Inference under Stability of Risk Preferences*

Levon Barseghyan
Cornell University

Francesca Molinari
Cornell University

Joshua C. Teitelbaum
Georgetown University

Extended Abstract
April 3, 2013

Economists strive to develop models that can explain behavior across manifold domains. At a minimum, we ask that a model's explanatory power extend across contexts that are essentially similar. Stated more formally, we require that a model satisfy a criterion of *stability*: a single parameterization of the model should be consistent with observed behavior in closely related contexts.

One can treat stability as a testable hypothesis. In the case of models of consumer choice, we can combine the stability criterion with revealed preference arguments to test the hypothesis that consumers' preferences—as represented by the model and revealed by their choices—are stable across similar decision contexts. Barseghyan et al. (2011) and Einav et al. (2012) take this approach to examine the stability of risk preferences, investigating whether there exists a single parameterization of the expected utility model that is consistent with data on insurance choices across multiple lines of coverage. Both found that most consumers (more than two thirds) do not exhibit stable risk preferences under the expected utility model. Yet Einav et al. (2012) also found that consumers' choices are rank correlated across contexts, suggesting that their risk preferences have a domain-general component but are not well represented by the expected utility model.

Motivated in part by this work, Barseghyan et al. (2012) [hereafter, BMOT] used data on households' insurance choices to estimate a generalization of the expected utility model that allows for "generic" probability distortions. The probability distortions in the BMOT model are generic in the sense that they can arise from a number of sources, including systematic risk misperceptions, rank-dependent probability weighting (Quiggin 1982), Kőszegi-Rabin loss aversion (Kőszegi and Rabin 2006, 2007), and Gul disappointment aversion (Gul 1991); see BMOT for a discussion. Based on their estimates, BMOT concluded that probability distortions—in the form of substantial overweighting of claim probabilities—play an important role in explaining the data.

We take a different approach in this paper. Rather than treat stability as a testable hypothesis, we exploit the stability criterion to conduct inference on the structure of households' risk preferences, as represented by the probability distortion model and revealed by deductible choices in three lines of property insurance. We take a partial identification approach (Manski 2003), making minimal additional assumptions and adding these assumptions sequentially in order to transparently show the role that each plays in sharpening the inference. Under this approach, we need make no assumptions about the relationship between observed heterogeneity and risk preferences, nor about the distribution of unobserved heterogeneity in risk preferences. The idea is simply to use revealed preferences arguments to bound the model parameters and to exploit the stability criterion and other minimal assumptions—all but one of which amount to shape restrictions on the utility and probability distortions functions—to sharpen the inference. The more choices we observe per household, the more precise the inference we can make about the household's risk preferences.

*Supported by NSF Grant SES-1031136.

In addition to stability, we make four additional assumptions. The first is constant absolute risk aversion, our only restriction on the shape of the utility function. The second assumption is plausibility: we require that there exists a single coefficient of absolute risk aversion and three distorted probabilities (one for each context) that can rationalize a household's choices. Given CARA, plausibility cannot be rejected for 87 percent of households. The last two assumptions—monotonicity and linearity—are shape restrictions on the probability distortion function. Monotonicity requires that the distortion function is increasing, and linearity requires that the function is linear. Monotonicity cannot be rejected for 85 percent of "rationalizable" households (i.e., households that satisfy plausibility), and monotonicity and linearity cannot be rejected for 81 percent of rationalizable households. By contrast, only 40 percent of households are consistent with expected utility, which entails two additional restrictions: unit slope and zero intercept. Naturally, as we add shape restrictions, the model can rationalize the choices of fewer households. However, the inference about the probability distortion function becomes much sharper: with monotonicity and linearity the implied bounds on the distortion function become 42 percent tighter.

In the next part of our analysis, we address two questions: (1) What single probability distortions function comes closest to rationalizing the choices of all households? and (2) What is the fraction of households whose behavior can be rationalized by this single probability distorting function? To answer these questions we propose and estimate a best linear point predictor that minimizes the expected Euclidean distance to each household's set of "stable" probability distortion functions. We prove that under mild conditions (satisfied in our data) the parameters of the predictor are point identified, and we establish consistency and asymptotic normality of our sample analog estimator. We find that (1) our "minimum distance" probability distortion function is remarkably similar to the "maximum likelihood" probability distortion function estimated in BMOT and (2) all three choices of 18 percent of "linear" households can be fully rationalized by this single probability distortion function.

We conclude our analysis by addressing three issues. First, we demonstrate a close connection between stability of preferences and rank correlation of choices. In short, we find that households who satisfy linearity (i.e., who satisfy stability and each of our additional assumptions through and including linearity) drive up the rank correlations across choices. Second, we use stability as a model selection criterion, comparing our model with the expected utility model and leading alternatives such as the probability weighting models of Tversky and Kahneman (1992) and Prelec (1998) and the loss aversion model of Kőszegi and Rabin (2006, 2007). According to the stability criterion, our model outperforms the others, including when we further restrict our model to have a linear utility function and a linear probability distortion function with unit slope. Finally, we exclude the possibility that unobserved heterogeneity in claim probabilities is driving our results.

References

Barseghyan, L., F. Molinari, T. O'Donoghue, and J. C. Teitelbaum (2012): "The Nature of Risk Preferences: Evidence from Insurance Choices," SSRN Working Paper 1646520. Forthcoming, *American Economic Review.*

Barseghyan, L., J. Prince, and J. C. Teitelbaum (2011): "Are Risk Preferences Stable Across Contexts? Evidence from Insurance Data," *American Economic Review*, 101, 591–631.

Einav, L., A. Finkelstein, I. Pascu, and M. R. Cullen (2012): "How General are Risk Preferences? Choice under Uncertainty in Different Domains," *American Economic Review*, 102, 2606–2638.

Gul, F. (1991): "A Theory of Disappointment Aversion," *Econometrica*, 59, 667–686.

Kőszegi, B. and M. Rabin (2006): "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, 121, 1133–1166.

——— (2007): "Reference-Dependent Risk Attitudes," *American Economic Review*, 97, 1047–1073.

Manski, C. F. (2003): *Partial Identification of Probability Distributions*, New York: Springer Verlag.

Prelec, D. (1998): "The Probability Weighting Function," *Econometrica*, 66, 497–527.

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

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

Inference under Stability of Risk Preferences[1]

L. Barseghyan F. Molinari J. Teitelbaum

Cornell and Georgetown

April 2013

[1]Supported by NSF Grant SES - 1031136.

Introduction

- We start with the assumption that households use the same model of decision making under risk when faced with similar lotteries.
- What type of inference about each households' underlying risk preferences can one make relying solely on this stability assumption (and the model)?
- We answer this question within a model that allows for "generic" probability distortions as in Barseghyan, Molinari, O'Donoghue, and Teitelbaum (forthcoming).
  - This model nests the standard expected utility model as a special case.
- We propose an estimator that yields the probability distortion function which is closest in the Euclidean sense to generating distorted probabilities consistent with the households' choices across our sample.

Motivation

- Recent studies find that under the standard expected utility model households do not exhibit stable risk preferences.
- Barseghyan, Prince, and Teitelbaum (2011)
  - The stability of risk preferences is rejected for a sample of households choosing auto and home deductibles.
- Einav, Finkelstein, Pascu, and Cullen (2012)
  - The stability of risk preferences is rejected for a sample of households choosing health, drug, dental, and disability insurances and 401(k) investments (about 30% appear stable);
  - Yet (rank) correlations across choices are positive.
- Hence a conjecture: Preferences are stable, but the model is "wrong."

Preview of Findings

- About 13% of households make choices that cannot be "rationalized" under any model considered.
- 85% of the rationalizable households make decisions consistent with an increasing distortion function.
- 82% – consistent with an increasing quadratic distortion function.
- 81% – consistent with an increasing linear distortion function.
- 62% – consistent with an increasing linear unit slope distortion function.
- A single linear distortion function (our estimator obtained by minimizing the Euclidean distance to the distorted probabilities implied by the households choices in our sample) can explain all choices of 18% of the households whose choices can be rationalized by a linear distortion function.

Overview of the Data

- Full data set comprises yearly information on more than 400,000 households who held auto or home policies between 1999 and 2006.
- Data include for each household:
  - full policy information
  - full menu of premium-deductible combinations for each coverage
  - full claims history under each coverage
  - rich set of demographic information
- We focus on three choices:
  - auto collision deductible
  - auto comprehensive deductible
  - home all perils deductible
- Sample: We restrict attention to the initial deductible choices of households who hold both auto and home policies and who first purchased their auto and home policies from the company in the same year, in either 2005 or 2006. This yields 4170 observations.

Table 1: Descriptive Statistics

Sample (4170 Ꮶouseholds)

Variable	**Mean**	**Std Dev**	**1st Pctl**	**99th Pctl**
Driver 1 age (years)	54.5	15.4	26	84
Driver 1 female	0.37			
Driver 1 single	0.24			
Driver 1 married	0.51			
Driver 1 ǈyãdzǋl·yǑŜ score	т66	м13	p30	ф87
Driver 2 indicator	0.42			
Home value (thousands of dollars)	191	125	10	619

Note: Omitted category for driver 1 marital status is divorced or separated.

Table 2: Summary of Deductible Choices

Sample (4170 households)

Deductible	Collision	Comp	Home
$50		5.2	
$100	1.0	4.1	0.9
$200	13.4	33.5	
$250	11.2	10.6	29.7
$500	67.7	43.0	51.9
$1000	6.7	3.6	15.9
$2500			1.2
$5000			0.4

Note: Values are percent of households.

Table 3: Summary of Premium Menus

Sample (4170 households)

Coverage	Mean	Std Dev	1st Pctl	99th Pctl
Auto collision premium for $500 deductible	180	100	50	555
Auto comprehensive premium for $500 deductible	115	81	26	403
Home all perils premium for $500 deductible	679	519	216	2511
Cost of decreasing deductible from $500 to $250:				
Auto collision	54	31	14	169
Auto comprehensive	30	22	6	107
Home all perils	56	43	11	220
Savings from increasing deductible from $500 to $1000:				
Auto collision	41	23	11	127
Auto comprehensive	23	16	5	80
Home all perils	74	58	15	294

Note: Annual amounts in dollars.

Claim Probabilities

- We estimate each household's claim probability μ for each coverage.
  - $\mu =$ probability that the household will experience at least one claim for such coverage during the policy period
- We begin by estimating how claim rates depend on observables.
  - We use the full data set: 1,347,461 household-year records for auto and 1,265,370 household-year records for home.
- We assume household i's claims under coverage j in year t follow a Poisson distribution with arrival rate λ_{ijt}.
- For each coverage, we estimate a Poisson panel regression model with random effects.
- We use the regression results to generate predicted annual claim rates $\widehat{\lambda}_{ij}$ for each household i and coverage j, conditional on the household's (ex ante) characteristics and (ex post) claims experience.
- We use $\widehat{\lambda}_{ij}$ to calculate the predicted claim probability $\widehat{\mu}_{ij}$ for each household in our sample: $\widehat{\mu}_{ij} = 1 - \exp(-\widehat{\lambda}_{ij})$.

Table 4: Predicted Claim Probabilities

Sample (4170 households)

	Collision	Comp	Home
Mean	0.069	0.021	0.084
Standard deviation	0.024	0.011	0.044
1st percentile	0.026	0.004	0.024
5th percentile	0.035	0.007	0.034
25th percentile	0.052	0.013	0.053
Median	0.066	0.019	0.076
75th percentile	0.083	0.027	0.104
95th percentile	0.114	0.041	0.163
99th percentile	0.139	0.054	0.233

- We assume the choice of deductible involves a choice among lotteries of the following form:

$$L_d \equiv (-p_d, 1-\mu; -p_d - d, \mu) .$$

- Key underlying assumptions:
  - household treats its deductible choices as independent decisions
  - full insurance against covered losses in excess of the deductible
  - household experiences at most one claim during the policy period
  - choice of deductible does not influence μ (no moral hazard)
  - every claim exceeds the highest available deductible

The Model
Probability Distortions

- We incorporate probability distortions into the standard expected utility model:

$$U(L_d) = (1 - \Omega(\mu)) \cdot u(w - p_d) + \Omega(\mu) \cdot u(w - p_d - d),$$

where w is the household's wealth.

- Note if $\Omega(\mu) = \mu$ the model reduces to the standard model:

$$EU(L_d) = (1 - \mu) \cdot u(w - p_d) + \mu \cdot u(w - p_d - d).$$

- As shown in Barseghyan, Molinari, O'Donoghue and Teitelbaum (forthcoming), probability distortions arise if we incorporate rank-dependent probability weighting, Kőszegi-Rabin loss aversion, or Gul's disappointment aversion into the standard model:

Rank-dependent probability weighting:	$\Omega(\mu) = \pi(\mu)$
Kőszegi-Rabin loss aversion:	$\Omega(\mu) = \mu\left[1 + \Lambda\left(1 - \mu\right)\right]$
Gul's disappointment aversion:	$\Omega(\mu) = \frac{(1+\beta)\mu}{1+\beta\mu}$

- Of course, probability distortions also can be interpreted as systematic risk misperceptions.

Bounds on Distorted Probabilities

- Per our model:

$$U(L_d) = (1 - \Omega(\mu)) \cdot u(w - p_d) + \Omega(\mu) \cdot u(w - p_d - d)$$

- Hence, for a choice k revealed preference arguments yield:

$$LB \leq \Omega(\mu) \leq UB$$

$$LB = \max\left\{0, \max_{h>k}\left[\frac{\Delta_{h,k} u(w - p_d)}{\Delta_{h,k}\left\{u(w - p_d) - u(w - p_d - d)\right\}}\right]\right\}$$

$$UB = \min\left\{1, \min_{h<k}\left[\frac{\Delta_{h,k} u(w - p_d)}{\Delta_{h,k}\left\{u(w - p_d) - u(w - p_d - d)\right\}}\right]\right\}$$

where $\Delta_{h,k} x \equiv x_h - x_k$.

Bounds on Distorted Probabilities

An example with linear $u(\cdot)$:

$$U(L_d) = -p_d - \Omega(\mu)d$$

- Menu consists of three deductible options:

  $$\{\$250, \$500, \$1000\}$$

  and the choice is $500.
- Then, *LB* (*UB*) is the distorted probability that makes the household indifferent between $500 and $1000 (*between $500 and $250*):

  $$\frac{p_{500} - p_{1000}}{1000 - 500} \leq \Omega(\mu) \leq \frac{p_{250} - p_{500}}{500 - 250}$$
- Note: The bounds on $\Omega(\mu)$ are functions of prices and deductibles.
  - They do not depend on claim probabilities.

Choice of the Utility Function

- For our purposes, CRRA and CARA are effectively equivalent.
  - For realistic levels of wealth and ρ (coef. of relative risk aversion) there exists an r (coef. of absolute risk aversion) such that utility differences implied by the lotteries in consideration are virtually identical under CRRA and CARA.
- Another alternative is the class of utility functions with a Negligible Third Derivative (Cohen and Einav 2007). NTD generates results very similar to CRRA/CARA.
- In what follows, we assume CARA.

Bounds on Distorted Probabilities and Stability

- Under the standard expected utility multiple contexts imply multiple intervals for the same object - the risk aversion parameter (BPT 2011; EFPC 2012)
  - Stability implies that these intervals should intersect - refining inference for r.

- A fundamental difference here:
  For a given r, multiple contexts imply multiple intervals for different objects, $\Omega(\mu_{ij})$, $j \in \{Coll, Comp, Home\}$, because μ_{ij}'s (generically) are not the same across contexts.

- Hence, stability of risk preferences in our model has no identifying power without shape restrictions on the probability distortion function $\Omega(\cdot)$.
  - Without shape restrictions each interval simply gives the admissible values for the function $\Omega(\cdot)$ at a given μ.

First Step - Plausibility

- Plausibility $\equiv$ *non-empty interval:*

$$LB \leq UB$$

- Of course, LB and UB depend on the degree of risk aversion.

- Hence plausibility amounts to asking whether there exists
  - a single coef. of absolute risk aversion r and
  - three distorted probabilities (one for each context)

  that can rationalize a household's choices.

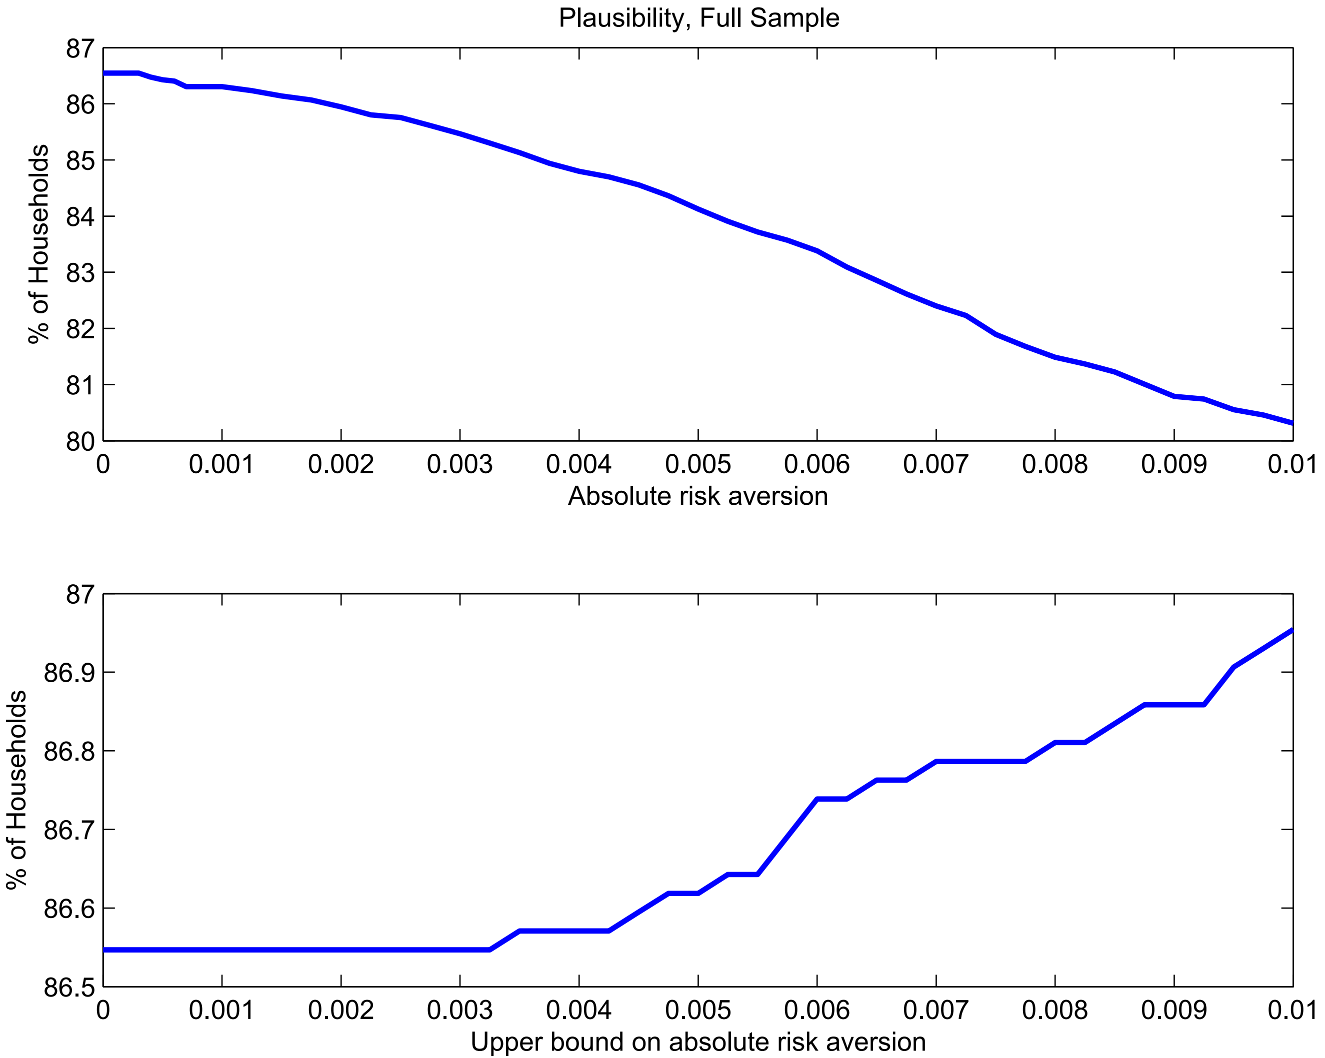

Plausibility, Full Sample
% of Households
87
86
85
84
83
82
81
80
0
0.001
0.002
0.003
0.004
0.005
0.006
0.007
0.008
0.009
0.01
Absolute risk aversion
% of Households
87
86.9
86.8
86.7
86.6
86.5
0
0.001
0.002
0.003
0.004
0.005
0.006
0.007
0.008
0.009
0.01
Upper bound on absolute risk aversion

Plausibility: Who Fails and Why

- Essentially all households that fail plausibility (about 13%) have chosen an auto collision deductible of \$200, which is a dominated choice for them:

Pricing Rule for Collision

deductible	\$100	\$200	\$250	\$500	\$1000
base price multiplier	1.15	1.00	0.85	0.65	0.50

With linear $u(\cdot)$, for any base price $\bar{p}$:

$$LB = \frac{p_{200} - p_{250}}{250 - 200} = \bar{p} \cdot \frac{0.15}{50}$$

and

$$UB = \frac{p_{100} - p_{200}}{200 - 100} = \bar{p} \cdot \frac{0.15}{100}$$

- With concave $u(\cdot)$ these patterns may change for only a tiny fraction of households...
- Going forward we *drop* households that have chosen \$200 in auto collision.

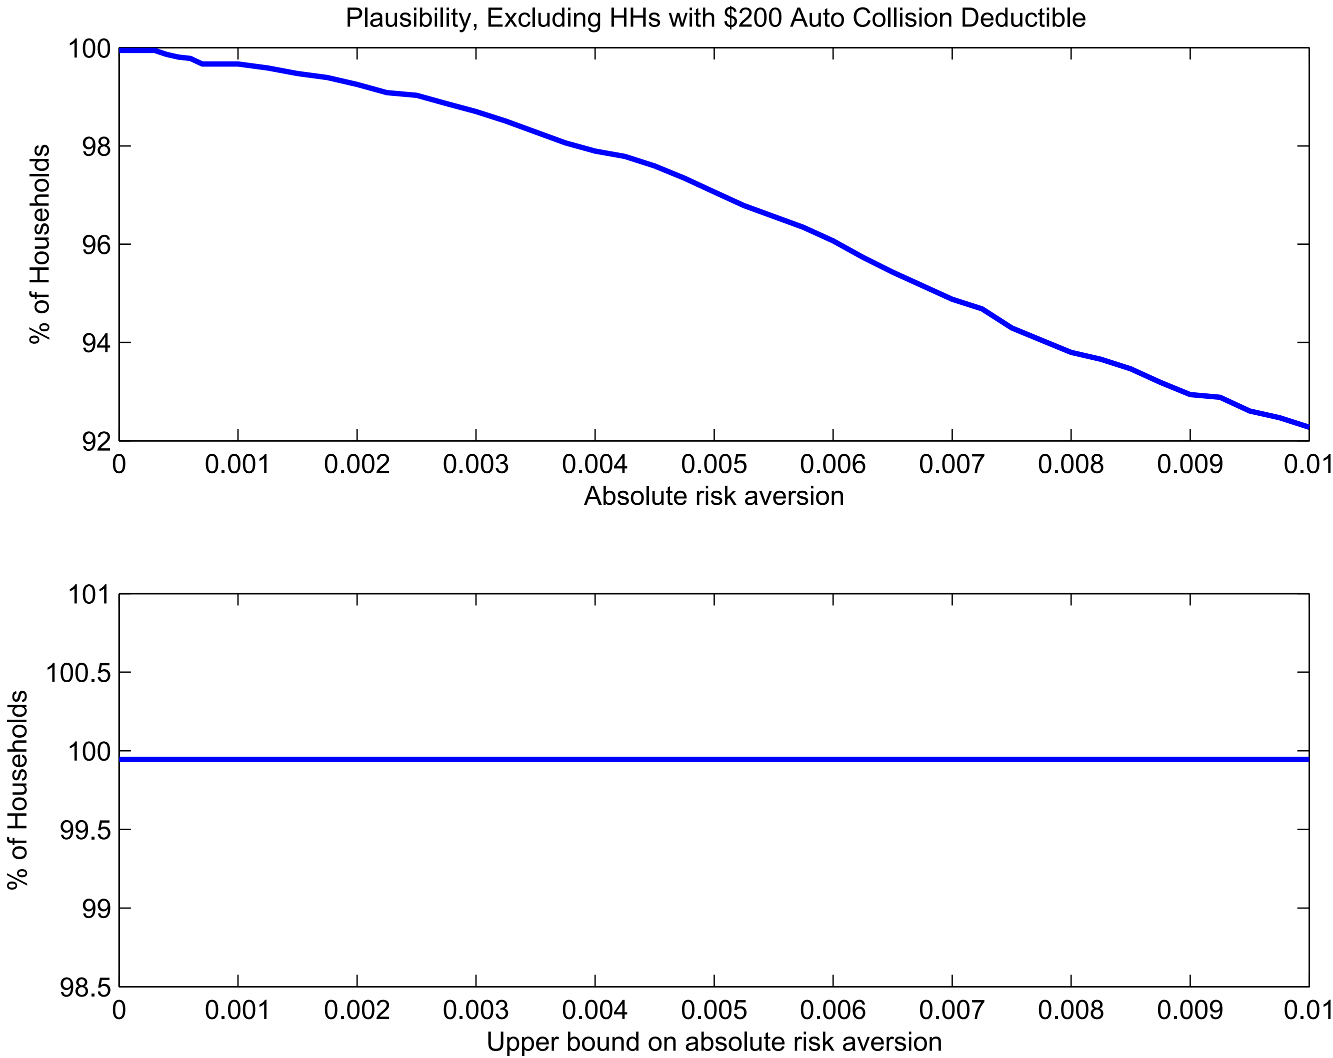

Plausibility, Excluding HHs with $200 Auto Collision Deductible
% of Households
Absolute risk aversion
100
98
96
94
92
0
0.001
0.002
0.003
0.004
0.005
0.006
0.007
0.008
0.009
0.01
% of Households
Upper bound on absolute risk aversion
101
100.5
100
99.5
99
98.5

Shape Restrictions on Probability Distortion Function

- Monotonicity:

$$\text{If } \mu \leq \tilde{\mu} \text{ then } \Omega(\mu) \leq \Omega(\tilde{\mu}).$$

- Monotone + Quadratic Shape:

$$\Omega(\mu) = a + b\mu + c\mu^2.$$

- Monotone + Linear Shape:

$$\Omega(\mu) = a + b\mu.$$

- Monotone + Linear Shape, Unit Slope:

$$\Omega(\mu) = a + \mu.$$

- Monotone + Linear Shape, Unit Slope, Zero Intercept = EUT:

$$\Omega(\mu) = \mu.$$

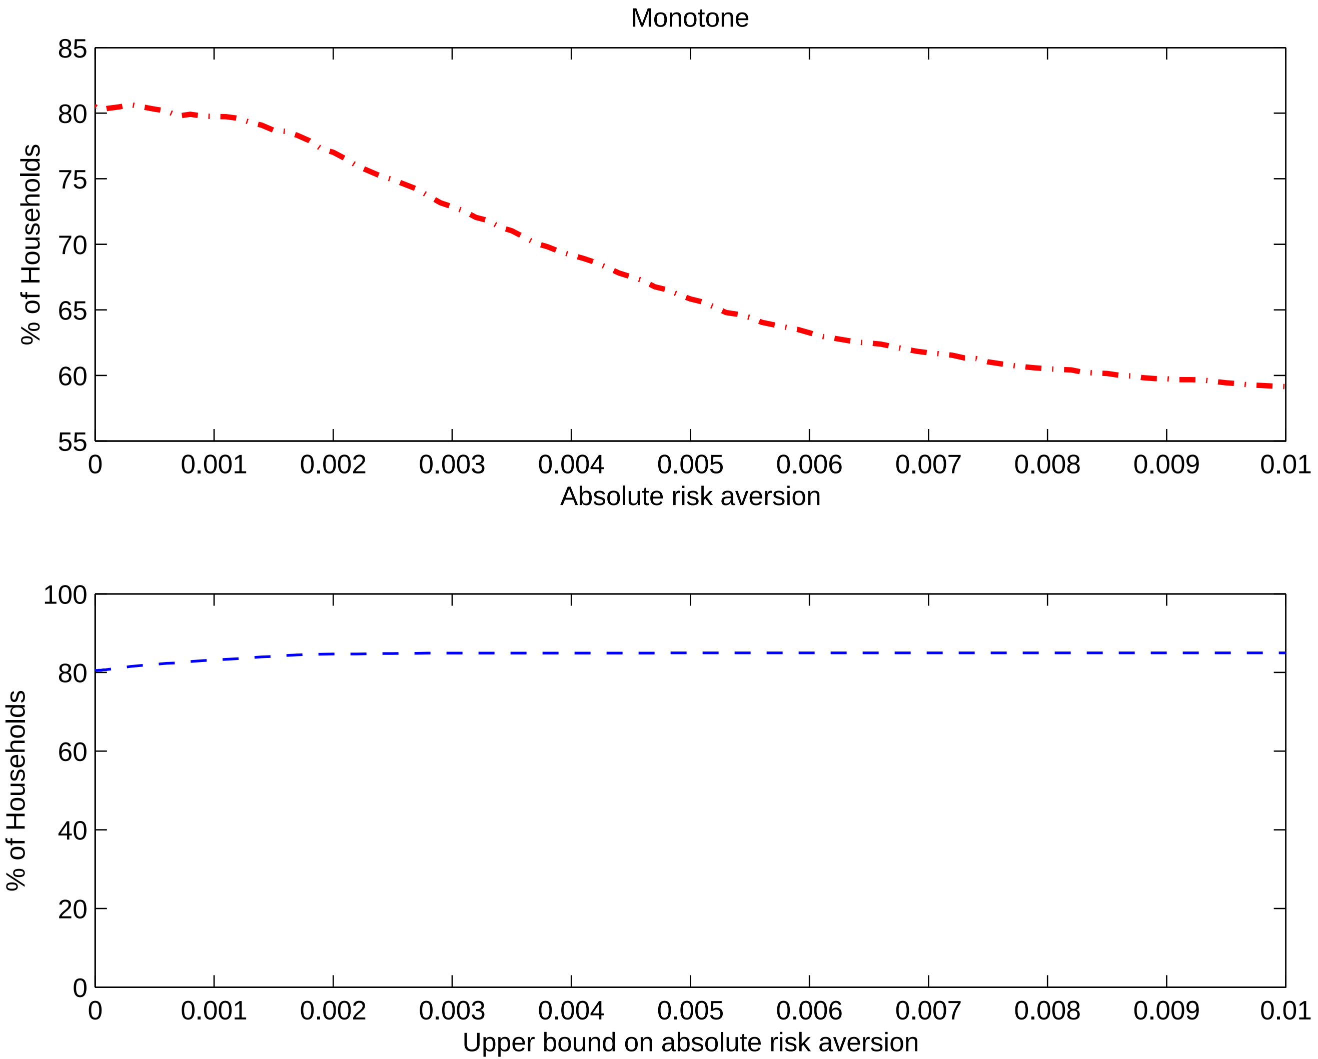
Monotone
85
80
75
70
65
60
55
% of Households
0
0.001
0.002
0.003
0.004
0.005
0.006
0.007
0.008
0.009
0.01
Absolute risk aversion
100
80
60
40
20
0
% of Households
0
0.001
0.002
0.003
0.004
0.005
0.006
0.007
0.008
0.009
0.01
Upper bound on absolute risk aversion

Monotonicity: What is the Extent of the Failure?

- One way to measure it is as follows.
- Suppose HHs use $\tilde{\Omega}(\mu)$ rather than $\Omega(\mu)$ when making decisions:

$$\tilde{\Omega}(\mu) = \Omega(\mu) + \varepsilon_{\Omega}.$$

- What is the distribution of the smallest (in absolute value) ε_{Ω} such that every household's $\Omega(\mu)$ satisfies monotonicity?

unconditional	conditional on $\lvert\varepsilon_{\Omega}\rvert > 0$		
average $\lvert\varepsilon_{\Omega}\rvert$	average $\lvert\varepsilon_{\Omega}\rvert$	median $\lvert\varepsilon_{\Omega}\rvert$	std. dev. of $\lvert\varepsilon_{\Omega}\rvert$
0.008	0.110	0.055	0.135

- For comparison, the average mid-points of the $\Omega(\mu)$ intervals at $\mu = 0.02$, 0.07, and 0.10 are 0.12, 0.19, and 0.23, respectively.

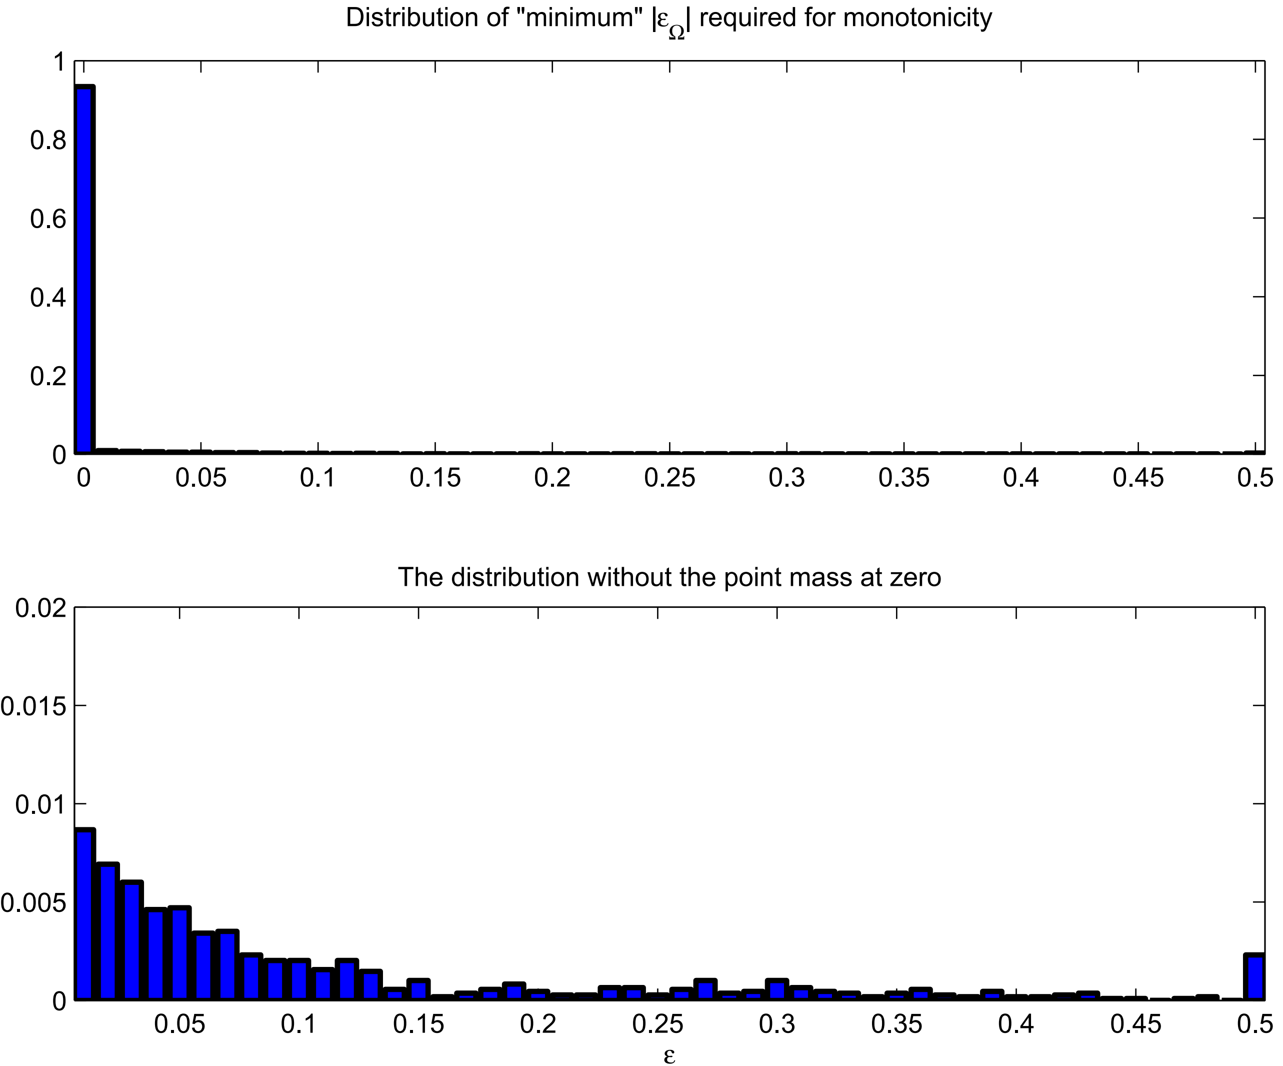

Distribution of "minimum" |ε_Ω| required for monotonicity
1
0.8
0.6
0.4
0.2
0
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
0.4
0.45
0.5
The distribution without the point mass at zero
0.02
0.015
0.01
0.005
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
0.4
0.45
0.5
ε

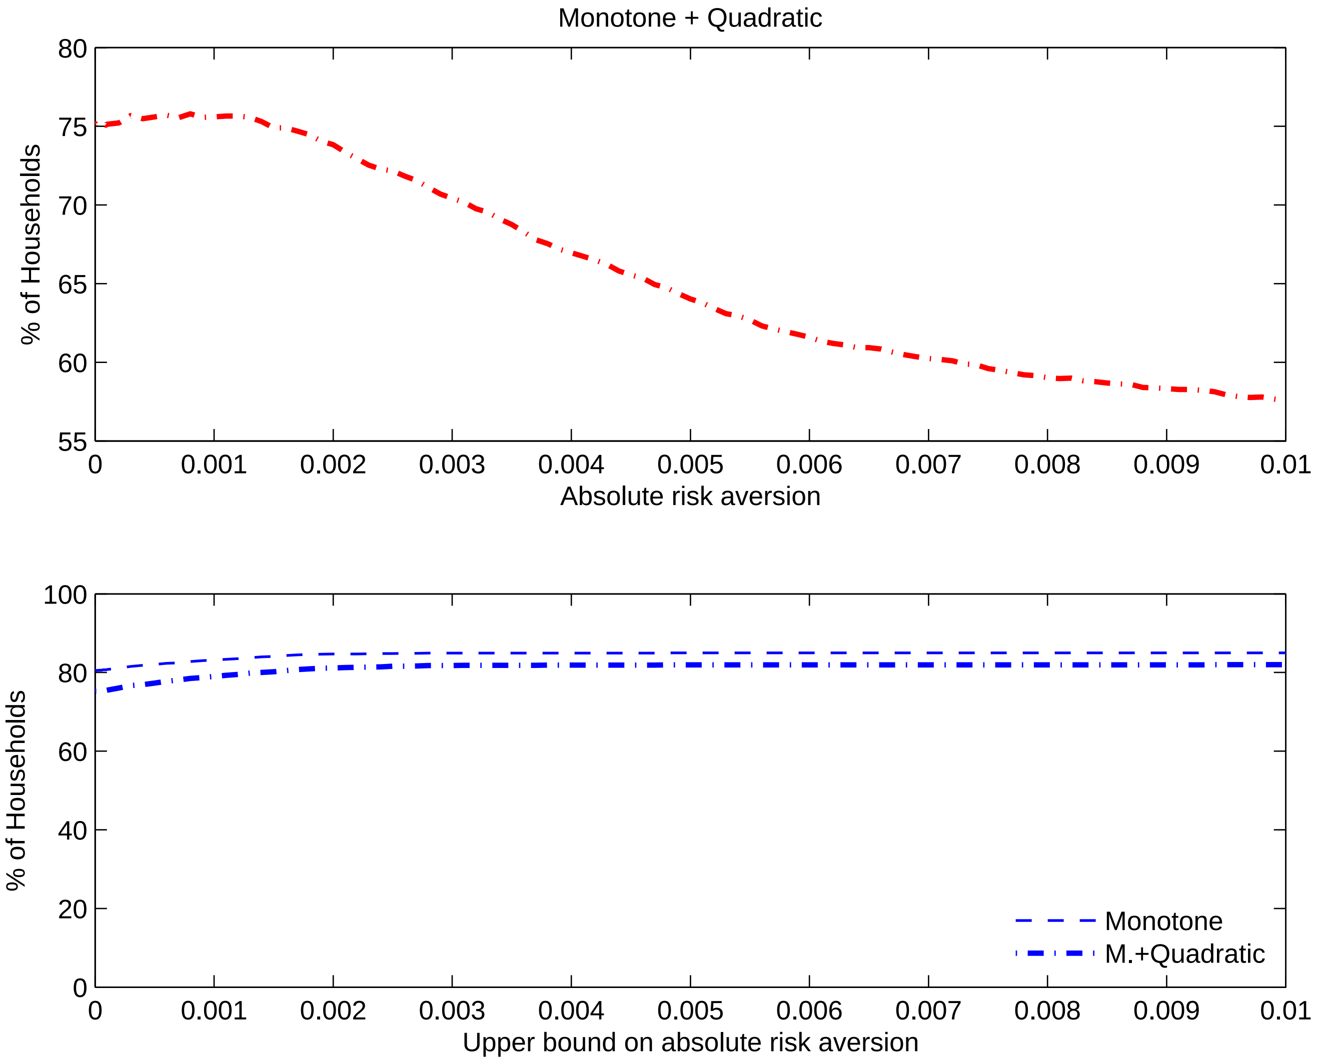
Monotone + Quadratic
80
75
70
65
60
55
% of Households
0
0.001
0.002
0.003
0.004
0.005
0.006
0.007
0.008
0.009
0.01
Absolute risk aversion
100
80
60
40
20
0
% of Households
Monotone
M.+Quadratic
Upper bound on absolute risk aversion

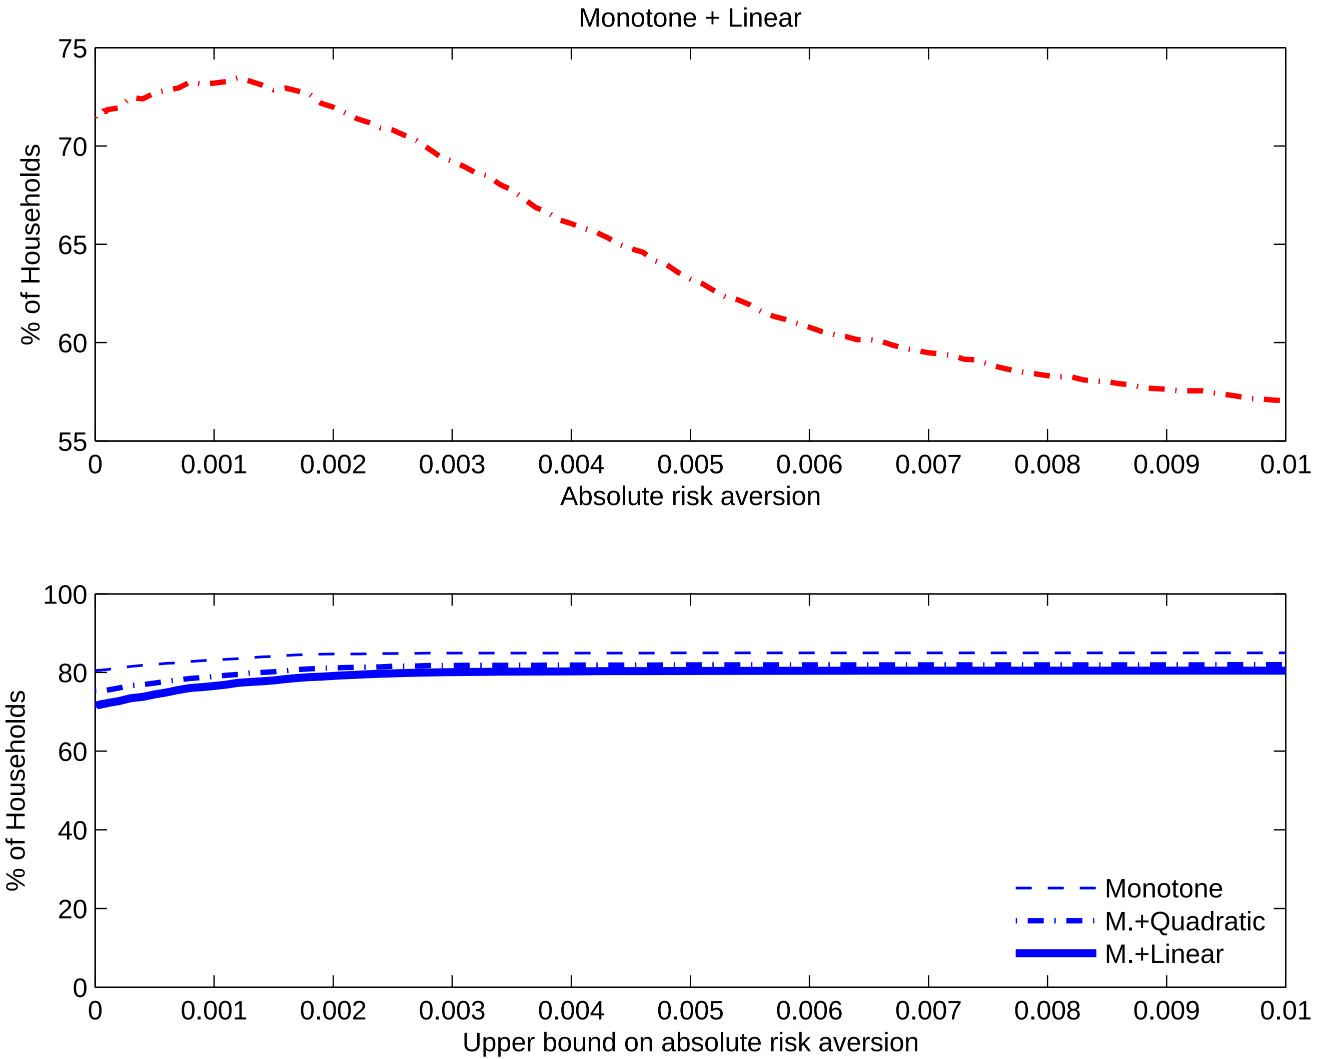

Monotone + Linear
% of Households
Absolute risk aversion
75
70
65
60
55
0
0.001
0.002
0.003
0.004
0.005
0.006
0.007
0.008
0.009
0.01
% of Households
Upper bound on absolute risk aversion
100
80
60
40
20
0
Monotone
M.+Quadratic
M.+Linear

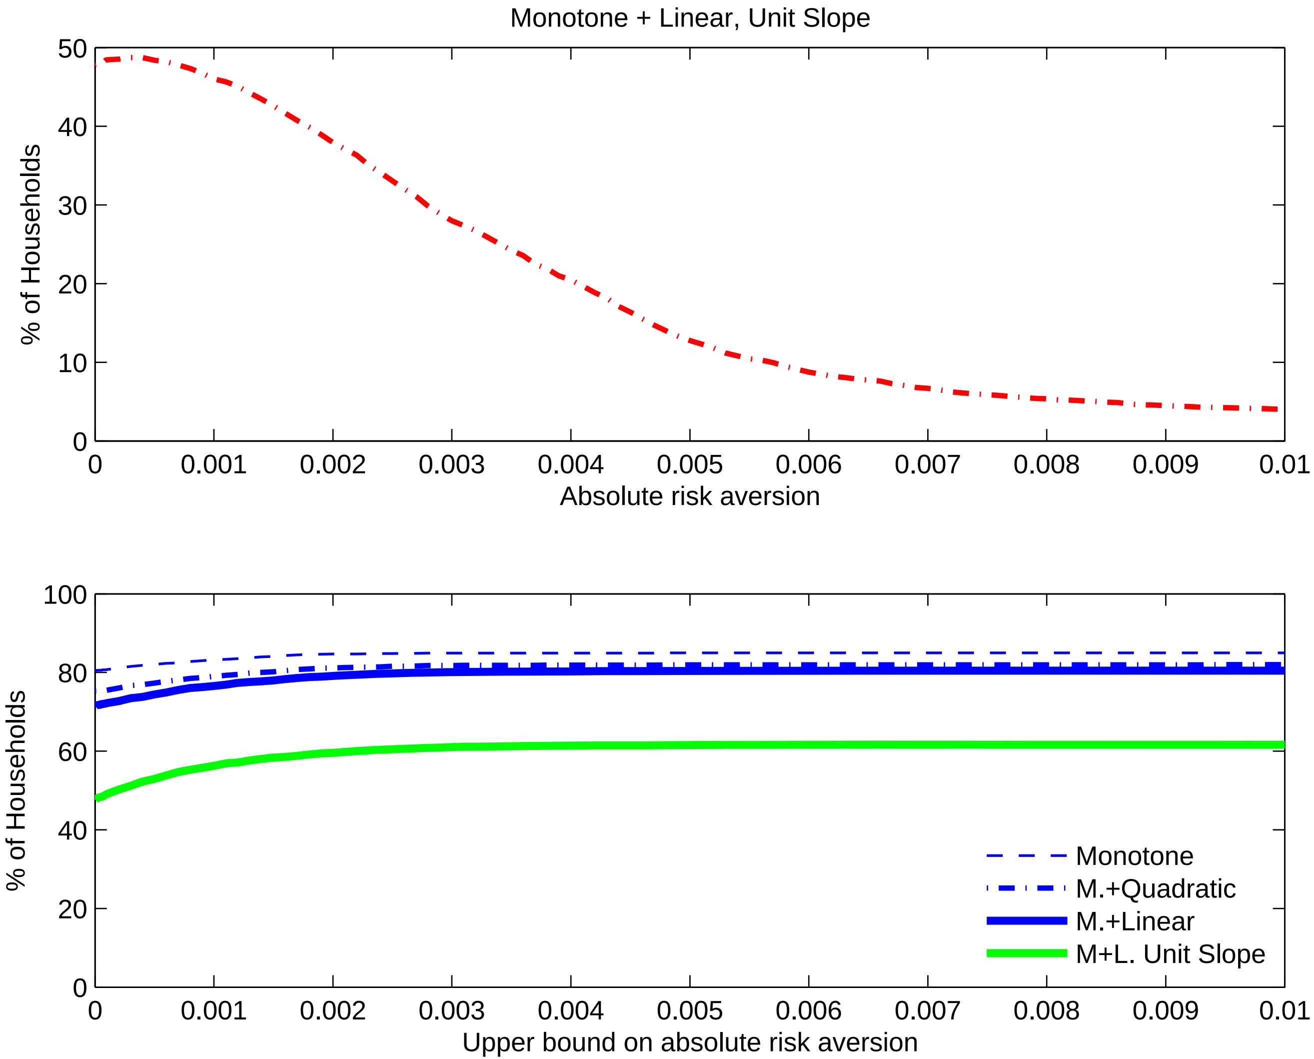

Monotone + Linear, Unit Slope
% of Households
0
10
20
30
40
50
0
0.001
0.002
0.003
0.004
0.005
0.006
0.007
0.008
0.009
0.01
Absolute risk aversion
% of Households
0
20
40
60
80
100
Monotone
M.+Quadratic
M.+Linear
M+L. Unit Slope
Upper bound on absolute risk aversion

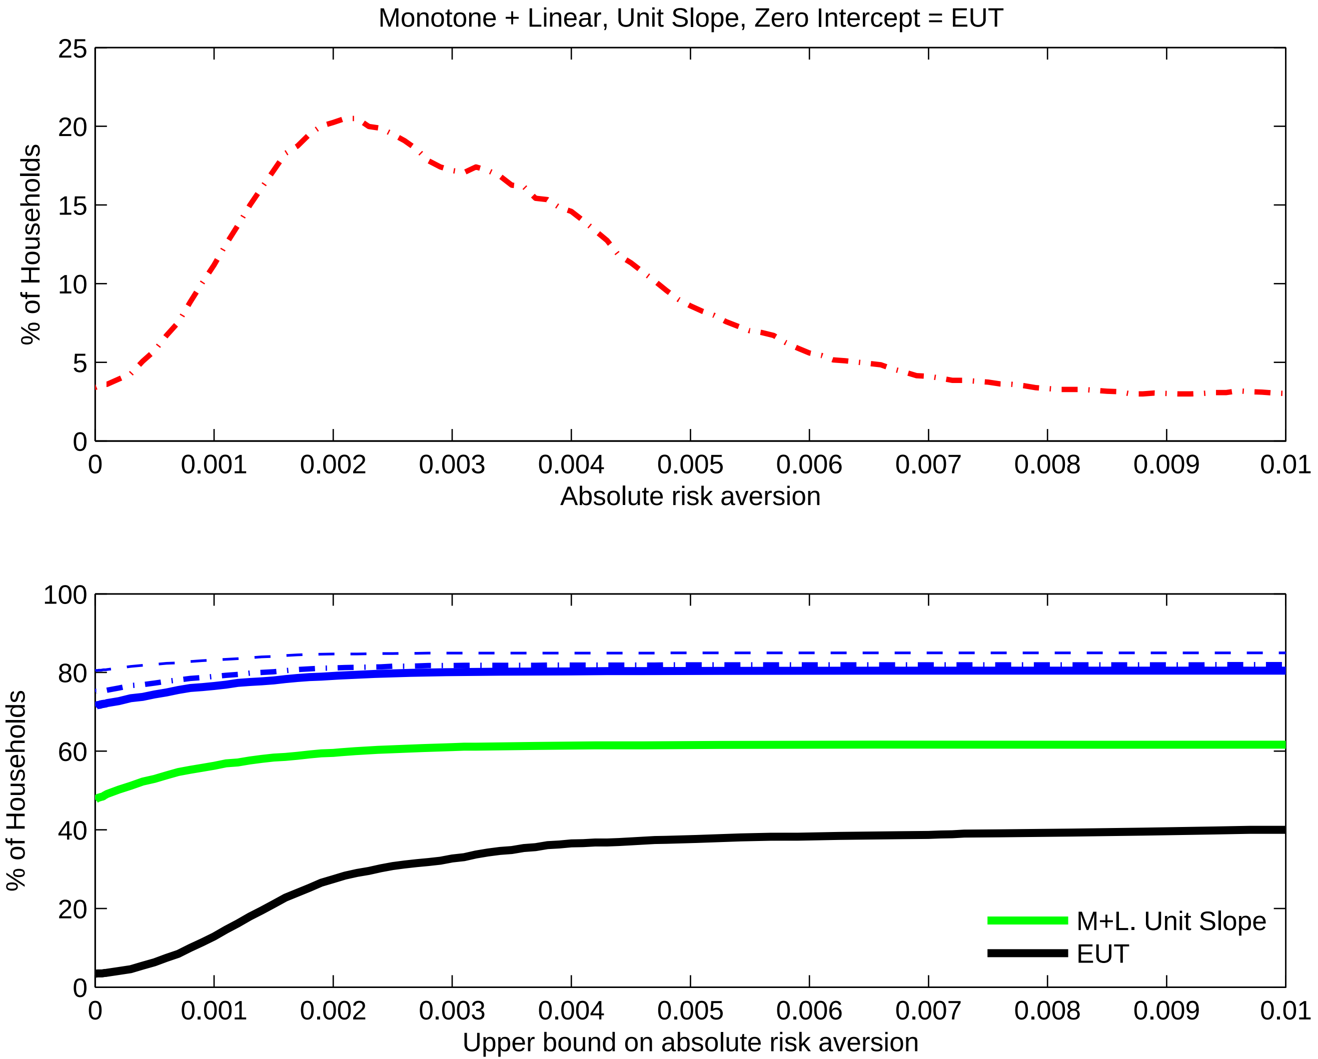

Monotone + Linear, Unit Slope, Zero Intercept = EUT
% of Households
25
20
15
10
5
0
0
0.001
0.002
0.003
0.004
0.005
0.006
0.007
0.008
0.009
0.01
Absolute risk aversion
% of Households
100
80
60
40
20
0
M+L. Unit Slope
EUT
Upper bound on absolute risk aversion

Choosing Coef. of Risk Aversion

- For most households there is more than one r for which a given property is satisfied, so the question is how to choose r.
- Note, however, that the larger is r, the harder it is for the model to satisfy any given property.
- Moreover, high r's in our context might seem implausible.
- Solution: For each household choose the smallest r under which a given desired property (e.g., linearity) is satisfied.

average	median	std. dev.
$1.25 \cdot 10^{-4}$	0	$4.75 \cdot 10^{-4}$

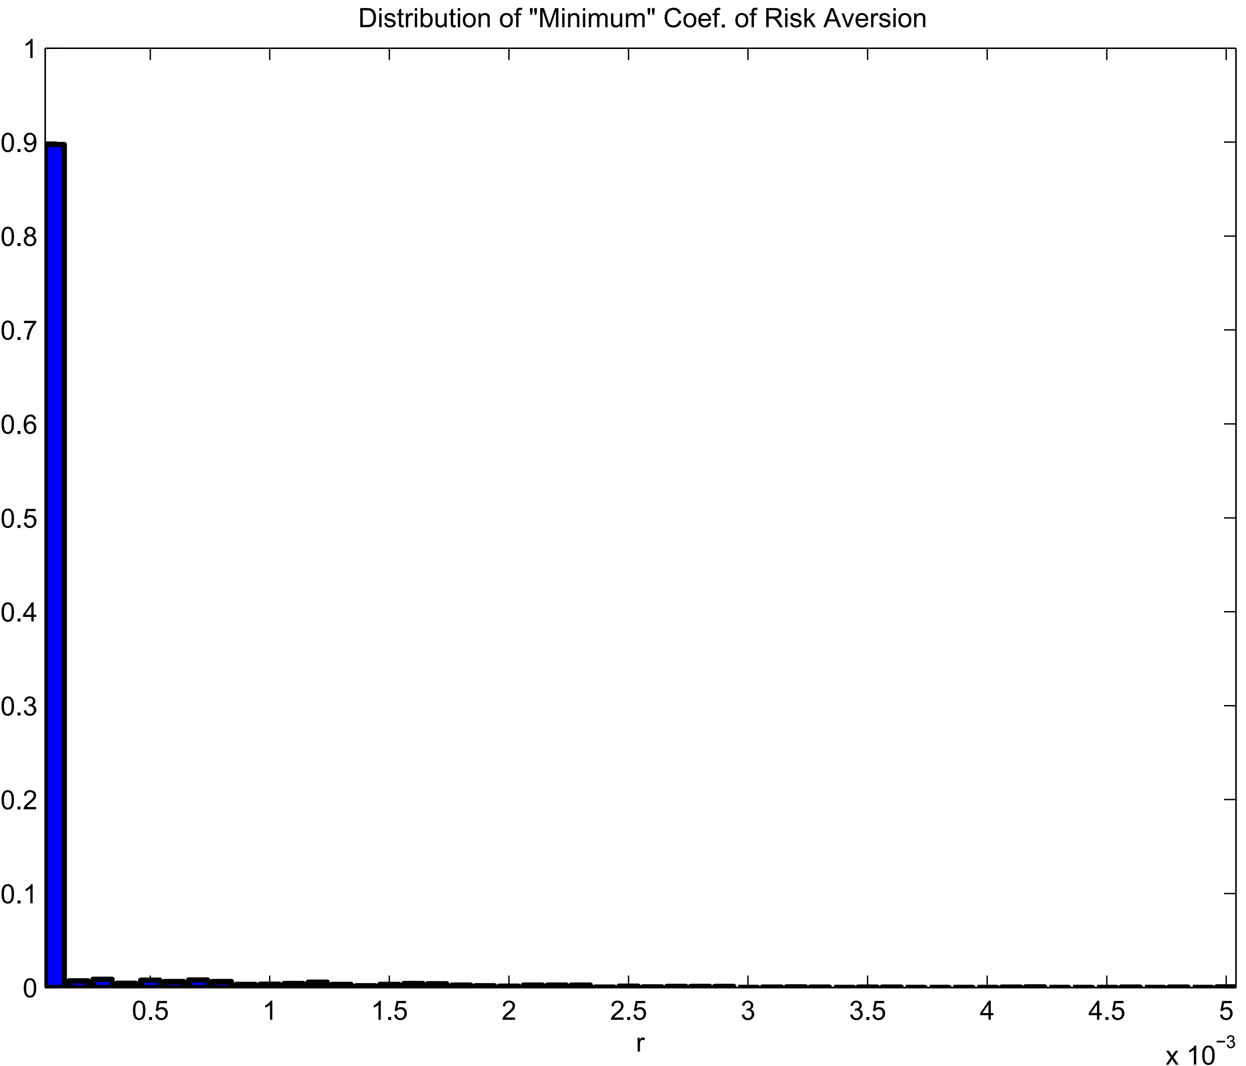

Distribution of "Minimum" Coef. of Risk Aversion
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0.5
1
1.5
2
2.5
3
3.5
4
4.5
5
r
x 10^−3

Shape Restrictions and Inference

As we add shape restrictions:

- The model can rationalize the choices of a smaller fraction of households.
- However, inference about the probability distortion function becomes sharper:
  - The implied intervals for $\Omega(\mu)$ become tighter.
  - With monotonicity+linearity the intervals shrink by 42%.

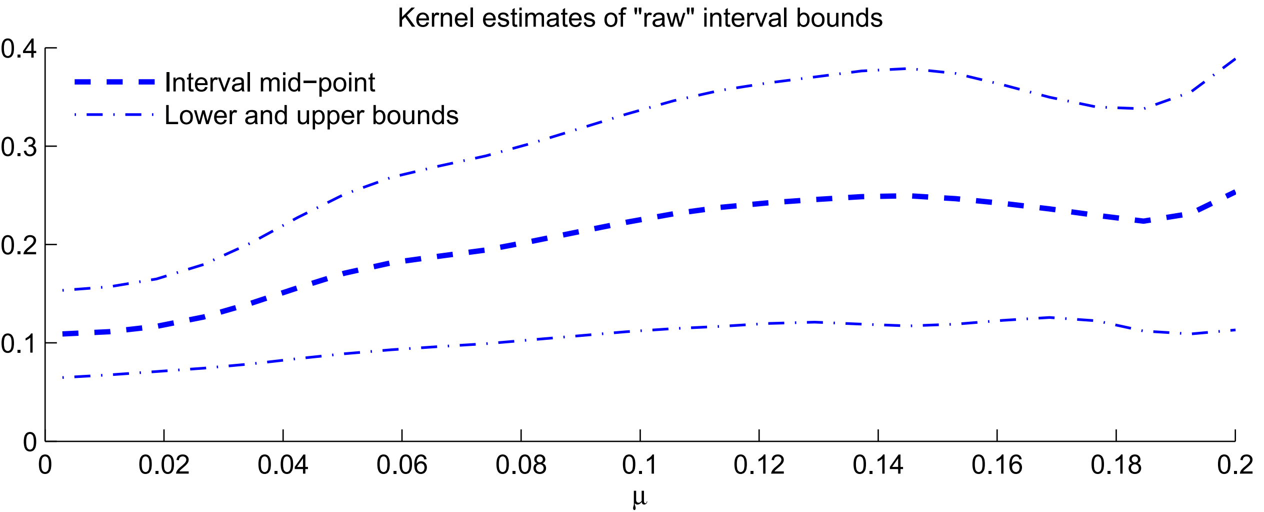
Kernel estimates of "raw" interval bounds
Interval mid−point
Lower and upper bounds
0
0.1
0.2
0.3
0.4
0
0.02
0.04
0.06
0.08
0.1
0.12
0.14
0.16
0.18
0.2
μ

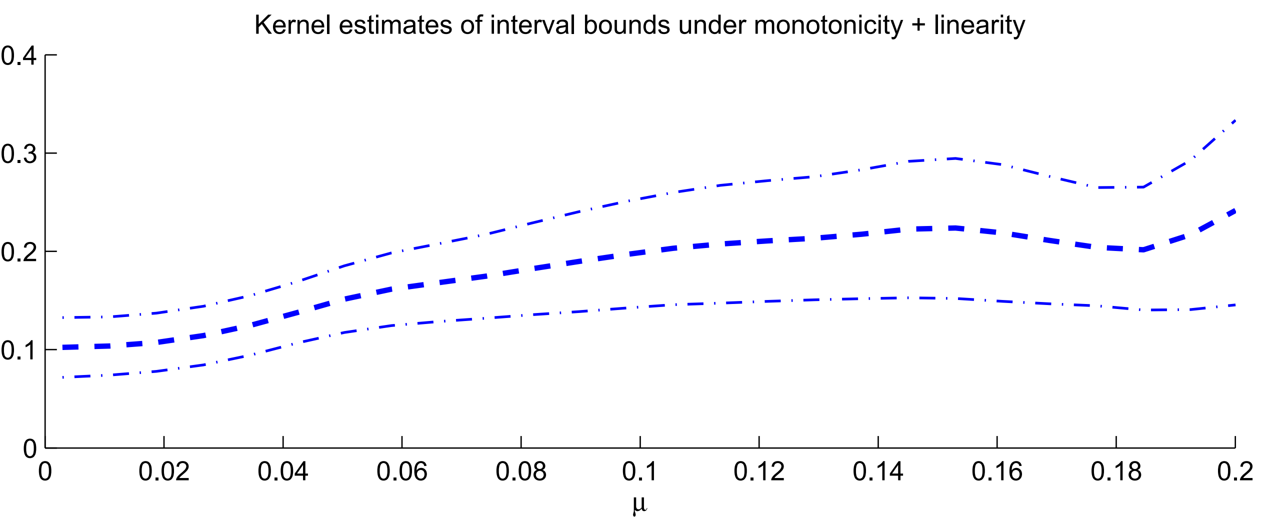
Kernel estimates of interval bounds under monotonicity + linearity
0
0.1
0.2
0.3
0.4
0
0.02
0.04
0.06
0.08
0.1
0.12
0.14
0.16
0.18
0.2
μ

Stability in Sample under Linearity

- What would be the preferences that are "stable" across the sample?
- Denote $\Psi_i^{stable} \equiv \{(a, b)\}_i$ the collection of all intercepts (a's) and slopes (b's) that are consistent with household i's preferences being stable.
- Suppose $\bar{\Psi}^{stable} = \cap_{i=1:N} \Psi_i^{stable}$ is not empty. Then any point in $\bar{\Psi}$ would give preferences that are stable across households in our sample.
- But the set $\bar{\Psi}$ is empty:
  1. There are "unstable" households.
  2. For households who are "stable," there is substantial variation in their Ψ_i^{stable} sets.

Stability in Sample under Linearity

- Hence, we need to pose a more modest question: What are the preferences that are most common in the sub-sample of households that have stable preferences?

- Or alternatively, what are the preferences that are as close as possible to being homogeneous?

- We **propose and estimate** a best linear point predictor:

$$\Omega^{LP}(m) = \alpha + \beta \cdot m,$$

  that minimizes the expected Euclidean distance to the random intervals $\Omega_i(m)$'s, where $\Omega_i(m)$'s are constructed based on Ψ_i^{stable}.

- In other words, we seek an (α, β) that generates preferences as close as possible to being homogeneous, where "close" is defined in terms of Euclidean distance.

- We call it the minimum distance Ω.

Stability in Sample under Linearity

- We prove that under mild conditions (satisfied in our data) the pair (α, β) is point identified.
- We establish consistency and asymptotic normality of our sample analog estimator.
- The asymptotic distribution of the estimator is consistently approximated by non-parametric bootstrap.

Stability in Sample under Linearity

- We find:

$$\hat{\alpha} = .087 \text{ (s.e. } .0017)$$
$$\hat{\beta} = .706 \text{ (s.e. } .026)$$

- More general version allows for α and β to be functions of observables.

- That is, conditional on observables, we seek an intercept and a slope that generate preferences as close as possible to being homogeneous, where "close" is defined in terms of Euclidean distance.

- We do this by imposing that for household i :

$$\log \alpha_i = X_i \alpha_x$$
$$\log \beta_i = X_i \beta_x$$

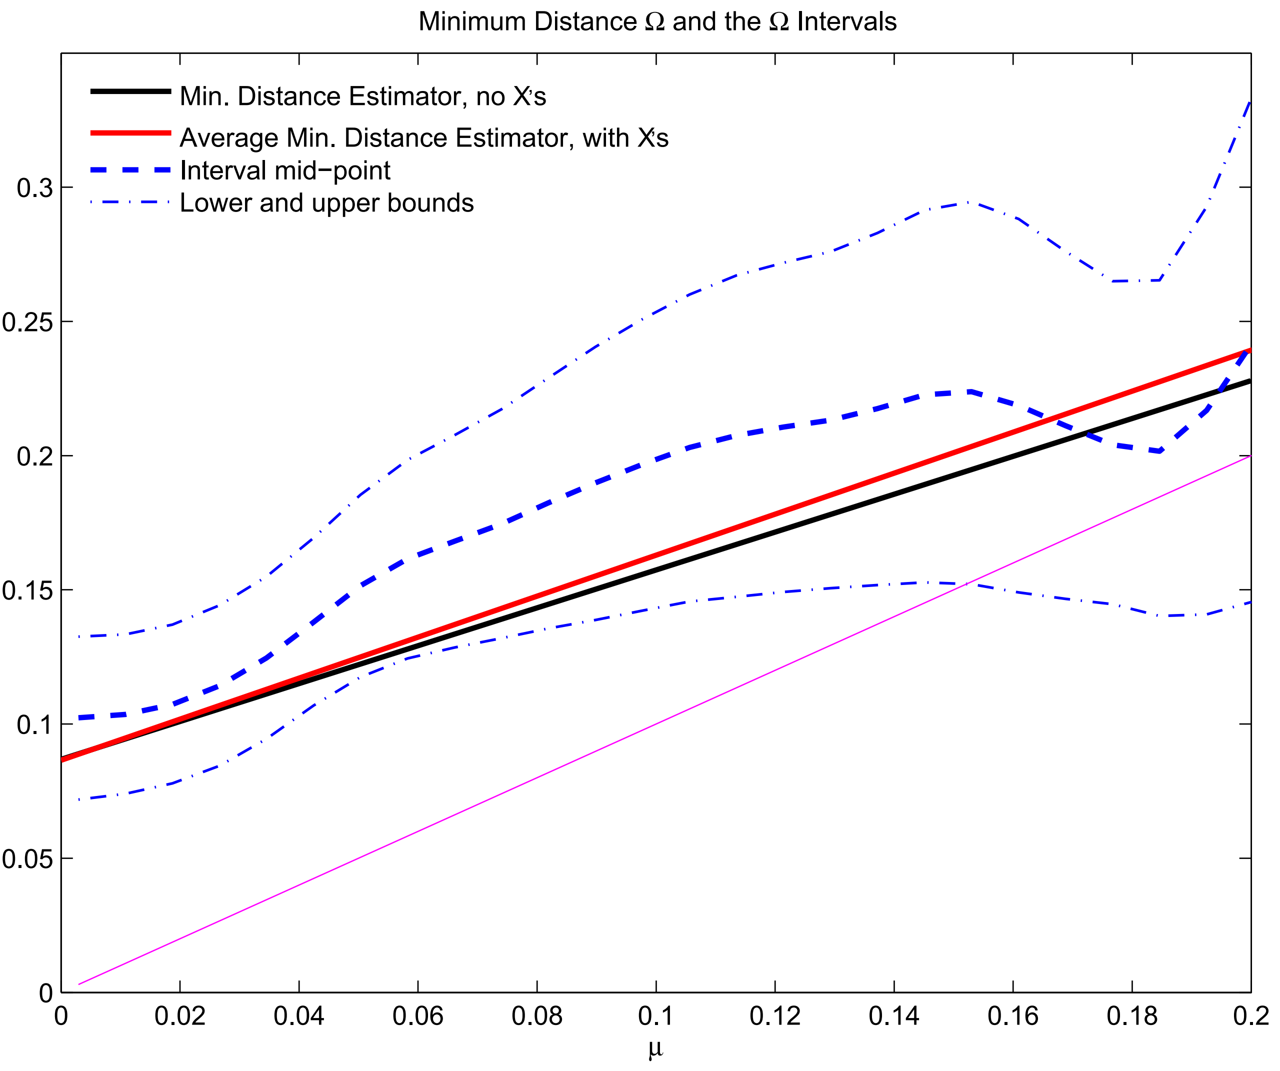

Minimum Distance Ω and the Ω Intervals
Min. Distance Estimator, no X's
Average Min. Distance Estimator, with X's
Interval mid−point
Lower and upper bounds
0.3
0.25
0.2
0.15
0.1
0.05
0
0
0.02
0.04
0.06
0.08
0.1
0.12
0.14
0.16
0.18
0.2
μ

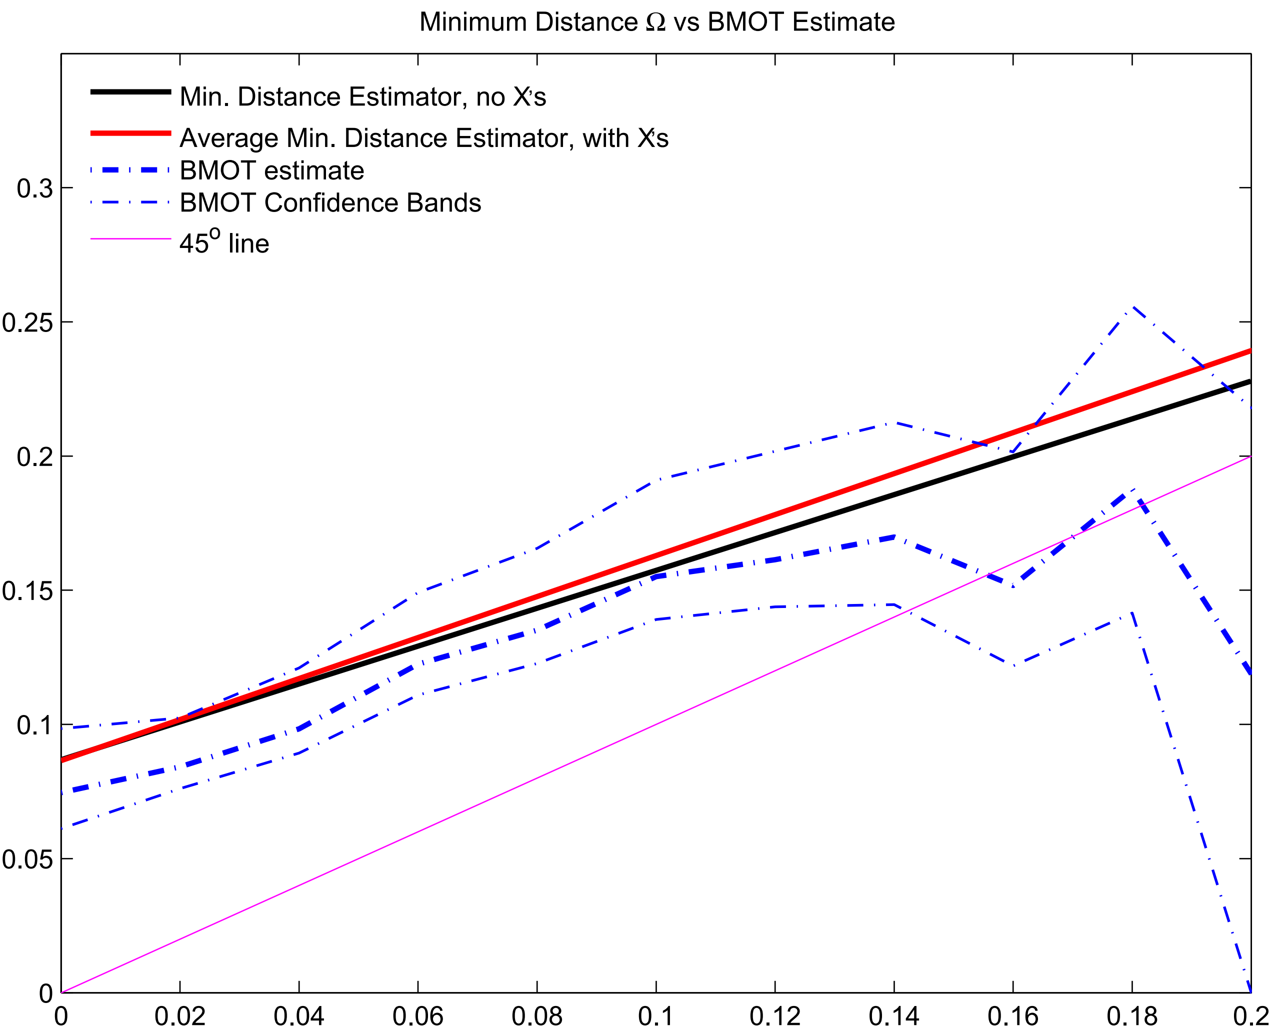

Minimum Distance Ω vs BMOT Estimate
Min. Distance Estimator, no X's
Average Min. Distance Estimator, with X's
BMOT estimate
BMOT Confidence Bands
45° line
0
0.05
0.1
0.15
0.2
0.25
0.3
0
0.02
0.04
0.06
0.08
0.1
0.12
0.14
0.16
0.18
0.2

Explanatory Power of the Minimum Distance Estimator

- The function $\Omega(\mu) = 0.087 + .706\mu$ (at the minimum r) is consistent with:
  - all three choices of 18% of "linear" households;
  - at least two choices of roughly a third of "linear" households;
  - at least one choice of roughly 60% of "linear" households.

- We measure how "close" this function comes to being consistent with choices of the remaining households.
  - Define close (within some "tolerance") as follows: the (maximum across the three coverages) distance between Ω^{LP} and the household's $\Omega(\mu)$ intervals is $\leq$ tolerance.

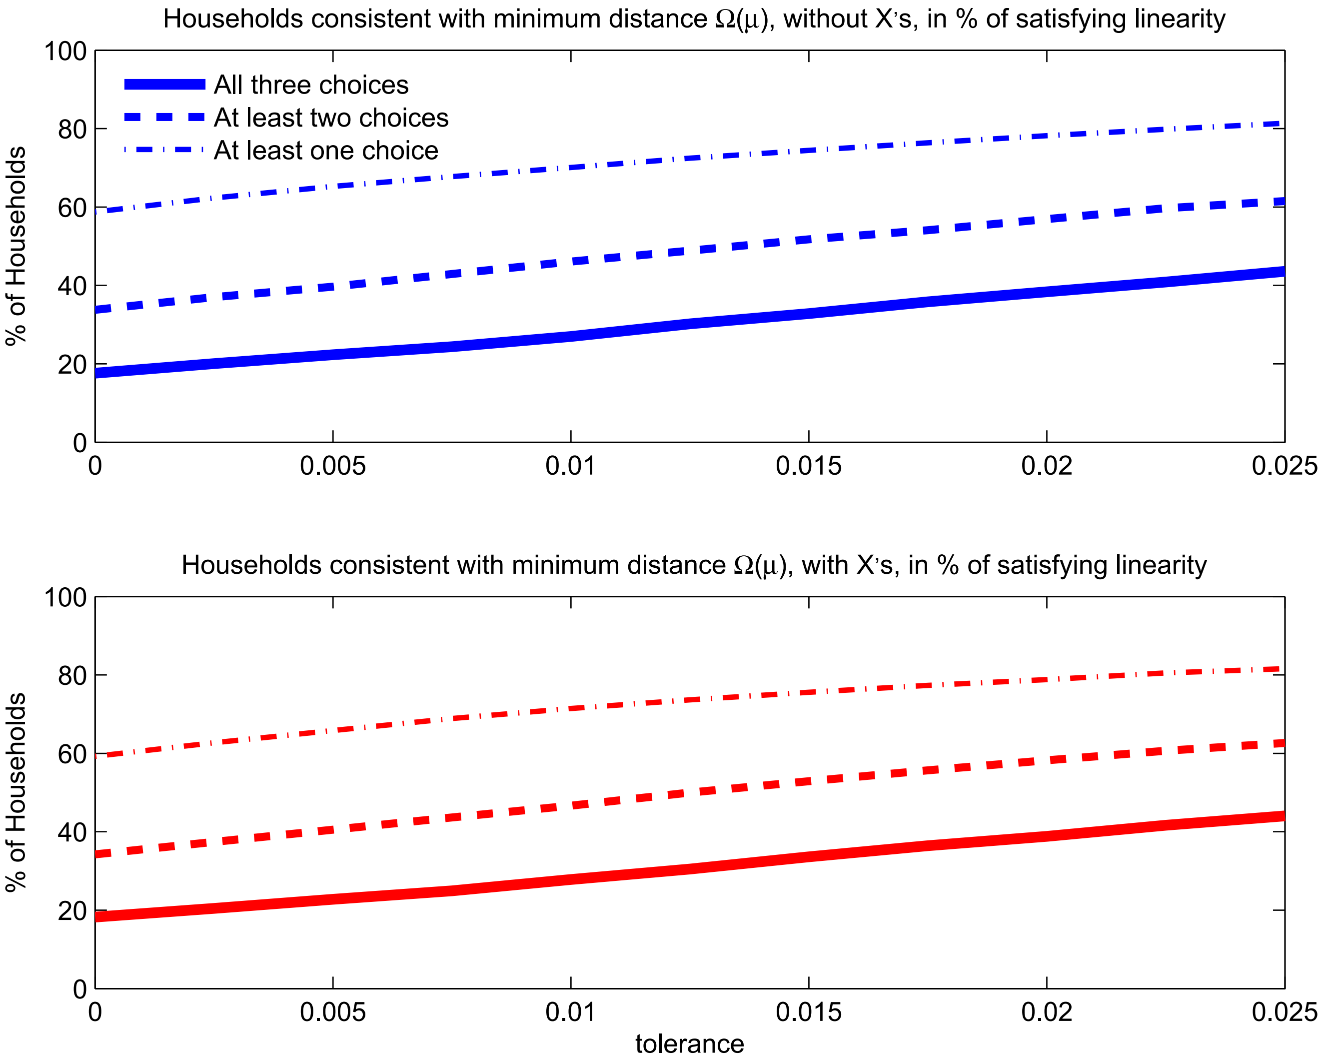
Households consistent with minimum distance Ω(μ), without X's, in % of satisfying linearity
All three choices
At least two choices
At least one choice
% of Households
0
20
40
60
80
100
0
0.005
0.01
0.015
0.02
0.025
Households consistent with minimum distance Ω(μ), with X's, in % of satisfying linearity
% of Households
0
20
40
60
80
100
0
0.005
0.01
0.015
0.02
0.025
tolerance

Stability (Linearity) and Rank Correlations

- Households who satisfy linearity drive up rank correlations across choices:

Coverage	Total	Stable	Unstable
Collision and Comprehensive	0.48***	0.56***	0.36***
Collision and Home	0.29***	0.41***	-0.05
Comprehensive and Home	0.28***	0.36***	-0.06

Stability: Alternative Models

How does our model compare to others through the prism of stability criterion?

Our Model				
Quad.+ varying r	Linear.+ varying r	Quad. $r = 0$	Linear $r = 0$	Unit Slope $r = 0$
82%	81%	74%	71%	48%

Alternatives				
Other P-Distortions		EUT		Kőszegi
Prelec	KT-power	CARA	NTD	Rabin
44%	40%	40%	33%	25%

Unobserved Heterogeneity in Claim Probabilities

- The failure of some households to satisfy a given property may be the result of not accounting for unobserved heterogeneity in risk.
- However, the intervals for Ω per se do not depend on μ_{ij}'s, only their relative order does.
  - For unobs. heterogeneity to matter it should "switch" the relative order of claim probabilities – the unobserved heterogeneity in claim prob. should be negatively correlated across contexts. This is not the case in the data (Barseghyan, Molinari, Morris, and Teitelbaum 2012).

Unobserved Heterogeneity in Claim Probabilities

- To assess the potential effect of unobserved heterogeneity, we compute the *expected* fraction of households that satisfy a given property:
  - For each coverage and each household we randomly draw an unobserved heterogeneity term, ε_{ij}, from the distribution we estimated in the claim rate regressions and use $\tilde{\lambda} = \hat{\lambda} \cdot \exp(\varepsilon_{ij})$ as the claim rate.
  - Using these claim rates we compute the fraction of households that satisfy a given property.
  - We repeat the above 200 times.

	At $\hat{\lambda}$	Distribution based on simulated λ's						
		5%	10%	25%	50%	75%	90%	95%
Linearity	80.5	80.4	80.6	80.8	80.9	81.1	81.3	81.5
EUT	40.4	31.8	32.1	32.4	32.9	33.3	33.7	33.8

Conclusion

- Choices of majority of households conform to decision making with distorted probabilities.
  - Their choices across three different contexts can be rationalized by a single distortion function (and a realistically low curvature of the utility function).
- We show how restrictions on the shape of the distortion function sharpen the inference.
- We construct an estimator that yields the distortion function which comes closest in the Euclidean sense to generating distorted probabilities consistent with the households' choices across our sample. In doing so we rely only on a revealed preference argument (and the model of decision making under risk).