Federal Reserve Bank of Minneapolis
Research Department

Reconciling Consumption Inequality
with Income Inequality[*]

Vadym Lepetyuk and Christian A. Stoltenberg

Working Paper 705

April 2013

ABSTRACT

The rise in within-group consumption inequality in response to the increase in within-group income inequality over the last three decades in the U.S. is puzzling to expected-utility-based incomplete market models. The two-sided lack of commitment models exhibit too little consumption inequality while the standard incomplete markets models tend to predict too much consumption inequality. We show that a model with two-sided lack of commitment and chance attitudes, as emphasized by prospect theory, can explain the relationship and can avoid the systematic bias of the expected utility models. The chance attitudes, such as optimism and pessimism, imply that the households attribute a higher weight to high and low outcomes compared to their objective probabilities. For realistic values of risk aversion and of chance attitudes, the incentives for households to share the idiosyncratic risk decrease. The latter effect endogenously amplifies the increase in consumption inequality relative to the expected utility model, thereby improving the fit to the data.

Keywords: Consumption Inequality; Prospect Theory; Limited Enforcement; Risk Sharing
JEL classification: E21, D31, D52

1 Introduction

The increase in within-group consumption inequality in response to the increase in within-group income inequality in the United States over the last three decades is hard to explain in standard macroeconomic models (Blundell et al., 2008, Krueger and Perri, 2006).[1] Krueger and Perri (2006) found that the standard incomplete markets models (Aiyagari, 1994) by two-fold overstated the increase in consumption inequality, whereas models with two-sided lack of commitment or debt-constrained markets (Alvarez and Jermann, 2000) by four-fold understated the rise in consumption inequality. Rationalizing the link between income and consumption inequality in a macroeconomic model is, however, essential for government policies, such as an optimal design of progressive income taxation (Krueger and Perri, 2011).

In macroeconomic research, economists have almost universally relied on expected utility theory as axiomatized by von Neumann and Morgenstern (1947) and Savage (1954).[2] According to this theory, the objective probabilities serve as the decision weights of the stochastic outcomes. Then, people's attitudes toward risk are captured solely by outcome sensitivity (curvature of utility), whereas people's sensitivity toward risk, such as their optimism or pessimism, is ignored. However, empirical studies of people's attitudes toward risk have found that chance attitudes are at least as important as the curvature of utility, and they have identified a critical regularity: people tend to be risk-seeking for low-probability high payoffs and risk-averse for low-probability low payoffs. Between the two effects the latter is stronger, that is people are relatively more pessimistic than optimistic (see e.g., Kahneman and Tversky, 1992, Wakker, 2010, Wu and Gonzalez, 1996). In this paper, we ask whether a proper accounting for the observed attitudes toward risk can reconcile the relationship between consumption and income inequality.

As the key innovative idea, we integrate chance attitudes into an environment with exogenous idiosyncratic income risk and two-sided lack of commitment (Kocherlakota, 1996). At the heart of our model is prospect theory for risk and uncertainty developed by Kahneman and Tversky (1992).

---

[1] Within-group inequality is referred to the differences across households that cannot be attributed to their fixed and observable characteristics such as occupation or education. Our entire analysis is devoted to the within-group component of inequality.

[2] The exceptions are Barberis et al. (2001), Hansen et al. (2006), who analyzed the implications of loss aversion for asset prices and for optimal policy decisions, respectively. Epstein and Zin (1989) and Weil (1990) disentangled risk aversion from intertemporal substitution. The application of non-expected utility theories to dynamic settings is a challenging task as we discuss in Section 2.

The theory provides a more general specification of preferences than expected utility theory by accounting for output sensitivity as well as for chance attitudes. In particular, the decision weights of stochastic outcomes are no longer given by their probabilities but depend crucially on the ranking of outcomes. Relatively high outcomes and relatively low outcomes that both appear with low probabilities are found to have higher decision weights, compared with their objective probabilities. Correspondingly, outcomes that are close to the mean are assigned lower decision weights.

To reduce their consumption risk, households in our model can voluntarily participate in risk-sharing arrangements with other households, or alternatively, they can walk away from the arrangement and consume their endowments as outside option. The participation constraints are important for two reasons. First and conventionally, they allow for a positive correlation between individual consumption and individual income. Second, and this is the novel implication here, the recursive and binary nature of participation constraints in our model render the optimal choices dynamically consistent. The choices made according to prospect theory can in principle be dynamically inconsistent, which is the major challenge for applying the theory to a dynamic setting (Machina, 1989).

We find that the equilibrium consumption risk is significantly different between the prospect theory model and the expected utility model. The main empirical finding of the paper is that our prospect theory model with relatively pessimistic households is capable to substantially better track consumption inequality in the U.S. from the year 1980 to the year 2003 – measured in a variety of ways – than the corresponding expected utility model.

Empirical studies on chance attitudes have simultaneously estimated the functional form of a weighting function, which captures the chance attitudes, and a utility function. The coefficient of relative risk aversion of the utility function is commonly found to be less than one and individuals tend to be relatively more pessimistic than optimistic (Wakker, 2010). For a coefficient of risk aversion less than one, if outcomes were weighed with the objective probabilities, then the risk aversion would be unreasonably low, and the response of consumption inequality to an increase in income inequality would be implausibly high. We analytically show that the relative pessimism captured in the decision weights can diminish the increase in consumption inequality and we lay down sufficient conditions for this case. In the model, relatively pessimistic high-income households place an additional weight on low-income states in the future which dampens the response of

consumption inequality. Overall, we demonstrate that the willingness of households to share a high income nevertheless decreases compared to the expected utility model. The latter effect is reflected in the optimal consumption allocation by a higher increase in consumption inequality.

The remainder of the paper is organized as follows. In the next section, we provide a review of prospect theory, and in Section 3, we lay down the economic environment. We highlight the effect of chance attitudes in a simplified environment in Section 4. We present our calibration strategy and our quantitative results in Section 5. The final section concludes the paper.

2   Review of Prospect Theory

Prospect theory for risk and uncertainty, as proposed by Kahneman and Tversky (1992), is motivated by examples that reveal that the expected utility theory, as axiomatized by von Neumann and Morgenstern (1947), has descriptive shortcomings. Perhaps the most famous example for risky choices is the Allais (1953) paradox, which illustrates that actual human behavior contradicts the independence axiom of expected utility also known as the sure-thing principle.[3] This principle postulates that the preference relationship between lotteries is not affected by addition or subtraction of a common consequence. The sure-thing principle is obeyed by the expected utility theory which makes the theory unsuitable for a comprehensive study of economic decisions under uncertainty and risk.[4]

Prospect theory provides a theoretical description for the behavior of economic agents under risk and uncertainty. Restricted to the domain of gains, the theory is equivalent to the rank-dependent utility theory that was first axiomatized by Quiggin (1982) under the name of *Anticipated Utility Theory*. Under the theory, individual attitudes toward risk are captured not only by the curvature of utility, as under the expected utility theory, but also by decision weights derived from a nonlinear function of outcome probabilities.

Consider a lottery $(\{x_j\}, \{p_j\})$ that pays x_j with probability p_j. The prospect theory value

---

[3] For choices under uncertainty (unknown probabilities), Ellsberg (1961) also provides strong evidence that cannot be captured by expected utility theory.

[4] Prelec (1990) reported a typical finding that violates the sure-thing principle. The vast majority of individuals prefer gamble $S = (0.02, \$20,000)$, which yields \$20,000 with probability 0.02, over gamble $R = (0.01, \$30,000)$, which yields \$30,000 with probability 0.01. The majority also prefers $R' = (0.01, \$30,000; 0.32, \$20,000)$ over $S' = (0.34, \$20,000)$. The sure-thing principle is violated because S' and R' are designed by adding a 0.32 chance to win \$20,000 to both of the original gambles, S and R.

of the lottery for a given probability weighting function, $w : [0,1] \rightarrow [0,1]$ with $w(0) = 0$ and $w(1) = 1$ is defined as follows. First, the lottery outcomes are ranked from the worst to the best, $x_1 \leq x_2 \leq ... \leq x_n$. Second, for each outcome x_j the decision weight π_j is computed as the marginal w-contribution of the outcome probability to the rank:

$$\pi_j = w\left(\sum_{k=j}^{n} p_k\right) - w\left(\sum_{k=j+1}^{n} p_k\right),$$

where the rank of an outcome is the probability of yielding something that is strictly better than this outcome. Finally, the prospect theory value is computed as the weighted sum of the outcomes

$$u^{PT}\left(\{x_j\}, \{p_j\}\right) = \sum_{j=1}^{n} \pi_j x_j.$$

The weighting function is commonly found to have an inverse-S shape (see Figure 7 in the Appendix). For the concave part of the weighting function, an increase in the rank, which corresponds to the decrease in the argument of the function, increases the decision weight and thus generates optimism and risk seeking. Similarly, convexity results in pessimism and risk aversion. In the major part of its domain, the weighting function is below the 45-degree line and convex, thus reflecting the empirical realistic property of relative pessimism (Wakker, 2010).[5]

Non-expected utility theories like prospect theory have been applied to analyze dynamic decisions. A fundamental question in dynamic settings is about a way in which preferences of households are updated with the flow of information. Non-expected utility theories are commonly axiomatized to respect one of the two principles: dynamic consistency and consequentialism (see e.g., Eichberger et al., 2007, Machina, 1989). Dynamic consistency requires that substrategies of optimal strategies are optimal; consequentialism demands that only outcomes that are still possible can matter for the updated preferences (see Appendix A.1 for details). Hammond (1989) demonstrated that if both dynamic consistency and consequentialism are preserved together with other standard assumptions, the preferences should have the expected utility form.

---

[5] Coming back to the example by Prelec (1990), the choices of the individuals can be explained by cumulative prospect theory and by a typical inverse-S shaped weighting function, which is concave for lower values of the argument such that $w(0.02) - w(0.01) > w(0.34) - w(0.33)$. For the evidence on individual choices under risk that are captured by the convex part of the weighting function, see e.g., Kahneman and Tversky (1979).

No consensus has so far emerged in the decision-theoretical literature on the most plausible axiomatization. Machina (1989) argues to drop consequentialism and to preserve dynamic consistency, while other authors like Eichberger et al. (2007) tend to preserve consequentialism rather than dynamic consistency. The empirical literature provides mounting evidence that the majority of people violate dynamic consistency rather than consequentialism, both for the case of risk and for the case of uncertainty (e.g., Busemeyer et al., 2000, Cubitt et al., 1998, Dominiak et al., 2012). Hence, we rely on a preference specification that preserves consequentialism but not dynamic consistency. Similar to Zimper (2012), we address dynamic inconsistency by considering sustainable allocations. The notion of sustainability requires that the incentives to reverse original plans at every decision node are *ex-ante* explicitly taken into account.

3 Economic Environment

In this section, we describe our economic environment. As in Thomas and Worrall (1988), we consider voluntary risk-sharing arrangements. Following Kahneman and Tversky (1992), we develop a model based on prospect theory for risk and uncertainty (PT) as the decision-theoretical framework.

Consider an exchange economy with a continuum of *ex-ante* identical households. The time is discrete and indexed by t from zero onward. The income of household i is given by a stochastic process $\{y_t^i\}_{t=0}^{\infty}$, where the set of possible income realizations in each period is finite $y_t^i \in Y_t \equiv \{y_{t1}, ..., y_{tn}\}$. As in Krueger and Perri (2006, 2011), the income is independent across households and evolves across time according to a first-order Markov chain with time-varying exogenous states and time-invariant transition probabilities. We consider time-invariant transition probabilities that imply a unique stationary distribution and assume that the initial distribution of the Markov chain is given by the stationary distribution. While the income is risky for a household, it is not uncertain, i.e., the transition probabilities of the Markov chain are public knowledge.

As the principal novel feature, we take into account the empirically observed inclination of households toward optimism and pessimism as emphasized by prospect theory (Kahneman and Tversky, 1992). Consider an agent's stochastic consumption allocation $c \equiv \{c_t\}_{t=0}^{\infty}$ with a finite number of realizations in each period. We denote the realization j of period t consumption as c_{tj}

and the unconditional probability of the realization as p_{tj}. The utility of the consumption allocation is calculated according to prospect theory as follows. First, within each period, the realizations are ordered from the lowest to the highest according to the period utility function $u : \mathbb{R}^+ \to \mathbb{R}$, which is assumed to be increasing and strictly concave. Then, the period-t allocation is evaluated as a weighted sum of one-period utilities

$$u_t^{PT}(c_t) = \sum_{j=1}^{n} \pi_{tj} u(c_{tj}),$$

where the decision weights are defined as $\pi_{tj} = w(\sum_{k=j}^{n} p_{tk}) - w(\sum_{k=j+1}^{n} p_{tk})$, and $w : [0,1] \to [0,1]$ is the probability weighting function. The weighting function satisfies $w(0) = 0$ and $w(1) = 1$ and is assumed to be strictly increasing. Finally, the utility of consumption allocation is calculated as the following time-separable discounted sum:

$$\text{PT}(c) = \sum_{t=0}^{\infty} \beta^t u_t^{PT}(c_t), \tag{1}$$

where $0 < \beta < 1$ is the time discount factor. The *ex-ante* utility derived by prospect theory, $\text{PT}(c)$, naturally serves as the social welfare measure given that before any realization of uncertainty the households are identical. Expected utility (EU) is a particular case of the prospect theory that emerges for the identity weighting function $w(p) = p$.

The discounted sum in (1) is over one-period utilities evaluated according to prospect theory. That means that each period is evaluated separately, treating the prospect theory value of the other periods as certainty equivalents. The procedure thus preserves consequentialism. However, in general, the preferences defined by (1) do not satisfy the law of iterated expectations and thus, dynamic consistency is not preserved.[6] In principle, modeling of households' decisions under risk requires a probability space for which a (possibly non-additive) probability measure is defined over all possible consumption streams. Such construction is standard for additive probability measures of expected utility theory but much more complicated for non-additive probability measures of prospect theory. We avoid these complications by the additive specification of a prospect theory utility function (1).

---

[6] In some cases, dynamic consistency is still preserved, for example, when the weighting function is the identity function or when the new information is irrelevant for the future as it is in case of i.i.d. income process.

As in Zimper (2012), we restrict attention to bounded memory choices in which households' information set only contains their current income, but not their income history. The assumption is consistent with the time-invariant weighting function as in Cecchetti et al. (2000) and Abel (2002).[7]

We analyze risk-sharing possibilities under two-sided lack of commitment by introducing participation constraints. The constraints characterize the intertemporal trade-off between the current-period consumption and the value of risk sharing provided by the arrangement in future periods. Without eliminating the intertemporal trade-off, the limited-memory assumption implies that the current risk-sharing transfers do not depend on transfers received in the past and that current consumption is measurable only with respect to current income.[8] For expected utility, the resulting history-independent insurance contracts are modeled by Kimball (1988), Coate and Ravallion (1993), Ligon et al. (2002), and by Lepetyuk and Stoltenberg (2013b).[9]

A risk-sharing arrangement is sustainable if each household, after observing the current period income, at least weakly prefers to follow the arrangement to defect into autarky. If a household deviates into autarky, then the household is excluded from risk sharing with other households for T periods and in each period consumes only its income. After T periods, the household may reengage in the risk-sharing agreement. The notion of sustainable risk-sharing arrangements is formally introduced by the following definition.

Definition 1 *A consumption allocation $\{\{c_t^i\}_{t=0}^\infty\}$ is sustainable if*

(i) for each household i in each period t, the consumption allocation is weakly preferable to the outside option

$$\text{PT}_t(\{c_{t+\tau}^i\}_{\tau=0}^\infty) \geq \text{PT}_t(\{c_{at,t+\tau}^i\}_{\tau=0}^\infty) \quad \forall i, t, \tag{2}$$

where the allocation of $\{c_{t+\tau}^i\}_{\tau=0}^\infty$ on the left-hand side is prospect theory utility conditional

---

[7] Even if the parameters of income process were unknown to the households, the lack-of-memory assumption would prevent them from learning and would result in a time-invariant weighting-function. Zimper (2012) further shows that the bounded memory assumption is the key to establish a unique stationary equilibrium price function in a Lucas-fruit tree economy with a neo-additive weighting function.

[8] Because the households lack memory to condition their actions on past income, in the corresponding social planner's problem the planner faces the same information structure as the households and it cannot trustfully commit to promise-keeping. In Lepetyuk and Stoltenberg (2013a), we characterize efficient allocations that are contingent on the entire income history.

[9] For risk-sharing contracts that are based on the entire history of income and their explanatory power, see Alvarez and Jermann (2000), Cordoba (2008), Kocherlakota (1996).

on all the information available in period t[10]

$$\text{PT}_t(\{c_{t+\tau}^i\}_{\tau=0}^\infty) = \sum_{\tau=0}^\infty \beta^\tau u_t^{PT}(c_{t+\tau}^i),$$

and the allocation of $\{c_{at,t+\tau}^i\}_{\tau=0}^\infty = (\{y_{t+\tau}^i\}_{\tau=0}^{T-1}, \{c_{t+\tau}^i\}_{\tau=T}^\infty)$ on the right-hand side is evaluated according to

$$\text{PT}_t(\{c_{at,t+\tau}^i\}_{\tau=0}^\infty) = \sum_{\tau=0}^{T-1} \beta^\tau u_t^{PT}(y_{t+\tau}^i) + \sum_{\tau=T}^\infty \beta^\tau u_t^{PT}(c_{t+\tau}^i);[11] \ \text{and}$$

(ii) the allocation is resource-feasible

$$\int c_t^i \, \mathrm{d}i \le \int y_t^i \, \mathrm{d}i \quad \forall t. \tag{3}$$

The participation constraints (2) are the key element of the definition. In each period, the households decide to sustain an insurance arrangement after their current income has realized. Their decision weights for all possible further realizations are based on the true objective conditional probabilities of these realizations. Thus, a consumption allocation is sustainable if the utility of the consumption allocation computed according to prospect theory is at least as large as the utility of the outside option.

Among the sustainable allocations, we are interested in the optimal allocation that is defined as follows.

Definition 2 *A socially optimal arrangement under voluntary participation is a consumption allocation that provides the highest social welfare among the set of sustainable allocations.*

Despite the dynamic inconsistency in preferences, the decisions to sustain the socially optimal arrangement are dynamically consistent. The participation constraints (2) postulate that every period represents a decision node at which a household makes a binary decision between sustaining the arrangement or deviating to the outside option. Because the problem of computing the

---

[10] Contrary to the *ex-ante* prospect utility in (1) where the decision weights are based on unconditional probabilities, here the decision weights are based on conditional probabilities *after* period-t income has realized.

[11] The expression reflects that an agent after T periods weakly prefers a sustainable allocation to autarky.

socially optimal arrangement is recursive, the decisions are time invariant, $c_t = c(y_t)$, resulting in dynamically consistent choices.

Doubts remain about the most appropriate preference specification based on prospect theory in dynamic settings (Wakker, 2010). As an alternative specification of preferences, we consider the backward induction method applied, for example, in Post et al. (2008). The method preserves dynamic consistency by construction and is based on a recursive evaluation in which chance nodes are subsequently replaced by their certainty equivalents while folding back period by period. Let $Q = \{q_{jk}\}$ be a time-invariant transition matrix whose elements are conditional probabilities of next period's consumption state c_k given current period consumption state c_j. Conditional on all the information available in period t and given $c_t = c_j$ is known, the prospect theory utility of an allocation of $\{c_{t+\tau}\}_{\tau=0}^{\infty}$ is defined recursively by

$$\mathrm{PT}^{BI}(\{c_{t+\tau}\}_{\tau=0}^{\infty}|c_t = c_j) = u(c_j) + \beta \sum_{k=1}^{n} \pi_{jk} \, \mathrm{PT}^{BI}(\{c_{t+\tau+1}\}_{\tau=0}^{\infty}|c_{t+1} = c_k), \qquad (4)$$

where π_{jk} are decision weights computed by applying the prospect theory weighting function to the transition probabilities q_{jk}. The *ex-ante* prospect theory value of a stochastic allocation $c = \{c_t\}_{t=0}^{\infty}$ is then given by the following expression

$$\mathrm{PT}^{BI}(c) = \sum_{j=1}^{n} \pi_j \, \mathrm{PT}^{BI}(c|c_0 = c_j), \qquad (5)$$

where π_j are prospect theory decision weights that are based on stationary probabilities p_j of the time-invariant transition matrix. The value serves as the measure of social welfare. The participation constraints (2) of household i in period t are evaluated according to (4). The preference specifications (1) and (5) in general lead to different evaluations of a consumption allocation. There are two important exceptions. First, when the probability weighting function is the identity function, they both collapse to the expected utility. Second, if the income is an independent and identically distributed random process, the two evaluation methods coincide. The latter is the case we analyze in the next section.

9

4 Analytical Results

Empirical studies on within-group inequality find that consumption inequality is smaller than income inequality, and that they are positively correlated. In an expected utility model, Krueger and Perri (2006) demonstrated that models with two-sided lack of commitment tend to underpredict the increase in consumption inequality in response to the recent rise in income inequality in the U.S. from the year 1980 to the year 2003. In this section, we show that our model has the potential to reconcile the changes in consumption and income inequality found in the data when (1) households are relatively more pessimistic than optimistic, and (2) incentive constraints matter.

To illustrate the main mechanism underlying the result, we abstract in this section from a number of features. First, we assume that the income process consists of three time-invariant states $y_h = \bar{y} + \sigma_y \sqrt{3/2}$, $y_m = \bar{y}$, and $y_l = \bar{y} - \sigma_y \sqrt{3/2}$, where $\sigma_y > 0$ is the standard deviation of the income process. The states are equally likely and the income realizations are independent across time. When attention is restricted to evaluation of i.i.d. processes, the preference specifications (1) and (5) coincide and exhibit dynamical consistency for any weighting function. Moreover, the decision weights in the participation constraints are time invariant. Second, we suppose that once an agent deviates to autarky, that agent is precluded from risk sharing forever.

Without loss of generality we employ a neo-additive weighting function.[12] The neo-additive weighting function is linear in the interior of its domain $w(p) = \lambda + (1 - \lambda - \gamma)p$ for $p \in (0, 1)$ and is discontinuous at the endpoints, $w(0) = 0$ and $w(1) = 1$. The decision weights associated with the weighting function are the following: $\pi_h = \lambda + (1 - \lambda - \gamma)/3$ for the highest state, $\pi_l = \gamma + (1 - \lambda - \gamma)/3$ for the lowest state, and $\pi_m = (1 - \lambda - \gamma)/3$ for the middle state. The parameter $\lambda \geq 0$ denotes the degree of optimism, and $\gamma \geq 0$ denotes the degree of pessimism with $\lambda + \gamma \leq 1$.[13] Consumption inequality and income inequality are measured as the standard deviations of the consumption (σ_c) and income allocation (σ_y), respectively.

With the three time-independent income realizations, the problem can be formulated as follows.

---

[12] In the case of three independent and identically distributed states, the preference specification in (1) can be equivalently characterized for any weighting function by two neo-additive parameters: the degree of pessimism, γ, and the degree of optimism, λ, with $w(1/3) = (1 + 2\lambda - \gamma)/3$ and $w(2/3) = (2 + \lambda - 2\gamma)/3$.

[13] When the neo-additive weighting function is applied to models with uncertainty, $\zeta \equiv \lambda + \gamma$ represents the degree of ambiguity, which is the lack of confidence the households have in the expected utility measure (Ellsberg, 1961).

The social welfare is

$$\text{PT}(c) = \frac{1}{1-\beta} \left\{ \lambda \max[u(c_l), u(c_m), u(c_h)] + \gamma \min[u(c_l), u(c_m), u(c_h)] \right.$$
$$\left. + (1 - \lambda - \gamma)\frac{1}{3}[u(c_l) + u(c_m) + u(c_h)] \right\}, \quad (6)$$

where c_h, c_m and c_l are the consumption of households with current-period income y_h, y_m, and y_l, respectively. The participation constraints for households with high, medium, and low current-period income are

$$u(c_j) + \beta \text{PT}(c) \geq u(y_j) + \beta \text{PT}(y) \quad \forall j \in \{l, m, h\} \quad (7)$$

where

$$\text{PT}(y) = \frac{1}{1-\beta} \left[\frac{1 + 2\lambda - \gamma}{3} u(y_l) + \frac{1 - \lambda - \gamma}{3} u(y_m) + \frac{1 - \lambda + 2\gamma}{3} u(y_h) \right],$$

and the resource constraint is

$$c_l + c_m + c_h = y_l + y_m + y_h. \quad (8)$$

An optimal consumption allocation $\{c_l, c_m, c_h\}$ maximizes the social welfare (6) subject to the participation constraints (7), and the resource feasibility constraint (8).

The optimal allocation exists but may not be unique.[14] We restrict our attention to empirically plausible allocations such that households with higher income consume more, i.e., $c_l \leq c_m \leq c_h$. Without this restriction, it is straightforward to deduce that, for example, if $\{c_l, c_m, c_h\}$ with $c_l > c_h$ is an optimal allocation, then $\{\hat{c}_l, c_m, \hat{c}_h\}$, where $\hat{c}_l = c_h$ and $\hat{c}_h = c_l$, is also optimal.

The first-best allocation maximizes the ex-ante welfare (6) subject to the resource feasibility constraint (8). The optimal allocation coincides with the first-best when the participation constraints (7) are not be binding. With optimism and pessimism, the first-best allocation does not necessary exhibit perfect risk sharing. If households are very pessimistic and pay particular attention to the lowest consumption realization, they prefer to smooth consumption across states. On the other hand, if households are optimistic and risk-seeking, the optimal consumption allocation exhibits inequality. In the following proposition, we formally characterize the first-best allocation.

---

[14] The existence follows from continuity of the objective function and compactness of the set of sustainable allocations.

Proposition 1 *If $\gamma < 2\lambda$, the first-best allocation exhibits inequality. Otherwise, the first-best allocation is perfect risk sharing.*

Proof. At the first-best allocation, $c_l = c_m$. If not then $c_l < c_m$ contradicts the first order condition $(1 - \lambda - \gamma)u'(c_m) \geq (1 - \lambda + 2\gamma)u'(c_l)$. If $\gamma < 2\lambda$, the first order condition for the optimal choice between $c_l = c_m$ and c_h, $(1 + \lambda - \gamma/2)u'(c_m) \geq (1 + 2\lambda - \gamma)u'(c_h)$ implies that $c_l = c_m < c_h$. If $\gamma \geq 2\lambda$, then $c_l = c_m = c_h$ is the optimal allocation. If not then $c_l = c_m < c_h$ contradicts the first order condition $(1 + \lambda - \gamma/2)u'(c_m) = (1 + 2\lambda - \gamma)u'(c_h)$ that has to hold with equality. ∎

In the expected utility model, the first-best allocation always exhibits perfect risk sharing. In the prospect theory model, however, the first-best allocation exhibits inequality and risk sharing is partial unless households are very pessimistic – more pessimistic than commonly found in experimental studies. The possibility of inequality in the first-best allocation by itself suggests that this allocation may be used to match the positive consumption inequality observed in the data. However, the consumption allocation at the first best is independent of income inequality which contradicts the fact that both inequalities are positively correlated. Therefore, the empirically relevant cases are those with bindings participation constraints.

To explain the changes in consumption inequality found in the data, we proceed with constrained-efficient allocations. As a first step, we show that consumption inequality can be smaller than income inequality only when the participation constraints of low-income households are not binding at the optimal allocation.[15] This property resembles the findings under expected utility.

Proposition 2 *Consider the optimal consumption allocation. Consumption inequality can be smaller than income inequality only if the participation constraints for low-income households are not binding.*

Proof. By contradiction, if at the optimal allocation $\{\hat{c}_l, \hat{c}_m, \hat{c}_h\}$ the participation constraints for low-income households were binding, then from these constraints we would obtain that $\hat{c}_l \leq y_l$. Because the decision weight for the middle state is smaller than for the low state, the participation constraints for the middle state should also be binding, which implies $\hat{c}_m \leq \bar{y}$. By resource feasibility, $\hat{c}_h \geq y_h$. Hence, consumption inequality would be greater than income inequality. ∎

---

[15] We say that a constraint is binding when the Lagrange multiplier associated with the constraint is positive.

In the data, consumption inequality is smaller than income inequality. Thus, the two empirically relevant possibilities that we are left with are the case of binding participation constraints of high-income households and the case when the constraints of middle-income households are binding as well. To illustrate that our prospect theory model has the potential to reconcile consumption with income inequality, we will focus on the case when only participation constraints of high-income households are binding. In this case, the optimal consumption allocation is characterized by $c_l = c_m < c_h$. This relationship follows from efficient risk-sharing at the lower end of the income distribution and implies that $\mathrm{PT}(c) > \mathrm{PT}(y)$. From the binding participation constraint of high-income households, it is straightforward to obtain that $c_h = \bar{y} + \sqrt{2}\sigma_c < y_h$ and $\sigma_c < \sqrt{3}/2\sigma_y$. Hence, consumption inequality is smaller than income inequality. In the next proposition, we characterize the dependence of the optimal allocation on the degrees of optimism and pessimism.

Proposition 3 *Consider the optimal consumption allocation. If only the participation constraints for high-income agents are binding, consumption inequality is decreasing in the degree of pessimism, γ, and increasing in the degree of optimism, λ.*

The proof can be found in Appendix A.2. The finding resembles the intuition for the first-best allocation. While pessimistic households dislike risk, optimistic households appreciate it exactly in the same way.

In the main proposition, we will show that an increase in the degree pessimism diminishes the response of consumption inequality to a change in income inequality. As an intermediate step, we state a condition in the following lemma that ensures a positive correlation between consumption inequality and income inequality.

Lemma 1 *If the participation constraints under the socially optimal arrangement are binding only for high-income agents, consumption inequality is increasing with income inequality if and only if the outside option for high-income households is increasing with income inequality, i.e.*

$$\left(1 + \frac{\beta}{1-\beta}\frac{1 + 2\lambda - \gamma}{3}\right)u'(y_h) - \left(\frac{\beta}{1-\beta}\frac{1 - \lambda + 2\gamma}{3}\right)u'(y_l) > 0.$$

The proof can be found in Appendix A.3. The intuition is the following. The binding participation constraints for high-income households imply that these households obtain the same utility from the

13

outside option as from the socially optimal arrangement. With the increase in income inequality, the high-income households obtain a higher utility if and only if the value of their outside option also increases with the income inequality. The increase in utility for high-income households results in higher consumption inequality.

In the expected utility model with logarithmic preferences (constant relative risk aversion of one), Krueger and Perri (2006) found that two-sided lack of commitment or debt-constrained markets models tend to underpredict the increase in consumption inequality compared to the data. If the (constant) relative risk-aversion coefficient were smaller the increase in consumption inequality would be larger, thereby guiding the expected utility model in the right direction. However, as documented by Barsky et al. (1997) and more recently by Meyer and Meyer (2005), it is implausible to postulate the coefficient to be below one unless risk aversion is preserved in another way, i.e., by taking into account chance attitudes as in the prospect theory.[16] Nonetheless, for coefficients of relative risk aversion smaller than one, the expected utility model can even overpredict the increase in consumption inequality. When this is the case, an increase in pessimism that can be captured by prospect theory has a diminishing effect on consumption inequality. In the following proposition, we provide conditions for this result.

Proposition 4 *Assume the utility function exhibits constant relative risk aversion, a degree of pessimism is sufficiently large, $\gamma \geq \bar{\gamma}$ and the value of the outside option for high-income households is increasing in income inequality. The rate of change in consumption inequality with respect to the increase in income inequality is decreasing in the degree of pessimism, γ.*

The proof can be found in Appendix A.4. The intuition for this result is as follows. Consider an increase in the variance of income. As the dominant effect, the high-income state increases, which increase the value of the outside option for high-income households. The dominant effect is partially offset by a decrease in the ex-ante value of the outside option. With an increase in the degree of pessimism, households put additional weight on the low-income state in the future amplifying the decrease in the ex-ante value of the outside option for high-income households. Pessimism thus attenuates the increase in the continuation value of the outside option. The lower increase must

---

[16] Viewed from the perspective of the prospect theory, estimates of the constant coefficient of relative risk aversion based on expected utility models are biased upward because of omitted variables.

also be reflected in the optimal consumption allocation and the increase in consumption inequality is dampened. The condition $\gamma \geq \bar{\gamma}$ ensures that an increase in income inequality leads to a less than one-for-one increase in consumption inequality.

In this section, we have shown that the prospect theory model with relative pessimism has qualitatively the potential to constitute the missing link between income and consumption inequality. However, whether the prospect theory model can indeed reconcile the inequalities for realistic values of pessimism, optimism, and relative risk aversion is a quantitative question that we analyze in the following section.

5   Numerical Results

Our prospect theory model explains substantially better than the expected utility model the relationship between income and consumption inequality as observed in the United States from 1980 to 2003. In this section, we describe the data and model calibration, and we numerically characterize the optimal risk-sharing arrangement and confront it with the data.

5.1   Data

We calibrate our model to match the U.S. economy for the years 1980–2003. We use the public dataset prepared by Krueger and Perri (2006), which is based on quarterly data from the Consumer Expenditure Interview Survey (CEX). The dataset consists of cross-sectional consumption and income series that have been aggregated from the original survey data across a number of income and expenditure categories.[17]

The differences in individual income and consumption can be in part attributed to observable characteristics of households such as education, experience, occupation, race, and region of residence. The inequality attributed to these characteristics is known as between-group inequality (Katz and Autor, 1999). We consider the between-group income differences as not insurable. On the contrary, the residual component that cannot be explained by publicly observable characteristics can potentially be insured. Our analysis is devoted entirely to this component, implying

---

[17] Based on the Panel Study of Income Dynamics for 1978–1992 and on the Consumer Expenditure Interview Survey for 1967–2006, Blundell et al. (2008) and Heathcote et al. (2010) reported similar statistics on consumption inequality. Aguiar and Bils (2011), Attanasio et al. (2007, 2012) pointed out that the Consumer Expenditure Interview Survey might be plagued with non-classical measurement errors.

that consumption and income inequality in our model correspond to the within-group inequality of consumption and income as observed in the data.

5.2 Calibration

Our benchmark calibration is designed to highlight the differences between our model and the debt-constrained markets model in Krueger and Perri (2006). Therefore, we set a number of corresponding parameters to the same values. In particular, we consider a period utility function that exhibits constant relative risk aversion and we set the time-discount factor to $\beta = 0.96$. One period of the model corresponds to one year.

We model the income as a stochastic Markov process with nine time-varying states. It implies that the income levels that correspond to each state vary over time, while the probability of transition from one state to another is constant across time. By varying the income levels, we set the cross-sectional variance of the income process in each year to the variance observed in the data. Employing the autoregressive coefficient of the income process $\rho = 0.8014$ reported by Krueger and Perri (2011), we then compute time-invariant transition probabilities of the Markov chain following Tauchen (1986). Each period, the households only know the current realizations of income but neither the past (due to the lack of memory) nor the future realizations.[18] Consumption and income inequality are calculated as cross-sectional variances of the consumption and the income distribution, respectively.

In our baseline, we follow Kahneman and Tversky (1992) in choosing the functional form of the probability weighting function and the utility function which are the following:

$$w(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{1/\delta}}, \qquad u(c) = \frac{c^{1-\sigma} - 1}{1 - \sigma}. \tag{9}$$

The weighting function supports the property of relative pessimism commonly found in the experimental literature. We simultaneously set the parameter of the weighting function and the constant coefficient of risk aversion of the utility function according to values estimated in laboratory studies specifically designed to capture decisions under risk. As our baseline, we employ the averages

---

[18] In Section 5.4, we employ alternatively the perfect foresight assumption as in Krueger and Perri (2006) that the changes in income inequality observed over the entire sample are known to the households at the beginning of the sample.
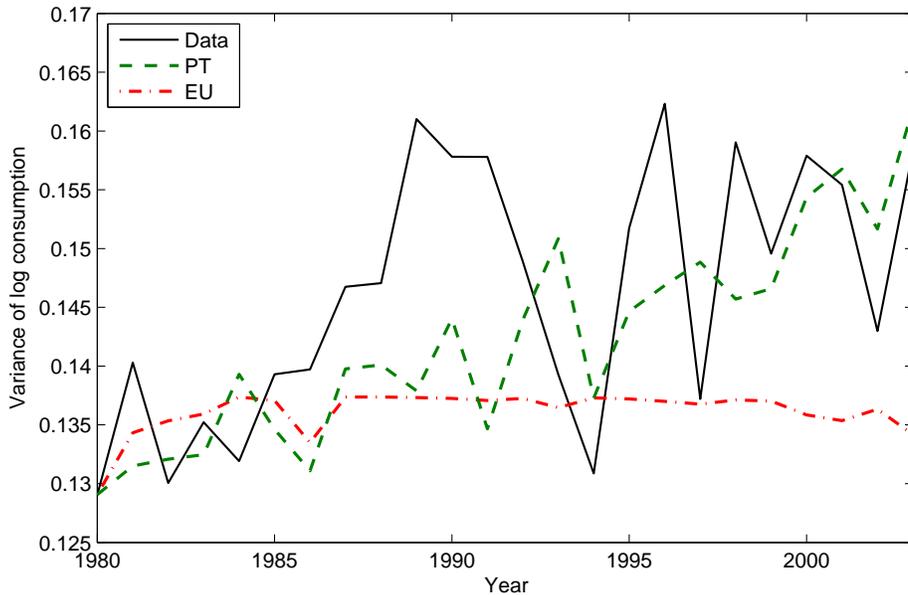
Figure 1: Consumption inequality over time in the data (solid) versus the prospect theory model (PT; dashed) and the expected utility model (EU; dash-dotted)

of the estimated values reported in seven studies (provided in Table 5 in the Appendix), including Kahneman and Tversky (1992) and Wu and Gonzalez (1996). The values are 0.74 for the Tversky–Kahneman parameter, δ, and 0.50 for the coefficient of relative risk aversion of the utility function, σ.

We match the consumption inequality in the initial year of our sample by calibrating the duration of exclusion from the risk-sharing arrangement, T. For the prospect theory model, the duration is equal to 10 years. For the expected utility model ($\delta = 1$), we employ the logarithmic period utility function, and the duration of exclusion is $T = 20$ years.

5.3 Model Simulation

In Table 1 and Figure 1, we report our principal numerical results on the capacity of the prospect theory model and the expected utility model in explaining the cross-sectional variance of consumption over time. As displayed in Figure 1, the expected utility model systematically understates the consumption inequality compared to the data. Between the years 1980 and 2003, the increase in consumption inequality that the expected utility model predicts is more than five times smaller than the actual increase. Though the lack of history-dependence that follows from the bounded

Table 1: Expected utility and prospect theory model: baseline results, 1980–2003

	EU model	PT model
Weighting function parameter, δ	1	0.74
Utility function parameter, σ	1	0.50
Increase in consumption inequality	0.51%	3.29%
Mean-square error, normalized	1.95	1
Akaike Information Criterion, AIC	-125.71	-139.72
Bayesian Information Criterion, BIC	-124.53	-137.36
Auto-correlation coefficient	0.77	0.71

memory assumption leads to contracts that exhibit less risk sharing, this property alone does not improve the fit to the data. The picture changes in our prospect theory model when we allow for non-linear probability weighting. The model is tracking particularly well the increase in consumption inequality in 1994–2003, while neither model is capturing the surge in consumption inequality in the late 1980's.

As the first novel result here, the prospect theory model is better at capturing the change in consumption inequality. Over the considered period, the model predicts an increase of 3.29 percentage points in the variance of log consumption which is close to the actual increase of 2.93 percentage points found in the data.[19] This observation is also reflected in the statistics that measure the explanatory power of both models for consumption inequality over all years in the sample (see Table 1). The mean squared error (MSE) in the expected utility model is 95 percent higher than in the prospect theory model. Furthermore, both information criteria, the Bayesian information criterion (BIC) and the Akaike information criterion (AIC) indicate a significant increase in fit for the prospect theory model. The magnitude of the increase indicates that the expected utility model fails to explain some substantial structural variations in the data.

To better understand the superior performance of the prospect theory model, consider the following two counterfactual experiments. First, assume that all parameters are given by our baseline prospect theory calibration except the weighting function parameter, δ, which is set to unity, thus eliminating pessimism and optimism. Such a model overshoots consumption inequality in all years. In particular, for 2003 consumption inequality is 25 percent higher than in the data.

---

[19] The prospect theory model is preferable to the expected utility model even if the increase in consumption inequality were higher as has been recently suggested by Aguiar and Bils (2011), Attanasio et al. (2012).

Alternatively, let the parameters be given by the prospect theory model except for the constant coefficient of risk aversion of the utility function, σ, which is fixed to one. In this counterfactual experiment, the model undershoots consumption inequality in all years, for example, in 2003, consumption inequality is five percent lower than in the data.

The prospect theory implies less risk aversion than the expected utility theory. The consumption equivalent of ax-ante utility of the optimal allocation under prospect theory is higher that the consumption equivalent of the same allocation under expected utility theory.

In the cross-section dimension, the optimal allocation under prospect theory prescribes the same consumption to households with lower levels of current-period income. Across all years in our sample, households with the five lowest income realizations out of nine receive the same level of consumption. The consumption smoothing at the lower end is also the feature of the expected utility model but only if the income process were not persistent. When the income process is persistent, the optimal allocation under expected utility prescribes gradually increasing levels of consumption with an increase in income, reflecting heterogeneous expectations of binding high-income households. In particular, complete redistribution at the lower end of income distribution is prevented by the binding constraints of high-income households. With pessimism, high-income households assign higher decision weights for lower consumption levels than under expected utility thus sustaining additional redistribution at the lower end of the consumption allocation.

In the cross-time dimension, our model channels income persistence into similar consumption persistence. This is a result of the lack-of-memory assumption which implies that consumption is a Markov process with the same transition probabilities as the income process. As a beneficial consequence, we find that the consumption persistence in the model is in line with the data. In the CEX data, the sample average of the annual autocorrelation coefficient of consumption is 0.71. In our baseline prospect theory model, the average autocorrelation coefficient is also 0.71. The apparent match is not a result of any calibration. The average autocorrelation coefficient of consumption in the expected utility model is 0.77, which is also reasonably close to the one observed in the data.

As can be seen from Figure 3, both the expected utility model (EU) and the prospect theory model capture the increasing trend in the insurance ratio, $1 - \text{var}_i(\ln c_t^i)/\text{var}_i(\ln y_t^i)$; however, the expected utility model systematically overstates the degree of private insurance relative to the

data over (almost) the entire sample. The prospect theory model does not suffer from the bias. The difference in the estimates of private insurance is significant for optimal tax policies. As pointed out by Krueger and Perri (2011), one motivation for redistribution policies implemented by a progressive tax system is to provide public insurance against the idiosyncratic income risk. Broer (2011) showed that a redistribution policy tends to increase welfare if the level of private insurance is high. For our analysis, it implies that the expected utility model with a high level of private insurance could lead policymakers to the wrong conclusion that more public insurance will be welfare improving, thereby justifying a higher progression in the tax scheme.

5.4 Robustness

In this subsection, we conduct several robustness exercises across different dimensions. First, we consider the Gini coefficient as an alternative measure of consumption and income inequality. Second, we use smoothed time series to analyze the implications of prospect theory and the expected utility model for the consumption inequality trend as in Krueger and Perri (2006). Third, we employ a specification of the income process that comprises both, a transitory and a persistent component as estimated by Storesletten et al. (1999, 2004). Fourth, we employ the perfect foresight assumption as in Krueger and Perri (2006), implying that the evolution of income inequality over the entire sample was perfectly known to the households at the beginning of the sample. Finally, we consider backward induction as an alternative specification of preferences.

Gini coefficient of consumption inequality As an alternative measure of inequality, we consider the Gini coefficient of consumption inequality. For this exercise, we keep the parameters of our baseline calibration and adjust T in both models to match the Gini coefficient in the year 1980. In Figure 5 in the Appendix, we plot the evolution of the Gini coefficient of consumption in the data versus the prospect theory model and the expected utility model. The primary message is similar to the analysis in which the inequality is measured by the variance of the logarithm of consumption. The prospect theory model captures substantially better than the expected utility model the changes in the Gini coefficient. In particular, the mean squared error in the prospect theory model is three times smaller than in the expected utility model (see Table 6).

Table 2: Expected utility and prospect theory model: smoothed time series, 1980–2003

	EU model	PT model
Weighting function parameter, δ	1	0.74
Utility function parameter, σ	1	0.50
Mean-square error, normalized	2.53	1
Akaike information criterion, AIC	-122.61	-142.90
Bayesian information criterion, BIC	-121.43	-140.55

Smoothed time series In this section, we focus on the long-term trend rather than capturing the year-to-year changes in the inequality measures. We apply a Hodrick–Prescott filter to the time series of the cross-sectional variances of the logarithm of income and the logarithm of consumption with smoothing parameter $\lambda_{\mathrm{HP}} = 400$.[20]

Compared to the baseline, the advantage of the prospect theory model with smoothed time series is even more significant. In Table 2, we report the statistics to assess the explanatory power of prospect theory and the expected utility model. The mean squared error in the prospect theory model is two and a half times smaller than in the expected utility model. Between 1980 and 2003, the prospect theory model predicts an increase in consumption inequality of 2.74 percentage points in the variance of log consumption which comes close to the actual increase in the data of 2.15 percentage points. The expected utility model tends to produce too much consumption insurance, predicting an increase in consumption inequality of only 0.13 percentage points. Moreover, the prospect theory model fits the data better than the standard incomplete markets model which predicts an increase of 4.5 percentage points (Krueger and Perri, 2006). As can be seen in Figure 6 in the Appendix, the prospect theory model performs substantially better than the expected utility model for values of the Tversky–Kahneman parameter between 0.725 and 0.765.

Persistent and transitory income components In this part, we employ an income specification that comprises transitory and persistent income components. The logarithm of income is modeled as

$$\ln(y_t) = z_t + \epsilon_t, \qquad z_t = \rho z_{t-1} + \eta_t,$$

---

[20] Alternatively, we also employ a value of $\lambda_{\mathrm{HP}} = 6.25$ for the smoothing parameter, which is the value proposed by Ravn and Uhlig (2002). The results on the performance of the two models are quantitatively similar.

Table 3: Persistent and transitory component of income, 1980–2003

	EU model	PT model
Weighting function parameter, δ	1	0.74
Utility function parameter, σ	1	0.50
Mean-square error, normalized	8.73	1
Akaike Information Criterion, AIC	-87.97	-137.97
Bayesian Information Criterion, BIC	-86.80	-135.611

where ϵ_t, η_t are independent, serially uncorrelated and normally distributed with variances $\sigma^2_{\epsilon_t}, \sigma^2_{\eta_t}$. The persistence parameter ρ is set to 0.95 which is the value found by Storesletten et al. (1999, 2004). Given the persistence parameter, we identify the variances $\sigma^2_{\epsilon_t}, \sigma^2_{\eta_t}$ from the cross-sectional within-group income variance and autocovariance in the CEX data. The persistent part of income is captured by a stochastic Markov process with seven time-varying exogenous states. The transitory part is modeled with two time-varying exogenous states. The results on the performance of prospect theory and the expected utility model are reported in Table 3 and confirm our findings in the baseline calibration. The mean square error for the prospect theory model is almost nine times smaller and both information criteria indicate a substantially better fit of the prospect theory model compared with the expected utility model.

Perfect foresight of changes in income inequality Here, we consider our baseline specification of the income process, however assuming that the actual changes in income inequality in 1980–2003 are known to the households in 1980. The optimal consumption allocation is computed backward. For example, the consumption in 2002 is computed taking as given the value of the insurance provided from 2003 onward. For 2003, the last year of the sample, we assume that the households expect income inequality to be constant from this year on. The results displayed in Table 4 confirm our baseline findings: the mean-square error for the prospect theory model is approximately two times smaller and both information criteria indicate a substantially better fit for prospect theory model.

Backward induction When stochastic outcomes are evaluated according to the backward-induction method, the increase in consumption inequality with respect to the increase in income

Table 4: Perfect foresight of changes in income inequality, 1980–2003

	EU model	PT model
Weighting function parameter, δ	1	0.74
Utility function parameter, σ	1	0.50
Mean-square error, normalized	2.10	1
Akaike Information Criterion, AIC	-122.83	-138.68
Bayesian Information Criterion, BIC	-121.65	-136.33

inequality is higher than under expected utility. The increase is however twice as high relative to our baseline prospect theory specification and thus too high compared to the data. The intuition is as follows. Ceteris paribus, backward-induction implies stronger risk aversion than our baseline preference specification as measured by a lower certainty equivalent for a given allocation. Under backward induction, the probability weighting affects decision weights accumulatively first for the conditional utility (4) and then again when computing the unconditional utility according to (5). The stronger risk aversion results in a lower value of the outside option and thus lower consumption inequality across all periods. As we calibrate the model to match consumption inequality in 1980, the value for T is adjusted downward which decreases the risk-sharing incentives. However, between 1980 and 2003 the value of the outside option for high-income households increases relatively more under backward induction than in our baseline prospect theory specification. The latter effect leads to a systematic overshooting of consumption inequality and a higher mean-square error than in our baseline specification. Summing up, the consumption inequality data for the U.S. tend to favor a dynamically inconsistent rather than dynamically consistent specification of prospect theory preferences.

6   Conclusion

The rise in within-group consumption inequality in response to the increase in within-group income inequality over the last three decades in the U.S. is puzzling to incomplete market models with expected utility. We have developed a model with two-sided lack of commitment and chance attitudes, as emphasized by prospect theory, which explains the relationship. Chance attitudes, such as optimism and pessimism, are captured by higher decision weights of high and low outcomes

relative to objective probabilities as in the expected utility model. With the proper accounting for chance attitudes as found in experimental studies, we have found that lenders' incentive to share are smaller than it would be implied by the expected utility models. The weaker willingness to share endogenously amplifies the increase in consumption inequality compared to the expected utility model and thereby improves the fit to the data.

An important question for future study is the decentralization of the optimal consumption allocation with prospect theory. On the one hand, Zimper (2012) provided conditions for existence and uniqueness of asset-price equilibria in a Lucas fruit-tree economy when chance attitudes are accounted for by a neo-additive weighting function. On the other hand, following Alvarez and Jermann (2000), the constraint-efficient allocation with history-dependent or non-history-dependent arrangements can be supported as an equilibrium allocation in the expected utility model with time- and state-contingent solvency constraints. The extension of these results to a prospect theory model with a general class of probability measures is a challenging task for future research.

Once the optimal consumption allocation with prospect theory is decentralized, one can also analyze the implications for the distribution of wealth. As documented by Cordoba (2008), the standard incomplete-market models with expected utility have also difficulties to reproduce the wealth inequality found in the data. A follow-up question is thus how the consumption and saving decisions of households with non-expected utility preferences affect the distribution of wealth.

References

ABEL, A. B. (2002): "An Exploration of the Effects of Pessimism and Doubt on Asset Returns," *Journal of Economic Dynamics and Control*, 26, 1075–1092.

AGUIAR, M. A. AND M. BILS (2011): "Has Consumption Inequality Mirrored Income Inequality?" NBER Working Papers 16807, National Bureau of Economic Research, Inc.

AIYAGARI, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109, 659–684.

ALLAIS, M. (1953): "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine," *Econometrica*, 21, 503–546.

ALVAREZ, F. AND U. J. JERMANN (2000): "Efficiency, Equilibrium, and Asset Pricing with Risk of Default," *Econometrica*, 68, 775–798.

ATTANASIO, O., E. BATTISTIN, AND H. ICHIMURA (2007): "What Really Happened to Consumption Inequality in the United States?" in *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, National Bureau of Economic Research, Inc, NBER Chapters, 515–543.

ATTANASIO, O., E. HURST, AND L. PISTAFERRI (2012): "The Evolution of Income, Consumption, and Leisure Inequality in the US, 1980-2010," NBER Working Papers 17982, National Bureau of Economic Research, Inc.

BARBERIS, N., M. HUANG, SANTOS, AND T. SANTOS (2001): "Prospect Theory and Asset Prices," *Quarterly Journal of Economics*, 116, 1–53.

BARSKY, R. B., T. F. JUSTER, M. S. KIMBALL, AND M. D. SHAPIRO (1997): "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *The Quarterly Journal of Economics*, 112, 537–79.

BATTALIO, R. C., J. H. KAGEL, AND K. JIRANYAKUL (1990): "Testing between Alternative Models of Choice under Uncertainty: Some Initial Results," *Journal of Risk and Uncertainty*, 3, 25–50.

BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98, 1887–1921.

BROER, T. (2011): "Crowding out and crowding in: When does redistribution improve risk-sharing in limited commitment economies?" *Journal of Economic Theory*, 146, 957–975.

BUSEMEYER, J. R., E. WEG, R. BARKAN, X. LI, AND Z. MA (2000): "Dynamic and consequential consistency of choices between paths of decision trees," *Journal of Experimental Psychology: General*, 129, 530–545.

CAMERER, C. F. (1989): "An Experimental Test of Several Generalized Utility Theories," *Journal of Risk and Uncertainty*, 2, 61–104.

CECCHETTI, S. G., P.-S. LAM, AND N. C. MARK (2000): "Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to Be True?" *American Economic Review*, 90, 787–805.

CHEW, S. H. AND W. S. WALLER (1986): "Empirical Tests of Weighted Utility Theory," *Journal of Mathematical Psychology*, 30, 55–72.

COATE, S. AND M. RAVALLION (1993): "Reciprocity without Commitment: Characterization and Performance of Informal Insurance Arrangements," *Journal of Development Economics*, 40, 1–24.

CONTE, A., J. D. HEY, AND P. G. MOFFATT (2011): "Mixture Models of Choice under Risk," *Journal of Econometrics*, 162, 79–88.

CORDOBA, J.-C. (2008): "U.S. inequality: Debt constraints or incomplete asset markets?" *Journal of Monetary Economics*, 55, 350–364.

CUBITT, R. P., C. STARMER, AND R. SUGDEN (1998): "Dynamic Choice and the Common Ratio Effect: An Experimental Investigation," *Economic Journal*, 108, 1362–1380.

DOMINIAK, A., P. DUERSCH, AND J.-P. LEFORT (2012): "A dynamic Ellsberg urn experiment," *Games and Economic Behavior*, 75, 625–638.

EICHBERGER, J., S. GRANT, AND D. KELSEY (2007): "Updating Choquet Beliefs," *Journal of Mathematical Economics*, 43, 888–899.

ELLSBERG, D. (1961): "Risk, Ambiguity and Savage Axioms," *Quarterly Journal of Economics*, 75, 643–669.

EPSTEIN, L. G. AND S. E. ZIN (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–69.

HAMMOND, P. J. (1989): "Consistent Plans, Consequentialism, and Expected Utility," *Econometrica*, 57, 1445–1449.

HANSEN, L. P., T. J. SARGENT, G. TURMUHAMBETOVA, AND N. WILLIAMS (2006): "Robust Control and Model Misspecification," *Journal of Economic Theory*, 128, 45–90.

HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006," *Review of Economic Dynamics*, 13, 15–51.

KAHNEMAN, D. AND A. TVERSKY (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263–291.

——— (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297–323.

KATZ, L. F. AND D. H. AUTOR (1999): "Changes in the Wage Structure and Earnings Inequality," *Handbook of Labor Economics*, 3, 1463–1555.

KIMBALL, M. S. (1988): "Farmers' Cooperatives as Behavior Toward Risk," *American Economic Review*, 78, 224–232.

KOCHERLAKOTA, N. (1996): "Implications of Efficient Risk Sharing without Commitment," *Review of Economic Studies*, 63, 595–609.

KRUEGER, D. AND F. PERRI (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *Review of Economic Studies*, 73, 163–193.

——— (2011): "Public versus Private Risk Sharing," *Journal of Economic Theory*, 146, 920–956.

LEPETYUK, V. AND C. A. STOLTENBERG (2013a): "On Efficient Distribution with Prospect Theory," University of Amsterdam, mimeo.

——— (2013b): "Policy Announcements and Welfare," *Economic Journal*, forthcoming.

LIGON, E., J. P. THOMAS, AND T. WORRALL (2002): "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies," *Review of Economic Studies*, 69, 209–244.

MACHINA, M. J. (1989): "Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty," *Journal of Economic Literature*, 27, 1622–1668.

MEYER, D. AND J. MEYER (2005): "Relative Risk Aversion: What Do We Know?" *Journal of Risk and Uncertainty*, 31, 243–262.

POST, T., M. J. VAN DEN ASSEM, G. BALTUSSEN, AND R. H. THALER (2008): "Deal or No Deal? Decision Making under Risk in a Large-Payoff Game Show," *American Economic Review*, 98, 38–71.

PRELEC, D. (1990): "A Pseudo-Endowment Effect, and its Implications for Some Recent Non-Expected Utility Models," *Journal of Risk and Uncertainty*, 3, 247–259.

QUIGGIN, J. (1982): "A Theory of Anticipated Utility," *Journal of Economic Behavior & Organization*, 3, 323–343.

RAVN, M. O. AND H. UHLIG (2002): "On Adjusting the Hodrick–Prescott Filter for the Frequency of Observations," *Review of Economics and Statistics*, 84, 371–375.

SAVAGE, L. J. (1954): *The Foundations of Statistics*, New York: Wiley.

STORESLETTEN, K., C. I. TELMER, AND A. YARON (1999): "The risk-sharing implications of alternative social security arrangements," *Carnegie-Rochester Conference Series on Public Policy*, 50, 213–259.

——— (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112, 695–717.

TAUCHEN, G. (1986): "Finite Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20, 177–181.

THOMAS, J. P. AND T. WORRALL (1988): "Self-Enforcing Wage Contracts," *Review of Economic Studies*, 55, 541–554.

VON NEUMANN, J. AND O. MORGENSTERN (1947): *Theory of Games and Economic Behavior*, First Edition, Princeton: Princeton University Press.

WAKKER, P. P. (2010): *Prospect Theory: For Risk and Ambiguity*, Cambridge University Press.

WEIL, P. (1990): "Nonexpected Utility in Macroeconomics," *The Quarterly Journal of Economics*, 105, 29–42.

WU, G. AND R. GONZALEZ (1996): "Curvature of the Probability Weighting Function," *Management Science*, 42, 1676–1690.

ZIMPER, A. (2012): "Asset Pricing in a Lucas Fruit-Tree Economy with the Best and Worst in Mind," *Journal of Economic Dynamics and Control*, 36, 610–628.

A Appendix

A.1 Example on dynamic consistency and consequentialism

Dynamic consistency and consequentialism can be illustrated with the following example. Consider the two-period decision tree in Figure 2. For the Kahneman–Tversky weighting function in Figure 7, the preference relationship at the initial node D_0 is $AC \succ AB \succ D$. At the subsequent decision node, D_1, however, the preference relationship will be reversed to $B \succ C$. Thus, in this example, the law of iterated expectations is violated and dynamic consistency is not preserved. In contrast, the decision weights under the expected utility are the objective probabilities of the outcomes and the weights are additive which results in dynamically consistent decisions.

Suppose that the preference relationship is such that the dynamic consistency is preserved, that is $AC \succ AB$ implies $C \succ B$. As the payoff of the lower outcome at A is increased, the preference relationship could eventually change to $AB \succ AC$ and therefore $B \succ C$. This result, however, illustrates the lack of consequentialism, as the preference relationship at D_1 does not exclusively depend on the possible outcomes at D_1. As pointed out by Machina (1989), the violation of dynamic consistency for non-expected utility roots in the principle of consequentialism that is implicitly imposed in the evaluation at the second decision node.
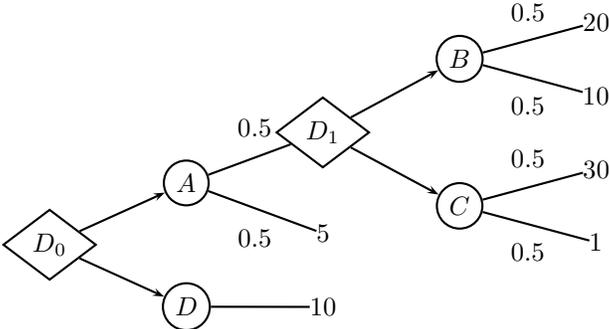

Figure 2: Decision tree

A.2 Proof of Proposition 3

Differentiating the participation constraint for high income households (7), we obtain

$$\frac{\partial \sigma_c}{\partial \gamma} = \frac{\beta}{1-\beta} \frac{1}{3C} \left(u(c_h) - u(c_l) - u(y_h) - u(\bar{y}) + 2u(y_l) \right) < 0,$$

where $c_h = \bar{y} + \sqrt{2}\sigma_c$ and $c_l = \bar{y} - \sigma_c/\sqrt{2}$, and σ_c is the standard deviation of the consumption allocation, and

$$C \equiv \left(1 + \frac{\beta}{1-\beta} \frac{1+2\lambda-\gamma}{3} \right) \sqrt{2} u'(c_h) - \left(\frac{\beta}{1-\beta} \frac{2-2\lambda+\gamma}{3} \right) \frac{1}{\sqrt{2}} u'(c_l).$$

Efficient risk sharing implies $c_h \leq y_h$ and by resource feasibility $y_l < c_l < \bar{y}$ that render the nominator of the derivative negative. As in Lemma 1, the denominator is positive at the optimal allocation. Similarly, taking the derivative with respect to λ we obtain

$$\frac{\partial \sigma_c}{\partial \lambda} = \frac{\beta}{1-\beta} \frac{1}{3C} \left(2u(y_h) - 2u(c_h) + 2u(c_l) - u(\bar{y}) - u(y_l) \right) > 0.$$

A.3 Proof of Lemma 1

Assuming $c_h \geq c_l$, the participation constraint for high-income households (7) is

$$\left(1 + \frac{\beta}{1-\beta} \frac{1+2\lambda-\gamma}{3} \right) u(c_h) + \left(\frac{\beta}{1-\beta} \frac{2-2\lambda+\gamma}{3} \right) u(c_l) =$$
$$\left(1 + \frac{\beta}{1-\beta} \frac{1+2\lambda-\gamma}{3} \right) u(y_h) + \left(\frac{\beta}{1-\beta} \frac{1-\lambda-\gamma}{3} \right) u(\bar{y}) +$$
$$\left(\frac{\beta}{1-\beta} \frac{1-\lambda+2\gamma}{3} \right) u(y_l) \quad \text{(A.1)}$$

Substituting for c_l from the resource-feasibility constraint (8), the left-hand side of (A.1) is a concave function of c_h. Hence, there are at most two solutions to (A.1) that are candidates for the optimal allocation. Expressing the allocations through deviations from the mean and employing

the implicit function theorem, from (A.1) we obtain

$$
\frac{\partial \sigma_c}{\partial \sigma_y} = \frac{\left(1 + \frac{\beta}{1-\beta}\frac{1+2\lambda-\gamma}{3}\right)\frac{\sqrt{3}}{\sqrt{2}}u'(y_h) - \left(\frac{\beta}{1-\beta}\frac{1-\lambda+2\gamma}{3}\right)\frac{\sqrt{3}}{\sqrt{2}}u'(y_l)}{\left(1 + \frac{\beta}{1-\beta}\frac{1+2\lambda-\gamma}{3}\right)\sqrt{2}u'(c_h) - \left(\frac{\beta}{1-\beta}\frac{2-2\lambda+\gamma}{3}\right)\frac{1}{\sqrt{2}}u'(c_l)}. \tag{A.2}
$$

If there are two solutions, the derivative of the left-hand side of (A.1) is positive at one solution and negative at the other. At the first best, $(1 + 2\lambda - \gamma)u'(c_h) = (2 - 2\lambda + \gamma)/2u'(c_l)$, such that the derivative of the left hand side of (A.1) with respect to c_h is positive. Because the objective function of the social planner is strictly concave, the optimal allocation corresponds to a solution that is closer to the first best, which is not sustainable, and the derivative of the left-hand side at the constrained-efficient allocation is positive.

A.4   Proof of Proposition 4

Using the participation constraint (7) of high-income households to differentiate σ_c with respect to σ_y and γ gives

$$
\frac{\partial \left(\partial \sigma_c / \partial \sigma_y\right)}{\partial \gamma} =
$$
$$
\frac{1}{C^2}\left\{ \frac{\beta}{1-\beta}\frac{-1}{3}\left(\frac{\sqrt{3}}{\sqrt{2}}u'(y_h) + \frac{2\sqrt{3}}{\sqrt{2}}u'(y_l)\right)C + \frac{\beta}{1-\beta}\frac{1}{3}\left(\sqrt{2}u'(c_h) + \frac{1}{\sqrt{2}}u'(c_l)\right)A \right.
$$
$$
\left. - \left(\left(1 + \frac{\beta}{1-\beta}\frac{1+2\lambda-\gamma}{3}\sqrt{2}\right)u''(c_h) + \left(\frac{\beta}{1-\beta}\frac{2-2\lambda+\gamma}{3}\right)\frac{1}{\sqrt{2}}u''(c_l)\right)\frac{\partial c_h}{\partial \gamma}A \right\},
$$

where
$$
A = \left(1 + \frac{\beta}{1-\beta}\frac{1+2\lambda-\gamma}{3}\right)\frac{\sqrt{3}}{\sqrt{2}}u'(y_h) - \left(\frac{\beta}{1-\beta}\frac{1-\lambda+2\gamma}{3}\right)\frac{\sqrt{3}}{\sqrt{2}}u'(y_l), \text{ and}
$$
$$
C = \left(1 + \frac{\beta}{1-\beta}\frac{1+2\lambda-\gamma}{3}\right)\sqrt{2}u'(c_h) - \left(\frac{\beta}{1-\beta}\frac{2-2\lambda+\gamma}{3}\right)\frac{1}{\sqrt{2}}u'(c_l).
$$

From Lemma 1, we obtain $A > 0$ and $C > 0$. Efficient risk sharing implies that $c_h < y_h$; combined together with $\gamma \geq \bar{\gamma}$, where $\bar{\gamma} \equiv (1 - \lambda)(2 - \sqrt{3})/(2\sqrt{3} - 1)$, we get $C > A$. Constant relative risk aversion results in convex marginal utilities and ensures that

$$
\frac{\sqrt{3}}{\sqrt{2}}u'(y_h) + \frac{2\sqrt{3}}{\sqrt{2}}u'(y_l) > \frac{\sqrt{3}}{\sqrt{2}}u'(c_h) + \frac{2\sqrt{3}}{\sqrt{2}}u'(c_l) > \sqrt{2}u'(c_h) + \frac{1}{\sqrt{2}}u'(c_l)
$$

and thus $\partial^2\sigma_c/(\partial\sigma_y\partial\gamma) < 0$.

A.5   Additional tables and figures

Table 5: Estimates of the prospect-theory parameters

Study	Weighting function, δ	Utility function, σ
Battalio et al. (1990)	0.72	–
Camerer (1989)	0.82	0.78
Chew and Waller (1986)	0.74	–
Conte et al. (2011)	$0.74 - 1.06$	$0.29 - 0.82$
Prelec (1990)	0.69	–
Kahneman and Tversky (1992)	0.61	0.12
Wu and Gonzalez (1996)	$0.61 - 0.81$	$0.38 - 0.62$
Average	0.74	0.50
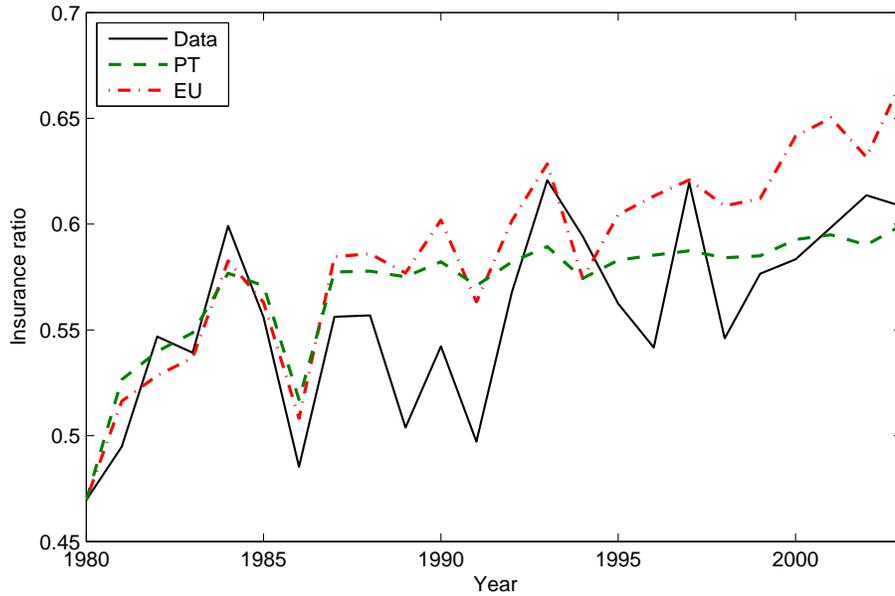
Figure 3: Insurance ratio $1 - \text{var}_i(\ln c_t^i)/\text{var}_i(\ln y_t^i)$ over time in the data (solid) versus the prospect theory model (PT; dashed) and the expected utility model (EU; dash-dotted)
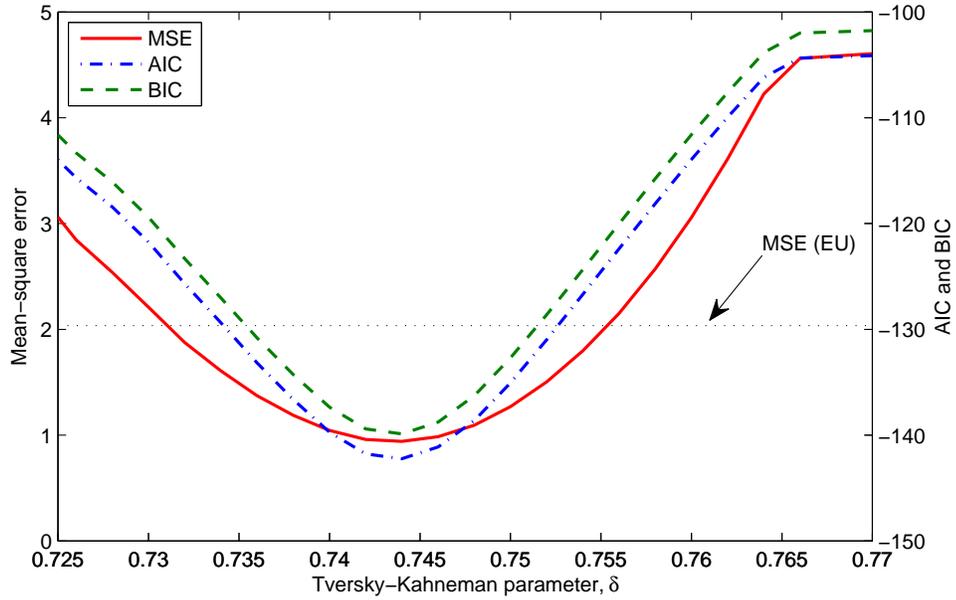
Figure 4: Mean squared error (MSE), Akaike information criterion (AIC), and Bayesian information criterion (BIC) in the prospect theory model for different values of the Tversky–Kahneman parameter

Table 6: Expected utility and prospect theory model: Gini coefficient, 1980–2003

	EU model	PT model
Weighting function parameter, δ	1	0.74
Utility function parameter, σ	1	0.50
Mean-square error, normalized	2.96	1
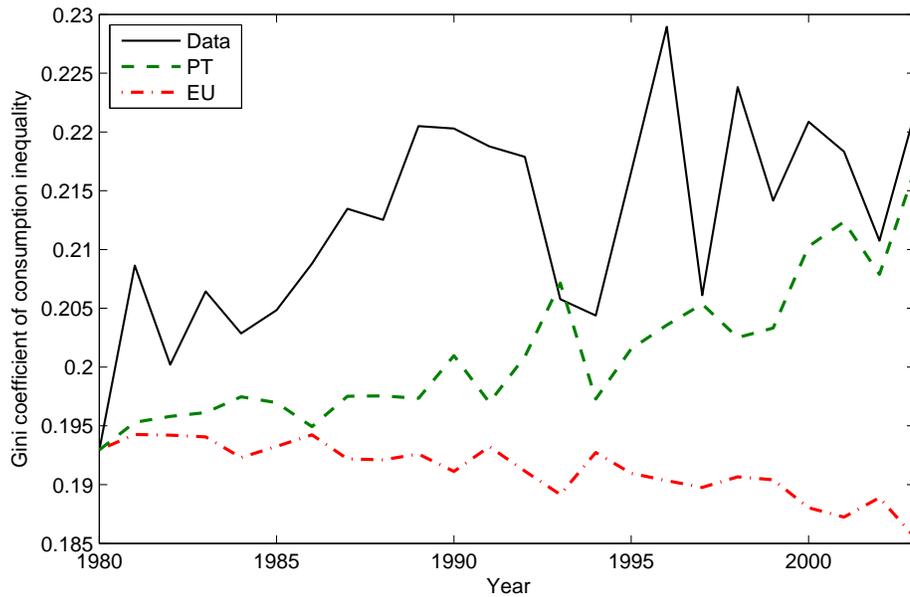
Figure 5: Gini coefficient of consumption inequality over time in the data (solid) versus the prospect theory model (PT; dashed) and the expected utility model (EU; dash-dotted)
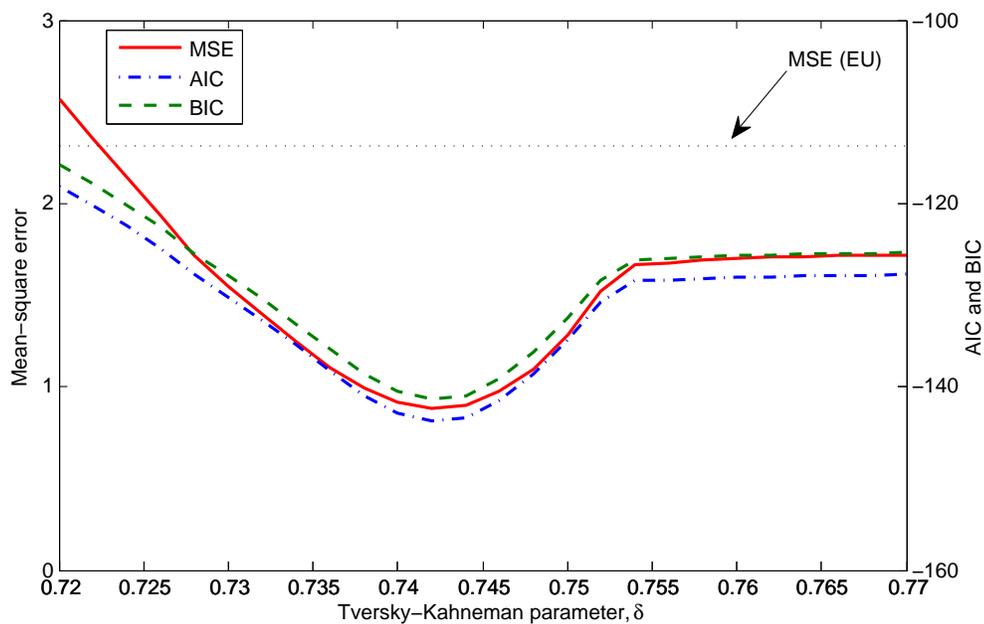

Figure 6: Smoothed time series. Mean squared error (MSE), Akaike information criterion (AIC), and Bayesian information criterion (BIC) in the prospect theory model for different values of the Tversky–Kahneman parameter
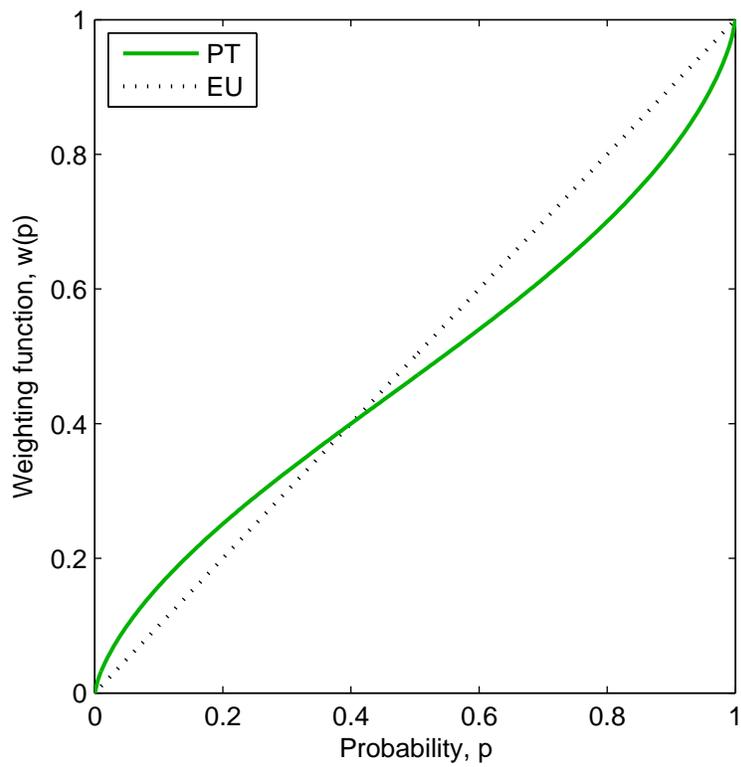
Figure 7: Tversky–Kahneman weighting function $w(p) = p^\delta/(p^\delta + (1-p)^\delta)^{1/\delta}$, for $\delta = 0.74$ (PT; solid) versus the expected utility (EU; dotted)