumer who discounts at $\hat{\beta}_i$ makes the plan look costly, at a rate Λ_i computed below; criteria off the locus can assign positive losses too. Quantity data give no way to choose among the three. Section 8.2 computes the cost under the twin's criterion and then adds the reverse comparison under consumer i 's own worst-case model.

8.1 A pessimist and an impatient twin

Define a hypothetical consumer i^{twin} who trusts the approximating model but has discount factor $\hat{\beta}_i$, facing the same interest rate $R = \beta^{-1}$. This twin never appears in the equilibrium; he is a device for interpreting consumer i 's behavior. His admissible plans are those with $\mathbb{E}_0 \sum_t \hat{\beta}_i^t \tilde{a}_t^2 < +\infty$ under the approximating model, square-summability at his own discount factor, as the robust consumer's are at his. Setting $\sigma = 0$ and $\hat{\beta} = \hat{\beta}_i$ in (34), the twin's marginal propensity to consume is

$$\Omega_i^{\text{twin}} = 1 - \frac{1}{\hat{\beta}_i R^2} = 1 - \beta \zeta_i. \quad (66)$$

By (35), the twin's net wealth and marginal utility grow under the approximating model at rate $R(1 - \Omega_i^{\text{twin}}) = (\hat{\beta}_i R)^{-1} = \zeta_i$. Because $\hat{\beta}_i R < 1$, the twin finds saving unattractive: he leans toward consuming now, and he plans a consumption path that drifts away from bliss at rate ζ_i .

Consumer i 's plan (8) is a martingale under the approximating model and drifts away from bliss at rate ζ_i under his own worst-case model. The twin plans the drift that consumer i fears. One rate ζ_i governs the twin's plan and consumer i 's worst-case forecast of his own consumption.

Define normalized robustness

$$\tau_i \equiv \frac{\sigma_i}{\underline{\sigma}} \in [0, 1), \quad \zeta_i = \frac{1}{1 - (1 - \beta)\tau_i}. \quad (67)$$

The twin's optimal plan satisfies the financial solvency restriction (3) only when $\tau_i < 1/(1 + \sqrt{\beta})$, which at $\beta = 0.95$ is $\tau_i < 0.506$. Beyond that threshold his value is approached, but not attained, by plans that satisfy (3).²⁶ The twin's problem has an interior solution only

²⁶The twin's problem does not impose (3). His optimal plan satisfies the present-value condition at the market discount factor, $\beta^T \mathbb{E}_0 \tilde{a}_T \rightarrow 0$, and it satisfies (3) only when $\zeta_i < \beta^{-1/2}$, equivalently when $\tau_i < 1/(1 + \sqrt{\beta})$. Beyond that threshold, follow the twin's rule for N periods and switch to (8) thereafter. Net wealth is a random walk after the switch, so each such plan satisfies (3). The difference between its value and the twin's is a discounted continuation value at date N , bounded by a constant times $\hat{\beta}_i^N (1 + \mathbb{E}_0 \tilde{a}_N^2)$, which vanishes because $\hat{\beta}_i \zeta_i^2 = \beta \zeta_i < 1$. The twin's value is therefore attained in his own admissible set and

when $\hat{\beta}_i R^2 > 1$, which on the locus is $\sigma_i > \underline{\sigma}$ (Proposition 9). As $\sigma_i \downarrow \underline{\sigma}$ the twin's marginal propensity to consume approaches zero. Statements below about that endpoint refer to limits, not to an admissible type.

8.2 The cost of a departure

Proposition 7 (The cost of consumer i 's plan under the twin's criterion). *Evaluate consumer i 's plan (8) and the twin's optimal plan (66) by the twin's criterion, and call the two values V_i^{act} and V_i^{opt} . Then*

$$\Lambda_i \equiv 1 - \frac{V_i^{opt}}{V_i^{act}} = \frac{\beta(\zeta_i - 1)^2}{\zeta_i(1 - \beta)^2} = \frac{\beta\tau_i^2}{1 - (1 - \beta)\tau_i} \in [0, 1). \quad (68)$$

Λ_i does not depend on the initial condition or on the bliss point, and it depends on σ_i and the endowment process only through τ_i . At fixed σ_i , a change in α^2 changes τ_i and therefore Λ_i . Λ_i is zero at $\tau_i = 0$ and approaches one as $\tau_i \uparrow 1$.

Proof. See Appendix C. □

Both values are shortfalls from bliss, which the twin's criterion sets at zero, so Λ_i is the fraction of the bliss-normalized utility shortfall of consumer i 's plan that re-optimizing would remove. Measured against V_i^{opt} the departure costs $\Lambda_i/(1 - \Lambda_i)$, which exceeds one once Λ_i passes a half. As $\tau_i \uparrow 1$, $\Lambda_i \rightarrow 1$, because the twin's optimal plan approaches the bliss plan, whose value is zero.

Λ_i is scale free, and it has a welfare-equivalent form.²⁷ Under the twin's criterion, contracting consumer i 's gap from bliss to

$$c_t^{adj} = \gamma - \sqrt{1 - \Lambda_i} (\gamma - c_t) \quad (69)$$

gives the value $(1 - \Lambda_i)V_i^{act} = V_i^{opt}$. Thus $1 - \sqrt{1 - \Lambda_i}$ is the proportional contraction of the consumption gap from bliss that is worth as much as re-optimizing. It is not a percentage increase in consumption, and the adjusted process need not be feasible with the same initial wealth and trading opportunities. Section 9 reports Λ_i for our parameter values.

Proposition 7 prices both plans by the twin's criterion. The reverse comparison prices both plans under consumer i 's worst-case model.

approached, but not attained, in the financial domain. The criterion is strictly concave, so no other plan attains it.

²⁷Cochrane (1989) prices departures from an optimal consumption rule in dollars per quarter and finds them small. Λ_i is a fraction of a shortfall from bliss, and converting it to dollars would require a scale for that shortfall.

Proposition 8 (The twin’s plan under consumer i ’s worst-case model). *Evaluate consumer i ’s plan (8) and the twin’s plan (66) under consumer i ’s worst-case model of Section 5.2, discounting at $\hat{\beta}_i$ and imposing no entropy penalty, and call the two values $\tilde{V}_i^{own}$ and $\tilde{V}_i^{twin}$. Then*

$$\tilde{\Lambda}_i \equiv 1 - \frac{\tilde{V}_i^{own}}{\tilde{V}_i^{twin}} = \frac{\Lambda_i}{1 + \Lambda_i} \in [0, \tfrac{1}{2}]. \quad (70)$$

$\tilde{\Lambda}_i$ does not depend on the initial condition or on the bliss point, and it depends on σ_i and the endowment process only through τ_i . It is zero at $\tau_i = 0$ and approaches one half as $\tau_i \uparrow 1$. Both $\tilde{\Lambda}_i$ and Λ_i are re-optimization losses.

Proof. See Appendix C.10. □

Λ_i is the fraction of consumer i ’s shortfall, priced by the twin’s criterion, that re-optimizing would remove. $\tilde{\Lambda}_i$ is the fraction of the twin’s shortfall, priced under consumer i ’s worst-case model, that adopting consumer i ’s plan would remove. By (70) the two comparisons are one number seen two ways, and they are asymmetric: $\tilde{\Lambda}_i < \Lambda_i$, with limits one half and one as $\tau_i \uparrow 1$. Consumer i ’s own multiplier criterion, which re-minimizes over distortions for whatever plan it prices, assigns the twin’s plan no finite value once τ_i exceeds about one quarter.²⁸ That divergence is not a general ordering of the two criteria. For a plan and an alternative law $\mathbb{Q}$, write

$$\mathcal{H}(\mathbb{Q}) \equiv \mathbb{E}_0^{\mathbb{Q}} \sum_{t=0}^{\infty} \hat{\beta}^{t+1} \mathcal{E}(\tilde{v}_{t+1}, \tilde{\Sigma}_{t+1}) \geq 0 \quad (71)$$

for the discounted relative entropy of $\mathbb{Q}$ with respect to the approximating model, with $\mathcal{E}$ from (29), so that the multiplier criterion is $V^{\text{mult}} = \inf_{\mathbb{Q}} \{ \mathbb{E}_0^{\mathbb{Q}} \sum_t \hat{\beta}^t u(c_t) + \theta \mathcal{H}(\mathbb{Q}) \}$. The multiplier value of a plan therefore equals its value under the minimizing law plus the entropy penalty, $V^{\text{mult}} = \tilde{V} + \theta \mathcal{H} \geq \tilde{V}$: on plans where both are finite the multiplier criterion assigns a *higher* value than the same worst-case law does with the penalty removed.

8.3 Pooling

Ljungqvist and Sargent (2026, section 6.5) observe that the equilibrium of Section 2.2 is inefficient. Before initial conditions are drawn, with the initial distribution of Proposition 4

²⁸Under that criterion the fixed-policy Riccati equation for the twin’s plan has no real solution once τ_i exceeds 0.2548 at $\beta = 0.95$. The twin lets net wealth drift at rate ζ_i , and a criterion that prices the worst case punishes drift without bound. That threshold depends on β alone: it is 0.2600 at $\beta = 0.90$ and 0.2509 at $\beta = 0.99$. It is an algebraic property of that fixed-policy equation, not a substitute for the admissibility argument of Appendix A.

and types independent of initial conditions, c_0^i has mean Y within every type. Pooling non-financial incomes and giving every consumer Y at every date is then feasible.²⁹ Every type prefers pooling under the approximating model, because the bond economy adds $\text{Var}(c_0^i)$ and future risk to the expected squared gap from bliss. The pooled allocation carries no risk, so the minimizing agent leaves it undistorted and it bears no entropy penalty. Every consumer therefore also ranks it above the bond economy's allocation under his multiplier criterion. Under consumer i 's own criterion the value of the bond economy's allocation includes the constant $\mathbf{p}_s$ of (82); the covariance distortion enlarges $\mathbf{p}_s$ and widens that preference. The ranking holds whatever the distribution Φ of types. The allocations are identical across the two economies of Corollary 2, and so is the ranking. The size of each consumer's gain from pooling differs, because it depends on his $(\sigma_i, \hat{\beta}_i)$. Equal pooling is a dominance benchmark and not the Pareto frontier. Characterizing the complete-markets allocation with heterogeneous multiplier preferences is a separate problem, which we do not take up.³⁰

9 How much is at stake

Section 6 bounds the freedom to choose a reference model and a distribution of types, a freedom that quantity data alone do not restrict. Here are the magnitudes.

Set $\beta = 0.95$, $\eta_1 = 0.15$, and $\eta_2 = 0.30$, the parameter values of Sargent and Stachurski (2026b), taking one period to be a year. Then $\alpha^2 = 0.022725$, of which the permanent shock contributes 99.0 percent. The breakdown point is $\underline{\sigma} = -0.110011$, at which $\hat{\beta} = \beta^2 = 0.9025$ and $\zeta = R = 1.052632$. By Corollary 1 the admissible discount factors are exactly $(0.9025, 0.95]$, corresponding to rates of time preference running from 5.3 percent up to, but not including, 10.8 percent per year. Figure 1 plots the locus at these values.

²⁹We do not claim that competitive complete markets implement equal pooling, and with heterogeneous discount factors they do not: constant deterministic consumption below bliss requires the one-period price to equal $\hat{\beta}_i$ for every unconstrained type, and one price cannot equal distinct discount factors. Complete markets opened after the draws insure subsequent shocks and leave initial wealth differences in place. In the homogeneous-discount benchmark of Sargent and Stachurski (2026a) they implement $c_t^i = c_0^i$, not $c_t^i = Y$. After the draws a benchmark consumer's bond-economy consumption is a random walk from c_0^i , and he prefers it to Y forever if and only if $(\gamma - c_0^i)^2 + \beta\alpha^2/(1 - \beta) < (\gamma - Y)^2$. Being nearer to bliss than Y at date zero is not enough, because future consumption risk also counts.

³⁰A planner must also choose whose worst case to use. Because the shocks are idiosyncratic, he can minimize agent by agent, which gives each consumer his own worst-case model and makes the Pareto frontier depend on the distribution of σ_i as well as on the allocation. We record the issue and do not settle it. Heterogeneous discount factors do not by themselves open gains from trading dates. By Corollary 4 every type on the locus values date- $(t + 1)$ goods at R^{-1} at the margin, so a deterministic transfer across dates adds to one group's net wealth exactly the present value at R that it takes from another's, and no such reallocation makes both groups better off. The transfers are deterministic, so the minimizing agent leaves them undistorted and they carry no entropy penalty; the gains from a planner must come from insuring shocks rather than from shifting dates.

Table 1 reports four admissible types spaced evenly across the admissible range, and the boundary limit.

τ_i	σ_i	$\hat{\beta}_i$	ζ_i	$\widehat{\text{DEP}}$			doubling horizon (years)	Λ_i	$\tilde{\Lambda}_i$
				$T=10$	$T=40$	$T=160$			
0.00	0.0000	0.9500	1.0000	0.500	0.500	0.500	∞	0.000	0.000
0.25	-0.0275	0.9381	1.0127	0.484	0.434	0.204	55.1	0.060	0.057
0.50	-0.0550	0.9262	1.0256	0.468	0.358	0.039	27.4	0.244	0.196
0.75	-0.0825	0.9144	1.0390	0.451	0.276	0.006	18.1	0.555	0.357
1.00 ⁻	-0.1100	0.9025	1.0526	0.433	0.199	0.001	13.5	1.000	0.500

Table 1: Four admissible types on the locus (39), evenly spaced from the benchmark to the breakdown point, and the boundary limit (last row), at the parameter values of Section 9. Detection error probabilities compare each type with the benchmark and start at $\mu_0 = 0$ (Section 7.1). The doubling horizon is $\log 2 / \log \zeta_i$. Λ_i and $\tilde{\Lambda}_i$ are (68) and (70), each a fraction of a shortfall from bliss rather than of consumption; in the last two rows both price a twin’s plan that violates (3) (Section 8.1).

As the robustness boundary is approached, discount factors fall by up to 0.0475 and the worst-case forecast of the gap between bliss and consumption doubles in thirteen and a half periods; the benchmark consumer’s forecast of the gap never changes. Priced by the twin’s criterion, that difference in views costs Λ_i . At a quarter of the way to the breakdown point the twin could remove 6.0 percent of consumer i ’s shortfall from bliss by re-optimizing; at three quarters of the way, 55.5 percent. Priced under consumer i ’s worst-case model the comparison runs the other way and is smaller: the twin’s plan wastes 5.7 percent of the twin’s shortfall a quarter of the way out and 35.7 percent three quarters of the way out.

The two constraints that bound σ , the breakdown point of Appendix A and the statistical detectability of Section 7, interact. At $T = 10$ the detection error probability stays above 0.43 on the admissible range. At $T = 40$ the smallest value on our grid, approached at the breakdown point, is 0.199, near the 0.2 threshold; with a Monte Carlo standard error of about 0.001 the value of σ at which the $T = 40$ curve crosses 0.2 is not sharply determined.³¹ For a consumer at bliss and samples shorter than about forty years, every concern about misspecification that the model can represent is difficult to distinguish from the approximating model under this criterion, and the breakdown point is the binding constraint. For longer samples detectability binds first. At $T = 160$ the detection error probability reaches 0.2 at $\sigma = -0.0278$, about a quarter of the way to the breakdown point ($\sigma/\underline{\sigma} = 0.253$). Figure 2

³¹A finer 21-point grid gives the same minima on that grid at $T = 10$ and $T = 40$. A worst-case model that distorts only the conditional mean gives 0.435 and 0.200.

draws the three curves.

The figures above set $\mu_0 = 0$, the initialization of Section 7.1. Table 2 reports what that choice costs and what it buys.

$ \mu_0 /\alpha$	0	4	8
$\widehat{\text{DEP}}$ at $T = 40$, $\sigma = 0.5 \underline{\sigma}$	0.358	0.302	0.193
$\Pr\{c_{40}^i > \gamma\}$ under the approximating model	0.500	0.264	0.103

Table 2: Sensitivity of the detection error probability to the initial distance from bliss, at the parameter values of this section and $\sigma = 0.5 \underline{\sigma}$. The second row reports the fraction of a cohort above bliss at $t = 40$ under the approximating model.

The first row confirms that the detection error probability falls as $|\mu_0|/\alpha$ rises, so the figures we report are upper bounds in every case we computed. The second row records what the near-bliss initialization costs in interpretation. Under the approximating model $\Pr\{c_t^i > \gamma\} = 1 - \mathcal{N}((\gamma - c_0^i)/(\alpha\sqrt{t}))$, where $\mathcal{N}$ is the standard normal cumulative distribution function, which is one half at every $t \geq 1$ when $c_0^i = \gamma$. The initialization that makes the two models hardest to tell apart is also the one at which half the cohort sits above bliss. The equilibrium of Proposition 4 gives consumers dispersed $\mu_0^i = \gamma - c_0^i$, and $|\mu_0^i|/\alpha$ is large whenever the bliss point sits well above mean consumption. An econometrician who studies a consumer far from bliss separates the two models faster than our figures say.

Table 3 drops the benchmark as a reference point and reports detection error probabilities between pairs of types.

The smallest values on our grid are approached at $\{0, \underline{\sigma}\}$: 0.199 at $T = 40$ and 0.001 at $T = 160$. For consumers at bliss and forty years of data, every pair of admissible types on the grid is difficult to distinguish under this criterion. At one hundred sixty years the only adjacent pair on the grid above the 0.2 threshold is the one containing the benchmark type.

Equal spacing in σ does not give equal detectability. The mean-distortion part of the discriminating information in (65) is proportional to $(\zeta_i - \zeta_j)^2 \sum_t \mu_t^2 / \alpha^2$, and with $\mu_0 = 0$ that sum grows like $T^2/2$ under a unit root and geometrically under an explosive one. Pairs farther out on the locus are therefore easier to tell apart: at $T = 160$ the five adjacent pairs in Table 3 run from 0.264 down to 0.003. Among equally spaced adjacent pairs on this grid, the hardest to distinguish is the one containing the benchmark type. In a continuous type space, arbitrarily close types near any point have detection error probabilities near one half, so no pair is hardest without a specified separation.

Priced by the benchmark criterion or by the consumer's own, a departure costs nothing. Priced by the twin's criterion, it costs from zero to 55.5 percent of the shortfall from bliss




Figure 2: Detection error probabilities (61) against the benchmark type for $T = 10, 40$, and 160 , at the parameter values of Table 1 and $\mu_0 = 0$. The horizontal line is the 0.2 threshold of Section 7.1. Each point is a Monte Carlo average over $100,000$ paths, with a standard error of about 0.001 , so the location of the $T = 160$ crossing is determined only to within a small interval.

across the admissible types of Table 1, and the cost approaches the whole shortfall at the boundary.

10 Concluding remarks

We have constructed a class of heterogeneous-preference, linear-quadratic incomplete-markets economies in which an equilibrium can be verified explicitly without solving a distributional fixed-point problem. A robust consumer's problem reduces to a scalar Riccati equation. Its solutions trace a locus along which every type chooses the marginal propensity to consume $1 - \beta$. Holding fixed the initial distribution and the law that generates endowments, the equilibrium interest rate and the joint distribution of consumption, debt, and income do not depend on the distribution of types across that locus. The equilibrium is nonstationary, and an imposed compensation between robustness and impatience makes the construction work. A constant death rate with fair annuities restores a stationary cross-section without disturbing any identification result (Remark 5 and Appendix B).

Even when the law generating realized income is known, consumption and saving choices do not separately identify patience and concern about model misspecification. Behavioral and welfare interpretations therefore require assumptions beyond the observed decision rule.

<i>Panel A: T = 40</i>					
$\tau_i \setminus \tau_j$	0.2	0.4	0.6	0.8	1.0 ⁻
0.0	0.448	0.390	0.326	0.260	0.199
0.2		0.439	0.370	0.297	0.226
0.4			0.428	0.345	0.262
0.6				0.411	0.314
0.8					0.390

<i>Panel B: T = 160</i>					
$\tau_i \setminus \tau_j$	0.2	0.4	0.6	0.8	1.0 ⁻
0.0	0.264	0.078	0.018	0.004	0.001
0.2		0.129	0.025	0.005	0.001
0.4			0.041	0.006	0.001
0.6				0.011	0.002
0.8					0.003

Table 3: Detection error probabilities $\widehat{\text{DEP}}(\zeta_i, \zeta_j; T)$ between pairs of types on the locus, at the parameter values of Table 1 and $\mu_0 = 0$, with $\tau_i = \sigma_i/\underline{\sigma}$. The detection error probability is symmetric in the pair, so each panel reports the distinct pairs once. The column $\tau_j = 1^-$ reports limits as $\sigma_j \downarrow \underline{\sigma}$ rather than an admissible type.

That is a statement about preference identification under one probability law. Statistical discrimination between the approximating law and a worst-case law is a different question, and data become more informative about it as samples grow.

Calling a behavior a bias requires a reference, and quantity data leave the analyst free to choose one. The benchmark consumer's criterion puts the cost of consumer i 's plan at zero. Consumer i 's own criterion puts it at zero as well, since he is optimizing. The criterion of a trusting consumer with discount factor $\hat{\beta}_i$ puts it at a fraction Λ_i of consumer i 's shortfall from bliss. Consumer i 's worst-case model prices the reverse comparison at $\Lambda_i/(1 + \Lambda_i)$. Section 9 reports magnitudes. The same freedom generates the three readings of one economy in Section 6.4 and the readings of σ_i in Section 6.6.

The model bounds that freedom. Admissible discount factors lie in $(\beta^2, \beta]$, so no consumer discounts one year more heavily than the benchmark discounts two. As the boundary is approached, the impatient twin of Section 8.1 approaches the bliss plan. In our parameterization, for a consumer at bliss, statistical detectability binds before the breakdown point in samples longer than about forty years. For a consumer far from bliss fewer observations separate his worst-case model from the approximating model.

Every robust consumer satisfies an Euler-equation restriction that links his discount factor to the persistence of his worst-case forecasts. Under an explicit assumption about reported

forecasts, elicited discount factors and forecasts can test it without data on savings or asset holdings. Satisfying it does not place a consumer on the locus; the locus adds the martingale restriction on consumption. Holding preferences and the interest rate fixed, a change in income risk breaks observational equivalence for robust consumers, so interventions offer another route to identification (Section 6.3).

Remark 3 records two limitations: every consumer is placed on the locus by construction, and we establish existence of an equilibrium at $R = \beta^{-1}$, not uniqueness. We record two more. First, the model exploits a linear-quadratic structure. Observational equivalence along (39) rests on it: the decision rule is linear in net wealth, and the locus makes its coefficient common across types. Certainty equivalence in the sense of Simon (1956) and Theil (1957) does not survive a concern about misspecification (Section 5.1); what survives on the locus is its appearance. Our exact results need not survive nonquadratic preferences. The offset between a precautionary motive and impatience is more general. Second, the non-identification of Corollary 2 applies to quantities in an economy with one traded asset. Observationally equivalent $(\sigma_i, \hat{\beta}_i)$ pairs imply different agent-specific shadow valuations of risky payoffs (Hansen et al., 1999, Barillas et al., 2009). Whether an expanded asset-market equilibrium would identify the type distribution is a conjecture, not a result. The shocks here are idiosyncratic, so a common traded payoff need not load on a particular household's consumption innovation; additional securities change portfolios and possibly the policies that support the equivalence; and marginal utility $\gamma - c_t$ can change sign, so the usual pricing kernel need not be positive. An expanded model would have to specify payoffs, participation, net supplies, admissibility, and prices before making an identification claim. Tallarini (2000)'s risk-sensitivity parameter shows the asymmetry between quantities and prices in a representative agent economy. It is nearly invisible in quantities and large in its effect on asset prices. Alvarez and Jermann (2004) recover a pricing kernel from asset returns rather than from an assumed utility function, and find the welfare cost of consumption fluctuations substantially larger than a consumption-based calculation suggests.

References

- Aiyagari, S. Rao. 1994. Uninsured Idiosyncratic Risk and Aggregate Saving. *Quarterly Journal of Economics* 109 (3):659–684.
- Alvarez, Fernando and Urban J. Jermann. 2004. Using Asset Prices to Measure the Cost of Business Cycles. *Journal of Political Economy* 112 (6):1223–1256.
- Anderson, Evan W., Lars Peter Hansen, and Thomas J. Sargent. 2003. A Quartet of Semi-

- groups for Model Specification, Robustness, Prices of Risk, and Model Detection. *Journal of the European Economic Association* 1 (1):68–123.
- Augenblick, Ned, Muriel Niederle, and Charles Sprenger. 2015. Working over Time: Dynamic Inconsistency in Real Effort Tasks. *Quarterly Journal of Economics* 130 (3):1067–1115.
- Bachmann, Rüdiger, Kai Carstensen, Stefan Lautenbacher, Manuel Menkhoff, and Martin Schneider. 2026. Uncertainty and Change: Survey Evidence of Firms’ Subjective Beliefs. *American Economic Review* 116 (8):2806–2849.
- Barillas, Francisco, Lars Peter Hansen, and Thomas J. Sargent. 2009. Doubts or Variability? *Journal of Economic Theory* 144 (6):2388–2418.
- Becker, Robert A. 1980. On the Long-Run Steady State in a Simple Dynamic Model of Equilibrium with Heterogeneous Households. *Quarterly Journal of Economics* 95 (2):375–382.
- Bewley, Truman. 1977. The Permanent Income Hypothesis: A Theoretical Formulation. *Journal of Economic Theory* 16 (2):252–292.
- . 1980. The Optimum Quantity of Money. In *Models of Monetary Economies*, edited by John H. Kareken and Neil Wallace, 169–210. Minneapolis, Minnesota: Federal Reserve Bank of Minneapolis.
- Bhandari, Anmol, Jaroslav Borovička, and Paul Ho. 2025. Survey Data and Subjective Beliefs in Business Cycle Models. *Review of Economic Studies* 92 (3):1375–1437.
- Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. Overreaction in Macroeconomic Expectations. *American Economic Review* 110 (9):2748–2782.
- Brunnermeier, Markus K. and Jonathan A. Parker. 2005. Optimal Expectations. *American Economic Review* 95 (4):1092–1118.
- Cochrane, John H. 1989. The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives. *American Economic Review* 79 (3):319–337.
- Coibion, Olivier and Yuriy Gorodnichenko. 2015. Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. *American Economic Review* 105 (8):2644–2678.
- Cubitt, Robin P. and Daniel Read. 2007. Can Intertemporal Choice Experiments Elicit Time Preferences for Consumption? *Experimental Economics* 10 (4):369–389.

- Deaton, Angus and Christina Paxson. 1994. Intertemporal Choice and Inequality. *Journal of Political Economy* 102 (3):437–467.
- Ellison, Martin and Thomas J. Sargent. 2012. A Defense of the FOMC. *International Economic Review* 53 (4):1047–1065.
- Ferland, Elyas, Camelia M. Kuhnen, Geng Li, and Itzhak Ben-David. 2024. Extrapolative Uncertainty and Household Economic Behavior. *Management Science* 70 (8):5607–5625.
- Frederick, Shane, George Loewenstein, and Ted O’Donoghue. 2002. Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature* 40 (2):351–401.
- Friedman, Milton. 1957. *A Theory of the Consumption Function*. Princeton, New Jersey: Princeton University Press.
- Gilboa, Itzhak and David Schmeidler. 1989. Maxmin Expected Utility with Non-Unique Prior. *Journal of Mathematical Economics* 18 (2):141–153.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker. 2002. Consumption over the Life Cycle. *Econometrica* 70 (1):47–89.
- Hall, Robert E. 1978. Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy* 86 (6):971–987.
- Hansen, Lars Peter and Thomas J. Sargent. 1995. Discounted Linear Exponential Quadratic Gaussian Control. *IEEE Transactions on Automatic Control* 40 (5):968–971.
- . 2001. Robust Control and Model Uncertainty. *American Economic Review* 91 (2):60–66.
- . 2005. Certainty Equivalence and Model Uncertainty. In *Models and Monetary Policy: Research in the Tradition of Dale Henderson, Richard Porter, and Peter Tinsley*, edited by Jon Faust, Athanasios Orphanides, and David Reifschneider, 17–38. Washington, D.C.: Board of Governors of the Federal Reserve System.
- . 2008. *Robustness*. Princeton, New Jersey: Princeton University Press.
- . 2022. Structured Ambiguity and Model Misspecification. *Journal of Economic Theory* 199:105165.
- . 2024. Risk, Ambiguity, and Misspecification: Decision Theory, Robust Control, and Statistics. *Journal of Applied Econometrics* 39 (6):969–999.

- Hansen, Lars Peter, Thomas J. Sargent, and Thomas D. Tallarini. 1999. Robust Permanent Income and Pricing. *Review of Economic Studies* 66 (4):873–907.
- Huggett, Mark. 1993. The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics and Control* 17 (5-6):953–969.
- Indarte, Sasha, Raymond Kluender, Ulrike Malmendier, and Michael Stepner. 2026. Consumption Wedges: Measuring and Diagnosing Distortions. NBER Working Paper 34891, National Bureau of Economic Research.
- Jacobson, David H. 1973. Optimal Stochastic Linear Systems with Exponential Performance Criteria and Their Relation to Deterministic Differential Games. *IEEE Transactions on Automatic Control* 18 (2):124–131.
- Kasa, Kenneth. 2006. Robustness and Information Processing. *Review of Economic Dynamics* 9 (1):1–33.
- Kimball, Miles S. 1990. Precautionary Saving in the Small and in the Large. *Econometrica* 58 (1):53–73.
- Krusell, Per and Anthony A. Smith. 1998. Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy* 106 (5):867–896.
- Lawrence, Emily C. 1991. Poverty and the Rate of Time Preference: Evidence from Panel Data. *Journal of Political Economy* 99 (1):54–77.
- Leland, Hayne E. 1968. Saving and Uncertainty: The Precautionary Demand for Saving. *Quarterly Journal of Economics* 82 (3):465–473.
- Ljungqvist, Lars and Thomas J. Sargent. 2026. *Recursive Macroeconomic Theory*. Cambridge, Massachusetts: MIT Press, 5 ed. Forthcoming.
- Lucas, Robert E., Jr. 1976. Econometric Policy Evaluation: A Critique. In *The Phillips Curve and Labor Markets*, vol. 1 of *Carnegie-Rochester Conference Series on Public Policy*, edited by Karl Brunner and Allan H. Meltzer, 19–46. Amsterdam: North-Holland.
- . 1987. *Models of Business Cycles*. Oxford: Basil Blackwell.
- Lucas, Robert E., Jr. and Thomas J. Sargent. 1981. *Rational Expectations and Econometric Practice*. Minneapolis, Minnesota: University of Minnesota Press.
- Luo, Yulei and Eric R. Young. 2010. Risk-Sensitive Consumption and Savings under Rational Inattention. *American Economic Journal: Macroeconomics* 2 (4):281–325.

- Maenhout, Pascal J. 2004. Robust Portfolio Rules and Asset Pricing. *Review of Financial Studies* 17 (4):951–983.
- Manski, Charles F. 2003. *Partial Identification of Probability Distributions*. Springer Series in Statistics. New York: Springer.
- . 2004. Measuring Expectations. *Econometrica* 72 (5):1329–1376.
- Nie, Jun, Yulei Luo, and Eric R. Young. 2019. Ambiguity, Low Risk-Free Rates, and Consumption Inequality. Research Working Paper RWP 15-14, Federal Reserve Bank of Kansas City. Revised January 2019.
- Rozsypal, Filip and Kathrin Schlafmann. 2023. Overpersistence Bias in Individual Income Expectations and Its Aggregate Implications. *American Economic Journal: Macroeconomics* 15 (4):331–371.
- Sargent, Thomas J. 1981. Interpreting Economic Time Series. *Journal of Political Economy* 89 (2):213–248.
- Sargent, Thomas J. and John Stachurski. 2026a. Consumption Smoothing with Incomplete and Complete Markets. QuantEcon lecture, Quantitative Economics with Python. https://python.quantecon.org/lq_bewley_complete_markets.html.
- . 2026b. A Robust LQ Bewley Model. QuantEcon lecture, Quantitative Economics with Python. https://python.quantecon.org/lq_robust_bewley.html.
- Simon, Herbert A. 1956. Dynamic Programming under Uncertainty with a Quadratic Criterion Function. *Econometrica* 24 (1):74–81.
- Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. Consumption and Risk Sharing over the Life Cycle. *Journal of Monetary Economics* 51 (3):609–633.
- Sun, Yeneng. 2006. The Exact Law of Large Numbers via Fubini Extension and Characterization of Insurable Risks. *Journal of Economic Theory* 126 (1):31–69.
- Tallarini, Thomas D. 2000. Risk-Sensitive Real Business Cycles. *Journal of Monetary Economics* 45 (3):507–532.
- Theil, Henri. 1957. A Note on Certainty Equivalence in Dynamic Planning. *Econometrica* 25 (2):346–349.
- Wang, Neng. 2003. Caballero Meets Bewley: The Permanent-Income Hypothesis in General Equilibrium. *American Economic Review* 93 (3):927–936.

- Warner, John T. and Saul Pleeter. 2001. The Personal Discount Rate: Evidence from Military Downsizing Programs. *American Economic Review* 91 (1):33–53.
- Whittle, Peter. 1981. Risk-Sensitive Linear/Quadratic/Gaussian Control. *Advances in Applied Probability* 13 (4):764–777.
- Zeldes, Stephen P. 1989. Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence. *Quarterly Journal of Economics* 104 (2):275–298.

A Admissible amount of robustness

Proposition 9 (Breakdown point). *Among rules of the form $\tilde{c} = \Omega \tilde{a}$, the robust consumer's auxiliary problem, posed on the integrability domain of standing assumption (iii), has an interior solution if and only if $\sigma > \underline{\sigma}$ and $\hat{\beta}R^2 > 1$, recalling $\underline{\sigma}$ from (38). That solution is $\tilde{c} = \Omega \tilde{a}$ with Ω given by (34), and it satisfies $\Omega \in (0, 1)$, $1 + \sigma\Omega s^2 > 0$, and $\hat{\beta}\zeta^2 < 1$. If $\hat{\beta}R^2 < 1$, the bliss plan $c_t = \gamma$ solves the auxiliary problem, though it violates (3) and is therefore not a feasible equilibrium choice. If $\hat{\beta}R^2 = 1$, neither the interior rule nor the bliss plan is integrability-admissible. Along the locus (39), $\hat{\beta}R^2 > 1$ if and only if $\sigma > \underline{\sigma}$. As $\sigma \downarrow \underline{\sigma}$ along the locus, $\zeta \rightarrow R = \beta^{-1}$ and $\hat{\beta} \rightarrow \beta^2$. The “only if” direction concerns this class of rules. The “if” direction does not: Appendix C.10 shows that the interior rule is the best response, among all plans with finite discounted second moments under the worst-case law, linear or not.*

Proof. A rule $\tilde{c} = \Omega \tilde{a}$ that solves the problem solves the Riccati equation (33), so by (35) $\zeta = (\hat{\beta}R)^{-1}$ and $\hat{\beta}\zeta^2 = (\hat{\beta}R^2)^{-1}$. Under the worst-case law

$$\widehat{\mathbb{E}}_0 \tilde{a}_t^2 = \zeta^{2t} \tilde{a}_0^2 + s^2 \tilde{\Sigma} \sum_{j=0}^{t-1} \zeta^{2j}, \quad (72)$$

so for $\hat{\beta} \in (0, 1)$ the discounted sum $\sum_t \hat{\beta}^t \widehat{\mathbb{E}}_0 \tilde{a}_t^2$ is finite if and only if $\hat{\beta}\zeta^2 < 1$, that is, if and only if $\hat{\beta}R^2 > 1$.

Let $\hat{\beta}R^2 > 1$. By (34), $1 - \Omega = (1 + \sigma s^2)/(\sigma s^2 + \hat{\beta}R^2)$, and by (33), $1 + \sigma\Omega s^2 = \hat{\beta}R^2(1 - \Omega)$. If $\sigma s^2 > -1$, then $\Omega \in (0, 1)$ and $1 + \sigma\Omega s^2 > 0$. If $-\hat{\beta}R^2 < \sigma s^2 \leq -1$, then $\Omega \geq 1$ and $1 + \sigma\Omega s^2 \leq 0$, so the worst-case variance (37) is unbounded and (25) diverges. If $\sigma s^2 < -\hat{\beta}R^2$, then $\Omega < 0$, and the candidate value $-\frac{1}{2}\Omega \tilde{a}^2$ is positive, above the upper bound of zero that the bliss point imposes, so it is not the value of the problem. Since $s^2 = \alpha^2/(1 - \beta)^2$, $\sigma s^2 > -1$ if and only if $\sigma > \underline{\sigma}$. For $\sigma > \underline{\sigma}$ and $\hat{\beta}R^2 > 1$, Appendix C.3 shows that the root (34) solves the Bellman equation (23).

If $\hat{\beta}R^2 < 1$, the bliss plan sets $\tilde{c}_t = 0$ at every date. The minimizing agent then leaves the distribution of the shocks undistorted, the objective (31) attains its upper bound of zero, and net wealth grows at rate R , which satisfies the admissibility condition because $\hat{\beta}R^2 < 1$. The bliss plan violates (3), which the auxiliary problem does not impose, so by standing assumption (iii) it is not a feasible equilibrium choice. If $\hat{\beta}R^2 = 1$, both the bliss plan and the root (34), which then equals zero or is undefined, give $\hat{\beta}\zeta^2 = 1$, and the strict admissibility condition fails.

On the locus $\hat{\beta}(\sigma)R^2 = [1 + \sigma\alpha^2/(1 - \beta)]/\beta$, which exceeds one if and only if $\sigma\alpha^2 >$

$-(1 - \beta)^2$, that is, if and only if $\sigma > \underline{\sigma}$. At $\sigma = \underline{\sigma}$ we have $\sigma\alpha^2 = -(1 - \beta)^2$, so $\zeta = [1 - (1 - \beta)]^{-1} = \beta^{-1}$ and $\hat{\beta} = \beta/\zeta = \beta^2$. $\square$

On the locus an equivalent statement avoids the minimizing agent altogether. By (66) the impatient twin's marginal propensity to consume is $\Omega_i^{\text{twin}} = 1 - \beta\zeta_i$, which is positive if and only if $\zeta_i < \beta^{-1}$, that is, if and only if $\hat{\beta}_i > \beta^2$. The twin's problem is an ordinary linear regulator, and it has an interior solution exactly on the admissible range. As the lower endpoint is approached along the locus, the robust consumer's marginal propensity to consume remains $1 - \beta$, the twin's marginal propensity to consume approaches zero, and the welfare loss (68) approaches one. The strict admissibility condition excludes the endpoint.

The two conditions of Proposition 9 have different sources. The condition $\hat{\beta}R^2 > 1$ is intertemporal: it makes the worst-case path satisfy the admissibility condition and rules out the bliss plan. The condition $\sigma > \underline{\sigma}$ is local to one period. Given $\hat{\beta}R^2 > 1$, it keeps the root of (33) where the log-exp formula (25) is finite, $1 + \sigma\Omega s^2 > 0$, with a positive marginal propensity to consume. As $1 + \sigma\Omega s^2 \downarrow 0$ the worst-case variance (37) grows without bound. This is the requirement in Jacobson (1973) and Hansen and Sargent (1995) that $I + \sigma C^\top V C$ be positive definite. The bound $\underline{\sigma}$ does not depend on $\hat{\beta}$, and the bound $\hat{\beta}R^2 > 1$ does not depend on σ . On the locus the two bounds meet at the corner $(\underline{\sigma}, \beta^2)$. As the corner is approached along the locus, $1 + \sigma\Omega s^2 = \zeta^{-1} \rightarrow \beta$ and $\tilde{\Sigma} = \zeta \rightarrow R$ stay finite. The one-period condition does not fail in that limit. The intertemporal condition fails: $\hat{\beta}\zeta^2 \rightarrow 1$.

The breakdown point tightens as the endowment becomes more volatile, because σ enters only through $\sigma\alpha^2$. When a consumer faces more income risk, a smaller $|\sigma|$ already buys a given amount of pessimism. It also tightens as β rises toward one. Corollary 1 restates the same bound in terms of discount factors $\hat{\beta}_i$: whatever the endowment process, the admissible interval is $(\beta^2, \beta]$.

B Mortality and a stationary cross-section

Proposition 10 (Mortality preserves the construction and restores stationarity). *Let each consumer survive to the next date with probability $1 - \chi$, $\chi \in (0, 1)$, let a mass χ of newborns enter every date with zero financial wealth and the average income state, and let annuities be fair, paying survivors the return $R/(1 - \chi)$ on balances carried into the period. With $\tilde{\beta}_i \equiv \hat{\beta}_i(1 - \chi)$ and $\tilde{R} \equiv R/(1 - \chi)$ in place of $\hat{\beta}_i$ and R throughout Sections 3 and 4:*

- (i) $R = \beta^{-1}$ remains an equilibrium gross interest rate;
- (ii) the locus, observational equivalence, and the invariance of the equilibrium to Φ continue to hold, with $\tilde{\alpha}^2$ and $1 - \tilde{\beta}$ in place of α^2 and $1 - \beta$;

(iii) *the cross-section of consumption is stationary rather than growing without bound, as is the cross-section of debt whenever the endowment has no explosive root; and the bound of Corollary 1 becomes $\tilde{\beta}_i \in (\tilde{\beta}^2, \tilde{\beta}]$, equivalently $\hat{\beta}_i > \beta^2(1 - \chi)$.*

Proof. Sections B.1–B.3 below. Section B.4 is a further consequence, not part of the statement above. $\square$

B.1 The mapping

Let each consumer survive from one date to the next with probability $1 - \chi$, and let a mass χ of newborns enter at every date, each endowed with zero financial wealth and the average income state. Let annuities be fair: an insurer collects the balance of every consumer who dies and pays survivors a return $R/(1 - \chi)$ on the balances they carry into the period. A survivor's problem is then exactly the one studied in Section 3, with

$$\tilde{\beta}_i \equiv \hat{\beta}_i(1 - \chi), \quad \tilde{R} \equiv R/(1 - \chi) \quad (73)$$

in place of $\hat{\beta}_i$ and R : utility accrues only while alive, so the effective discount factor is scaled by the survival probability, and the budget constraint carries the annuity-adjusted return $\tilde{R}$. Every reduced-form object of Sections 3 and 4 follows by substituting $\tilde{\beta}_i$ for $\hat{\beta}_i$ and $\tilde{R}$ for R .

Two modeling choices make the substitution exact. First, the entropy penalty that restricts the worst-case model is discounted at the same survival-adjusted rate as utility: the consumer guards against misspecification over his own lifetime, not beyond it. Second, the minimizing agent is allowed to distort the endowment but not the survival hazard. The second choice matters: flow utility is nonpositive and death delivers zero, so a minimizing agent free to twist the survival probability would find it worthwhile to do so, for reasons unrelated to the misspecification concern the paper studies.

One object does not simply carry over. Income accrues only while a consumer is alive, so human wealth is discounted at $\tilde{R}^{-1} = \beta(1 - \chi)$, and the innovation variance α^2 (19) must be recomputed at the survival-adjusted discount factor; call the result $\tilde{\alpha}^2$.

B.2 The equilibrium interest rate is unchanged

The benchmark condition $R = \beta^{-1}$ of Proposition 4 becomes, in the mortality economy, $\tilde{R} = \tilde{\beta}^{-1}$, that is, $R/(1 - \chi) = [\beta(1 - \chi)]^{-1}$, which again gives $R = \beta^{-1}$: mortality and the annuity return cancel. The locus (39) itself is evaluated at the survival-adjusted parameters, $\hat{\beta}(\sigma) = \beta(1 + \sigma\tilde{\alpha}^2/(1 - \tilde{\beta}))$ with $\tilde{\alpha}^2$ and $1 - \tilde{\beta}$ in place of α^2 and $1 - \beta$, so the locus shifts with χ even though it keeps the same functional form. What is exactly unchanged is worst-case

persistence as a function of the discount factor, $\zeta_i = \tilde{\beta}/\hat{\beta}_i = \beta/\hat{\beta}_i$: a consumer's worst-case persistence depends on $\hat{\beta}_i$ exactly as it does when $\chi = 0$.

B.3 Aggregation

Provided newborns enter with zero financial wealth and at the average income state, so that the cross-section conditions (11) and (12) apply to the inflow as well as to the initial cross-section, the argument of the proof of Proposition 4 (Appendix C.11) extends unchanged, market by market.

- (a) *Goods market.* Every type consumes the fraction $1 - \tilde{\beta}$ of wealth discounted at $\tilde{R}$; aggregate financial wealth is zero and aggregate human wealth is $Y/(1 - \tilde{\beta})$, so aggregate consumption remains Y .
- (b) *Bond market.* The dead are a random sample of the population, so their average balance equals the population average, zero; the annuity is self-financing, and newborns enter with zero balances.
- (c) *Invariance to Φ .* Unchanged: every type still uses one decision rule, and the common hazard χ creates no cross-subsidy between types.

The age distribution in the cross-section is geometric with mean $(1 - \chi)/\chi$, and within-cohort consumption variance grows at rate $\tilde{\alpha}^2$ a period, so the cross-section variance of consumption is stationary rather than growing without bound in t , in place of part (iii) of Proposition 4. The same calculation gives a stationary cross-section of debt whenever the endowment itself does not have an explosive root, as in the two-factor process of Section 9. The bound of Corollary 1 applies to the survival-adjusted factors: $\tilde{\beta}_i \in (\tilde{\beta}^2, \tilde{\beta}]$, with $\tilde{\beta}_i = \hat{\beta}_i(1 - \chi)$ and $\tilde{\beta} = \beta(1 - \chi)$. In terms of a consumer's own discount factor the admissible floor falls from β^2 to $\beta^2(1 - \chi)$, and the upper end stays at β : mortality widens the admissible range downward by the survival factor. At $\beta = 0.95$ the floor is 0.884 for $\chi = 0.02$ and 0.857 for $\chi = 0.05$, against 0.9025 with no mortality.

B.4 The share of the population above bliss

A newborn enters with zero financial wealth and the average income state, so by the decision rule he consumes exactly $(1 - \tilde{\beta}) \cdot Y/(1 - \tilde{\beta}) = Y$: his gap from bliss is $\gamma - Y$. It is not a free initial condition chosen for this calculation; it is implied by γ once γ is fixed, and $(\gamma - Y)/\alpha$ is the same quantity already tabulated as $|\mu_0|/\alpha$ in Table 2. (Quantity data do not identify γ itself; Section 4.) The knife-edge case $\gamma = Y$ is the one Remark 3 singles out,

where mean net wealth relative to bliss is zero within every type and stays zero; there every newborn enters exactly at bliss, and the stationary share of the cross-section strictly above bliss is close to one half for every χ , approaching it exactly as $\chi \rightarrow 0$. Away from that case, a positive χ pins the stationary share above bliss down to a value set by χ and $(\gamma - Y)/\alpha$, in contrast to its drift toward one half as $t \rightarrow \infty$ when $\chi = 0$ (Table 2, second row).

Table 4 reports the stationary share above bliss at the parameter values of Section 9 ($\alpha = \sqrt{0.022725} = 0.151$) and $\chi = 2.5$ percent.

$(\gamma - Y)/\alpha$	0	2	4	8
Stationary share above bliss	0.50	0.32	0.20	0.08

Table 4: Stationary share of the cross-section above bliss under mortality, at $\chi = 2.5$ percent and the parameter values of Section 9.

Reading mean income as one, $(\gamma - Y)/\alpha = 4$ and 8 put bliss 60 and 120 percent above mean income, respectively.

Nothing else in the body of the paper changes. The numerical illustration of Section 9, including the detection error probabilities of Section 7, is reported for the $\chi = 0$ case throughout. Mortality bounds how long a single consumer is typically observed, with expected lifetime $1/\chi$ periods, forty years at $\chi = 2.5$ percent, the same horizon as the $T = 40$ column of Table 1, which bears directly on the single-consumer detection error probability of Section 7.1: redoing that calculation for $\chi > 0$, alongside the panel alternative discussed there, is left for future work.

C The scalar reduction and proofs

Sections C.1 to C.4 derive the scalar reduction of Section 3.2 and Proposition 2. Sections C.5 to C.13 state and prove Lemma 1 and prove Propositions 6, 3, 5, 7, 8, 4, and Corollaries 5 and 1. Section C.14 describes the replication scripts. Throughout, $R = \beta^{-1}$ is the market interest rate and $\hat{\beta}$ is the discount factor of the consumer under consideration.

C.1 A scalar state

Let $H_t \equiv G(I - \beta A)^{-1} z_t$ be human wealth valued at the market discount factor. From the identity $(I - \beta A)^{-1} A = \beta^{-1} [(I - \beta A)^{-1} - I]$ and (4),

$$H_{t+1} = R(H_t - y_t) + \frac{1}{1 - \beta} h w_{t+1}, \quad (74)$$

where h is defined in (9). The budget constraint (2) is $b_{t+1} = R(b_t + c_t - y_t)$. Subtracting, net wealth $a_t = H_t - b_t$ of (7) obeys

$$a_{t+1} = R(a_t - c_t) + \frac{1}{1 - \beta} h w_{t+1}, \quad (75)$$

in which the endowment vector z_t has disappeared. Setting $c_t = \gamma$ for all t in the deterministic version of (75) gives the stationary point $\gamma/(1 - \beta)$, so with $\tilde{a}$ and $\tilde{c}$ defined in (27) and $\epsilon_{t+1} = h w_{t+1}/\alpha = \tilde{\epsilon}_{t+1} + \tilde{v}_{t+1}$,

$$\tilde{a}_{t+1} = R(\tilde{a}_t - \tilde{c}_t) + s \epsilon_{t+1}, \quad s = \frac{\alpha}{1 - \beta}, \quad (76)$$

and the objective (1) is $-\frac{1}{2}\mathbb{E}_0 \sum_t \hat{\beta}^t \tilde{c}_t^2$ for a consumer with $\sigma = 0$. Sections C.2 and C.3 add the distortion and replace this objective by (31).

C.2 The distortion collapses to a scalar

The shock w_{t+1} affects $\tilde{a}_{t+1}$ in (76) only through $\epsilon_{t+1} = h w_{t+1}/\alpha$. Write an alternative joint density of w_{t+1} as the product of a marginal density of ϵ_{t+1} and a conditional density of w_{t+1} given ϵ_{t+1} . The chain rule for relative entropy expresses the relative entropy of the joint density as the relative entropy of the marginal plus the expected relative entropy of the conditional. Only the marginal affects the objective, so the minimizing agent leaves the conditional density at its baseline and distorts only the marginal density of ϵ_{t+1} . Section C.3 shows that the worst-case marginal is normal. For a normal marginal with mean $\tilde{v}_{t+1}$ and variance $\tilde{\Sigma}_{t+1}$, relative entropy is (29), the scalar case of (26). Under the baseline, w_{t+1} given ϵ_{t+1} is normal with mean $\epsilon_{t+1} h^\top/\alpha$ and covariance matrix $I - h^\top h/\alpha^2$. The implied distortion of w_{t+1} therefore has mean and covariance matrix

$$v_{t+1} = \frac{\tilde{v}_{t+1}}{\alpha} h^\top, \quad \Sigma_{t+1} = I + \frac{\tilde{\Sigma}_{t+1} - 1}{\alpha^2} h^\top h, \quad (77)$$

with $v_{t+1}^\top v_{t+1} = \tilde{v}_{t+1}^2$. This converts (76) into (28). Equation (77) gives the claims of Lemma 1 about the direction of the distortion.

C.3 The Bellman equation

Guess a value function $-\frac{1}{2}\Omega\tilde{a}^2 - \mathbf{p}_s$. Writing $n \equiv R(\tilde{a} - \tilde{c})$, so that $\mathbb{E}^{\tilde{v}, \tilde{\Sigma}}(\tilde{a}')^2 = (n + s\tilde{v})^2 + s^2\tilde{\Sigma}$ under a normal alternative, the Bellman equation (23) in reduced form is, written for the

negative of the value function so that its sup inf becomes a min max,

$$\frac{1}{2}\Omega\tilde{a}^2 + \mathbf{p}_s = \min_{\tilde{c}} \max_{\tilde{v}, \tilde{\Sigma}} \left\{ \frac{1}{2}\tilde{c}^2 + \hat{\beta} \left[\frac{1}{2}\Omega \left((n + s\tilde{v})^2 + s^2\tilde{\Sigma} \right) - \theta \mathcal{E}(\tilde{v}, \tilde{\Sigma}) + \mathbf{p}_s \right] \right\}. \quad (78)$$

By Remark 2 the minimizing density ratio is proportional to $\exp(\sigma V(\tilde{a}'))$, the exponential of a quadratic function of ϵ' , so the worst-case law of ϵ' is normal and normal alternatives suffice. The inner objective is additively separable in $\tilde{v}$ and $\tilde{\Sigma}$. The first-order condition for $\tilde{v}$ is $\Omega s(n + s\tilde{v}) = \theta\tilde{v}$, hence

$$\tilde{v} = \frac{\Omega s n}{\theta - \Omega s^2} = -\frac{\sigma \Omega s n}{1 + \sigma \Omega s^2}, \quad n + s\tilde{v} = \frac{n}{1 + \sigma \Omega s^2}, \quad (79)$$

a maximum provided $1 + \sigma \Omega s^2 > 0$. The first-order condition for $\tilde{\Sigma}$ is $\Omega s^2 = \theta(1 - \tilde{\Sigma}^{-1})$, hence $\tilde{\Sigma} = (1 + \sigma \Omega s^2)^{-1}$, which is (37) and requires the same condition. Substituting,

$$\frac{1}{2}\Omega \left((n + s\tilde{v})^2 + s^2\tilde{\Sigma} \right) - \theta \mathcal{E}(\tilde{v}, \tilde{\Sigma}) = \frac{1}{2}\mathcal{D}(\Omega) n^2 + \frac{1}{2\sigma} \log(1 + \sigma \Omega s^2), \quad \mathcal{D}(\Omega) \equiv \frac{\Omega}{1 + \sigma \Omega s^2}. \quad (80)$$

For $\sigma < 0$ we have $\mathcal{D}(\Omega) > \Omega$: the minimizing agent raises the effective curvature of the continuation value. The outer minimization is then $\min_{\tilde{c}} \{ \frac{1}{2}\tilde{c}^2 + \frac{1}{2}\kappa(\tilde{a} - \tilde{c})^2 \}$ with $\kappa \equiv \hat{\beta} R^2 \mathcal{D}(\Omega)$, whose solution is $\tilde{c} = \kappa\tilde{a}/(1 + \kappa)$ with value $\frac{1}{2}\kappa\tilde{a}^2/(1 + \kappa)$. Matching coefficients,

$$\Omega = \frac{\kappa}{1 + \kappa} \quad \text{and} \quad \tilde{c}_t = \Omega \tilde{a}_t, \quad (81)$$

so Ω is the marginal propensity to consume out of $\tilde{a}$, which gives (32). Eliminating κ between (81) and $\kappa = \hat{\beta} R^2 \Omega / (1 + \sigma \Omega s^2)$ gives $\Omega \left[1 + \sigma \Omega s^2 - \hat{\beta} R^2 (1 - \Omega) \right] = 0$. The root $\Omega = 0$ is the bliss plan, under which net wealth grows at rate R ; since $\hat{\beta} R^2 > 1$ on the admissible range, it violates the integrability restriction of standing assumption (iii). The other root is (33), whose solution is (34). The guessed value function is quadratic, the rule $\tilde{c} = \Omega\tilde{a}$ makes $\hat{\mathbb{E}}_0 \tilde{a}_t^2$ grow at rate ζ^2 , and $\hat{\beta}\zeta^2 < 1$ sends $\hat{\beta}^t \hat{\mathbb{E}}_0 \tilde{a}_t^2$ to zero. That transversality condition verifies that the solution of (78) is the value of the problem over integrability-admissible plans. Appendix C.10 verifies the same thing directly, without the quadratic guess, by comparing an arbitrary plan with the equilibrium plan under a fixed worst-case law.

Matching constants in (78) gives

$$\mathbf{p}_s = \frac{\hat{\beta}}{2\sigma(1 - \hat{\beta})} \log(1 + \sigma \Omega s^2). \quad (82)$$

Restricting the minimizing agent to $\tilde{\Sigma} = 1$, as in Hansen et al. (1999, eq. 9), replaces the last

term of (80) by $\frac{1}{2}\Omega s^2$ and gives $\mathbf{p}_s = \hat{\beta}\Omega s^2/(2(1 - \hat{\beta}))$. Removing the shocks gives $\mathbf{p}_s = 0$. The three problems share Ω , the decision rule, and $\tilde{v}$. Because $\log(1 + x) < x$ for $x \neq 0$, (82) exceeds the mean-only constant for every $\sigma < 0$; the two coincide as $\sigma \rightarrow 0$. Luo and Young (2010, Proposition 11) use this ranking to show that a consumer who distorts the covariance matrix has lower utility than an observationally equivalent consumer who distorts only means. On the locus $1 + \sigma\Omega s^2 = \zeta^{-1}$, and (82) equals $\hat{\beta} \log \zeta / (2|\sigma|(1 - \hat{\beta}))$.

Equation (82) is the constant for the bliss-centered scalar state. The vector value function (22) normalizes the constant entry of $\mathcal{P}$ to zero, so its constant also carries the bliss term:

$$\mathbf{p}_{\text{vec}} = \mathbf{p}_s + \frac{\Omega\gamma^2}{2(1 - \beta)^2}. \quad (83)$$

The distinction leaves Ω , the locus, and the welfare ratios of Section 8 unchanged, since those compare values computed from consumption paths.

C.4 Worst-case dynamics

Combining (79) with $n = R(1 - \Omega)\tilde{a}$,

$$\hat{\mathbb{E}}_t \tilde{a}_{t+1} = \frac{R(1 - \Omega)}{1 + \sigma\Omega s^2} \tilde{a}_t, \quad \mathbb{E}_t \tilde{a}_{t+1} = R(1 - \Omega) \tilde{a}_t, \quad (84)$$

and the Riccati equation (33) turns the first coefficient into $R(1 - \Omega)/[\hat{\beta}R^2(1 - \Omega)] = (\hat{\beta}R)^{-1}$, which is (35). Since $\mu_t = -\tilde{c}_t = -\Omega\tilde{a}_t$, the same two persistences govern marginal utility. Under the worst-case law the conditional variance of ϵ_{t+1} is $\tilde{\Sigma} = (1 + \sigma\Omega s^2)^{-1}$, and the Riccati equation (33) writes it as $1/(\hat{\beta}R^2(1 - \Omega))$, which is (37).

C.5 Detection with the vector endowment

Lemma 1. *Along the locus (39) the minimizing agent's distortion of the conditional mean of w_{t+1} is*

$$v_{t+1} = -\frac{\zeta - 1}{\alpha^2} h^\top \mu_t, \quad (85)$$

and it raises the conditional covariance matrix of w_{t+1} from I to

$$I + \frac{\zeta - 1}{\alpha^2} h^\top h. \quad (86)$$

Both distortions lie in the one-dimensional subspace of shock space spanned by $h^\top$. The log-likelihood ratio between the approximating and the worst-case laws of the vector endowment

(4) therefore depends on the data only through hw_{t+1} , and it coincides with the log-likelihood ratio statistic used to test between the two hypotheses in (60).

Proof. On the locus, $\Omega = 1 - \beta$ and $R(1 - \Omega) = 1$, so $s\tilde{v}_{t+1} = (\zeta - 1)\tilde{a}_t = -(\zeta - 1)\mu_t/(1 - \beta)$, whence $\tilde{v}_{t+1} = -\alpha^{-1}(\zeta - 1)\mu_t$ and, by (77), equation (85). On the locus $\tilde{\Sigma} = \zeta$ (Proposition 3), and (77) gives (86). The worst-case law of w_{t+1} changes only the distribution of $\epsilon_{t+1} = hw_{t+1}/\alpha$ and leaves the conditional distribution of w_{t+1} given ϵ_{t+1} at its baseline (Section C.2). The likelihood ratio of the two laws of w_{t+1} therefore equals the likelihood ratio of the two laws of ϵ_{t+1} , a function of hw_{t+1} alone. Since $\mu_{t+1} - \mu_t = -\alpha\epsilon_{t+1}$ on the locus, it coincides with the statistic for (60). $\square$

C.6 Proof of Proposition 6

Both laws in (60) have Gaussian innovations, with variances α^2 and $\zeta\alpha^2$, so the log-likelihood ratio for a sample $\{\mu_t\}_{t=0}^T$ is the statistic S of (62), a function of ζ and the normalized data $\{\mu_t/\alpha\}$ alone. Under either law μ_{t+1}/α is a multiple of μ_t/α plus a multiple of $\tilde{w}_{t+1}$, with coefficients that depend only on ζ . Hence the distribution of S under either law is determined by ζ , T , and μ_0/α . Under a change of units $\mu \mapsto \lambda\mu$, $\alpha \mapsto \lambda\alpha$, every ratio in the expression is unchanged, and so is its sign. Since (61) depends on the data only through that sign, the detection error probability is unchanged.

C.7 Proof of Proposition 3 (Observational equivalence)

By (32), a consumer chooses (8) if and only if $\Omega = 1 - \beta$. This is also the only value of Ω at which the bliss point drops out of the decision rule. Substituting $\Omega = 1 - \beta$ into the Riccati equation (33) and using $s^2 = \alpha^2/(1 - \beta)^2$ and $R = \beta^{-1}$,

$$1 + \frac{\sigma\alpha^2}{1 - \beta} = \frac{\hat{\beta}}{\beta^2}\beta = \frac{\hat{\beta}}{\beta},$$

which is (39). Equation (40) is $\zeta = (\hat{\beta}R)^{-1}$ from (35). Since $\Omega = 1 - \beta$ makes $R(1 - \Omega) = R\beta = 1$, consumption is a martingale under the approximating model, as in (6). The same fact and (37) give $\tilde{\Sigma} = (\hat{\beta}R)^{-1} = \zeta$. The claim about $\underline{\sigma}$ is Proposition 9 in Appendix A.

C.8 Proof of Proposition 5 (Forecasts across the locus)

With a constant bliss point, $\mu_t^i = \gamma - c_t^i$. Iterating the worst-case law (35) gives $\hat{\mathbb{E}}_t\mu_{t+j}^i = \zeta_i^j\mu_t^i$. Substituting $\mu_{t+j}^i = \gamma - c_{t+j}^i$ and rearranging gives (48). That step uses only $\zeta_i = (\hat{\beta}_iR)^{-1}$ from (35), which holds at every interior $(\sigma_i, \hat{\beta}_i)$. The proposition's last clause restates (10):

on the locus $R(1 - \Omega) = 1$, so μ_t^i is a martingale under the approximating model and $\mathbb{E}_t c_{t+j}^i = c_t^i$ for every j . Off the locus $R(1 - \Omega) \neq 1$ and consumption is not a martingale.

C.9 Proof of Proposition 7

Under a rule $\tilde{c}_t = \Omega' \tilde{a}_t$ the state follows $\tilde{a}_{t+1} = \varrho \tilde{a}_t + s \epsilon_{t+1}$ with $\varrho = R(1 - \Omega')$, so $\mathbb{E}_0 \tilde{a}_t^2 = \varrho^{2t} \tilde{a}_0^2 + s^2(1 - \varrho^{2t})/(1 - \varrho^2)$ for $\varrho \neq 1$, and the value of that rule under the $(0, \hat{\beta}_i)$ criterion is

$$V(\Omega') = -\frac{\Omega'^2}{2} \sum_{t \geq 0} \hat{\beta}_i^t \mathbb{E}_0 \tilde{a}_t^2 = -\frac{\Omega'^2}{2(1 - \hat{\beta}_i \varrho^2)} \left[\tilde{a}_0^2 + \frac{\hat{\beta}_i s^2}{1 - \hat{\beta}_i} \right], \quad (87)$$

where the second equality uses

$$\frac{1}{1 - \varrho^2} \left[\frac{1}{1 - \hat{\beta}_i} - \frac{1}{1 - \hat{\beta}_i \varrho^2} \right] = \frac{\hat{\beta}_i}{(1 - \hat{\beta}_i)(1 - \hat{\beta}_i \varrho^2)}$$

and holds for $\varrho = 1$ by continuity. The actual plan has $\Omega' = 1 - \beta$ and $\varrho = 1$; the twin's optimal plan has $\Omega' = \Omega_i^{\text{twin}} = 1 - \beta \zeta_i$ and $\varrho = \zeta_i$, and $1 - \hat{\beta}_i \zeta_i^2 = 1 - \beta \zeta_i = \Omega_i^{\text{twin}}$. The bracket in (87) is common to the two, so

$$\frac{V_i^{\text{opt}}}{V_i^{\text{act}}} = \frac{(\Omega_i^{\text{twin}})^2 / (1 - \hat{\beta}_i \zeta_i^2)}{(1 - \beta)^2 / (1 - \hat{\beta}_i)} = \frac{\Omega_i^{\text{twin}} (1 - \hat{\beta}_i)}{(1 - \beta)^2} = \frac{(1 - \beta \zeta_i) (1 - \beta \zeta_i^{-1})}{(1 - \beta)^2}, \quad (88)$$

and $1 - (1 - \beta \zeta_i)(1 - \beta \zeta_i^{-1}) / (1 - \beta)^2 = \beta(\zeta_i - 1)^2 / (\zeta_i(1 - \beta)^2)$. Substituting $\zeta_i = 1 / (1 - (1 - \beta)\tau_i)$ from (67) gives $\zeta_i - 1 = (1 - \beta)\tau_i \zeta_i$ and hence the last expression in (68). The ratio is one at $\zeta_i = 1$ and approaches zero as $\zeta_i \uparrow \beta^{-1}$, so Λ_i rises from zero toward one.

C.10 Proof of Proposition 8

Under the worst-case model $\mathbb{Q}^*$ of Section 5.2, the equilibrium plan's net wealth follows $\tilde{a}_{t+1}^* = \zeta_i \tilde{a}_t^* + s \nu_{t+1}$ with $\nu_{t+1} \sim N(0, \zeta_i)$, which is (43). That law shifts the conditional mean of ϵ_{t+1} by $(\zeta_i - 1)\tilde{a}_t^*/s$, so a consumer who faces the same endowment process and follows $\tilde{c}_t = \Omega' \tilde{a}_t$ has

$$\tilde{a}_{t+1} = \varrho' \tilde{a}_t + (\zeta_i - 1) \tilde{a}_t^* + s \nu_{t+1}, \quad \varrho' = R(1 - \Omega'). \quad (89)$$

Let M_t be the second moment matrix of $(\tilde{a}_t^*, \tilde{a}_t)^\top$, let $\iota = (1, 1)^\top$, and let

$$A = \begin{pmatrix} \zeta_i & 0 \\ \zeta_i - 1 & \varrho' \end{pmatrix}, \quad N = s^2 \zeta_i \iota \iota^\top,$$

so that $M_{t+1} = AM_t A^\top + N$. The discounted sum $\mathcal{S} \equiv \sum_{t \geq 0} \hat{\beta}_i^t M_t$ solves

$$\mathcal{S} = M_0 + \hat{\beta}_i A \mathcal{S} A^\top + \frac{\hat{\beta}_i}{1 - \hat{\beta}_i} N, \quad (90)$$

which has a solution when $\hat{\beta}_i \max\{\zeta_i^2, (\varrho')^2\} < 1$. That holds for the two rules compared below: $\varrho' = 1$ gives $\hat{\beta}_i < 1$, and $\varrho' = \zeta_i$ gives $\hat{\beta}_i \zeta_i^2 = \beta \zeta_i < 1$ on the admissible range. The value of the rule is $\tilde{V}(\Omega') = -\frac{1}{2} \Omega'^2 \mathcal{S}_{22}$. With $M_0 = \tilde{a}_0^2 \mu^\top$, both inhomogeneous terms in (90) are multiples of $\mu^\top$, so $\mathcal{S} = \left[\tilde{a}_0^2 + \hat{\beta}_i s^2 \zeta_i / (1 - \hat{\beta}_i) \right] \hat{\mathcal{S}}$ with $\hat{\mathcal{S}} = \mu^\top + \hat{\beta}_i A \hat{\mathcal{S}} A^\top$. That factor cancels in (70), which is why $\tilde{\Lambda}_i$ depends on neither $\tilde{a}_0$ nor α^2 nor γ .

Write the entries of $\hat{\mathcal{S}}$ as $\hat{\mathcal{S}}_{11}$, $\hat{\mathcal{S}}_{12}$, $\hat{\mathcal{S}}_{22}$. Consumer i 's plan has $\Omega' = 1 - \beta$ and $\varrho' = 1$, under which (89) makes $\tilde{a}_t = \tilde{a}_t^*$ at every date, so $\hat{\mathcal{S}} = \hat{\mathcal{S}}_{11} \mu^\top$ with $\hat{\mathcal{S}}_{11} = (1 - \hat{\beta}_i \zeta_i^2)^{-1} = (1 - \beta \zeta_i)^{-1}$. The twin's plan has $\Omega' = 1 - \beta \zeta_i$ and $\varrho' = \zeta_i$, so, using $\hat{\beta}_i = \beta / \zeta_i$,

$$\hat{\mathcal{S}}_{12} = \frac{1 - \beta}{(1 - \beta \zeta_i)^2}, \quad \hat{\mathcal{S}}_{22} (1 - \beta \zeta_i) = 1 + \frac{\beta}{\zeta_i} \frac{(\zeta_i - 1)^2}{1 - \beta \zeta_i} + \frac{2\beta(1 - \beta)(\zeta_i - 1)}{(1 - \beta \zeta_i)^2}.$$

Then $\tilde{V}_i^{\text{own}}$ is proportional to $-\frac{1}{2}(1 - \beta)^2 \hat{\mathcal{S}}_{11}$ and $\tilde{V}_i^{\text{twin}}$ to $-\frac{1}{2}(1 - \beta \zeta_i)^2 \hat{\mathcal{S}}_{22}$. Multiplying out the expression for $\hat{\mathcal{S}}_{22}$ and using (68) gives

$$(1 - \beta \zeta_i)^3 \hat{\mathcal{S}}_{22} = (1 - \beta)^2 (1 + \Lambda_i),$$

so $\tilde{V}_i^{\text{own}} / \tilde{V}_i^{\text{twin}} = (1 - \beta)^2 \hat{\mathcal{S}}_{11} / \left[(1 - \beta \zeta_i)^2 \hat{\mathcal{S}}_{22} \right] = 1 / (1 + \Lambda_i)$ and $\tilde{\Lambda}_i = \Lambda_i / (1 + \Lambda_i)$, which is (70). Since Λ_i rises from zero toward one, $\tilde{\Lambda}_i$ rises from zero toward one half.

It remains to show that the equilibrium plan is the best response to $\mathbb{Q}^*$, so that $\tilde{\Lambda}_i$, like Λ_i , is a re-optimization loss. Fix $\mathbb{Q}^*$ and the initial state and let $\tilde{c}$ be any adapted plan with $\hat{\mathbb{E}}_0 \sum_t \hat{\beta}_i^t \tilde{a}_t^2 < +\infty$. Write $\delta \tilde{a}_t = \tilde{a}_t - \tilde{a}_t^*$ and $\delta \tilde{c}_t = \tilde{c}_t - \tilde{c}_t^*$ for its deviations from the equilibrium plan. Both plans face the same endowment realization, so (44) gives

$$\delta \tilde{a}_{t+1} = R(\delta \tilde{a}_t - \delta \tilde{c}_t), \quad \delta \tilde{a}_0 = 0, \quad (91)$$

and $\delta \tilde{a}_{t+1}$ is known at t . Quadratic utility gives exactly

$$u(c_t) - u(c_t^*) = \mu_t^* \delta \tilde{c}_t - \frac{1}{2} (\delta \tilde{c}_t)^2, \quad \mu_t^* = \gamma - c_t^*.$$

Substituting $\delta \tilde{c}_t = \delta \tilde{a}_t - \delta \tilde{a}_{t+1} / R$ and using the Euler equation $\hat{\beta}_i R \hat{\mathbb{E}}_t \mu_{t+1}^* = \mu_t^*$ under $\mathbb{Q}^*$,

the linear terms telescope:

$$\sum_{t=0}^N \hat{\beta}_i^t \hat{\mathbb{E}}_0 [\mu_t^* \delta \tilde{c}_t] = -\hat{\beta}_i^{N+1} \hat{\mathbb{E}}_0 [\mu_{N+1}^* \delta \tilde{a}_{N+1}].$$

By Cauchy–Schwarz the right side is bounded in absolute value by

$$\left(\hat{\beta}_i^{N+1} \hat{\mathbb{E}}_0 (\mu_{N+1}^*)^2 \right)^{1/2} \left(\hat{\beta}_i^{N+1} \hat{\mathbb{E}}_0 \delta \tilde{a}_{N+1}^2 \right)^{1/2},$$

and both factors vanish because the two plans have finite discounted second moments. Hence

$$\hat{\mathbb{E}}_0 \sum_{t \geq 0} \hat{\beta}_i^t [u(c_t) - u(c_t^*)] = -\frac{1}{2} \hat{\mathbb{E}}_0 \sum_{t \geq 0} \hat{\beta}_i^t (\delta \tilde{c}_t)^2 \leq 0, \quad (92)$$

with equality only if $\delta \tilde{c}_t = 0$ almost surely. The equilibrium plan is therefore the unique best response in this domain. The entropy penalty does not enter the comparison, because $\mathbb{Q}^*$ is held fixed. The bliss plan attains the upper bound of zero but is not integrability-admissible.

Inequality (92) also settles optimality under the multiplier criterion itself, without restricting attention to rules of the form $\tilde{c} = \Omega \tilde{a}$. For any plan c in this domain,

$$\begin{aligned} V^{\text{mult}}(c) &= \inf_{\mathbb{Q}} \{ \mathbb{E}_{\mathbb{Q}} u(c) + \theta \mathcal{H}(\mathbb{Q}) \} \leq \mathbb{E}_{\mathbb{Q}^*} u(c) + \theta \mathcal{H}(\mathbb{Q}^*) \\ &\leq \mathbb{E}_{\mathbb{Q}^*} u(c^*) + \theta \mathcal{H}(\mathbb{Q}^*) = V^{\text{mult}}(c^*), \end{aligned}$$

where the first inequality drops the infimum, the second is (92), and the last equality holds because $\mathbb{Q}^*$ attains the infimum at c^* . The equilibrium plan therefore maximizes the multiplier criterion among all plans with finite discounted second moments under $\mathbb{Q}^*$, linear or not, which is the “if” direction of Proposition 9.

C.11 Proof of Proposition 4 (Equilibrium invariant to Φ)

Part (i) is Proposition 3 applied consumer by consumer. The proposition fixes the interest rate at $R = \beta^{-1}$ and places $(\sigma_i, \hat{\beta}_i)$ on the locus (39).

For part (ii), because every consumer uses (8), the arguments of Section 2 apply. With (11) and (12) we obtain $\int_0^1 c_0^i di = Y$, and (9) with $\int_0^1 w_{t+1}^i di = 0$ gives $\int_0^1 c_t^i di = Y$ for all t . The goods market clears at every date and bonds are in zero net supply. Market clearing involves only the common rule (8) and the cross-section average of the shocks. Neither refers to Φ .

Parts (i) and (ii) use no assumption about which law generates realized endowments

beyond $\int_0^1 w_{t+1}^i di = 0$. Part (iii) follows because aggregate and cross-section objects are integrals of individual paths, and by part (i) the individual paths are the benchmark paths. Its statement about the cross-section distribution assigns probabilities to those paths under the approximating model.

C.12 Proof of Corollary 5

On the locus $\tilde{c}_t = (1 - \beta)\tilde{a}_t$ and $\tilde{a}_{t+1} = \zeta_i \tilde{a}_t + s \tilde{\epsilon}_{t+1}$ under the worst-case law, where $\tilde{\epsilon}_{t+1}$ has variance $\tilde{\Sigma} = \zeta_i$ (Proposition 3). Since $(1 - \beta)^2 s^2 = \alpha^2$, the one-period variance is $\zeta_i \alpha^2$ and the horizon- j variance accumulates the factors ζ_i^{2k} . The second equality in (59) is (56).

C.13 Proof of Corollary 1

By (38), $\sigma \alpha^2 \in (-(1 - \beta)^2, 0]$, and by (39), $\hat{\beta} = \beta(1 + \sigma \alpha^2 / (1 - \beta))$, so $\hat{\beta} \in (\beta(1 - (1 - \beta)), \beta] = (\beta^2, \beta]$. Equation (41) follows from $\beta^{-2} - 1 = (1 - \beta^2) / \beta^2$ and $\beta^{-1} - 1 = (1 - \beta) / \beta$.

C.14 Replication

Seven scripts in the replication files check this appendix and the numbers in the text and draw the figures, and an eighth runs them. The entry point `run_all.py` lists them in a manifest, runs them in order, and regenerates every table and figure. `make_figures.py` draws Figures 1 and 2. `check_decision_rules.py` also solves the stochastic Bellman equation (23) by Gauss–Hermite quadrature of (25). The value function with constant (82) satisfies it to 10^{-14} , $\tilde{c} = \Omega \tilde{a}$ attains the maximum, the worst-case law of ϵ_{t+1} has mean (79) and variance (37), and the worst-case covariance matrix of w_{t+1} is (86) to 10^{-10} . `check_welfare_loss.py` computes the value of each rule by summing the recursion for $\mathbb{E}_0 \tilde{a}_t^2$ directly rather than by (87), and maximizes over Ω' by golden-section search. The maximizer agrees with Ω_i^{twin} and the loss agrees with (68) to seven digits. `check_numbers.py` recomputes the numbers quoted in the text and the columns of Table 1 that do not come from simulation. `detection_tables.py` computes the detection error probabilities in Tables 1 and 3. It simulates 200,000 paths of length T from $\mu_0 = 0$ under each hypothesis, so the Monte Carlo standard error is about 0.001; `make_figures.py` uses 100,000 paths per point for Figure 2. The row of Table 3 for $\sigma_i = 0$ agrees with Table 1 at the two grid points the tables share, up to Monte Carlo noise. `check_reverse_loss.py` checks (70) against a numerical solution of (90), verifies that it does not depend on $\tilde{a}_0$ or α^2 , and locates the value of $\sigma_i / \underline{\sigma}$ beyond which the twin’s plan has no finite value under consumer i ’s multiplier criterion. `check_forecast_regression.py` simulates the regression of forecast errors on forecast revisions and checks (50).