NBER WORKING PAPER SERIES

RISK PREFERENCES IN SMALL AND LARGE STAKES: EVIDENCE FROM INSURANCE CONTRACT DECISIONS

Benjamin L. Collier
Daniel Schwartz
Howard C. Kunreuther
Erwann O. Michel-Kerjan

Working Paper 23579
http://www.nber.org/papers/w23579

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2017

We thank Levon Barseghyan, Georges Dionne, Glenn Harrison, Johannes Jasperson, Dylan Small and seminar participants in the National Bureau of Economic Research Insurance Workshop, Risk Theory Society, American Risk and Insurance Association, and Society for Judgment and Decision Making for their guidance and insightful comments. We thank Juan Zhang for her research assistance. This work is partially supported by the Zurich Insurance Foundation, the Center for Risk and Economic Analysis of Terrorism Events (CREATE) at the University of Southern California (U.S. Department of Homeland Security's Center of Excellence), the Travelers-Wharton Partnership for Risk Management Fund, and the Wharton Risk Management and Decision Processes Center. D. Schwartz was supported by the Complex Engineering Systems Institute (CONICYT – PIA – FB0816; ICM P-05-004-F). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Risk Preferences in Small and Large Stakes: Evidence from Insurance Contract Decisions
Benjamin L. Collier, Daniel Schwartz, Howard C. Kunreuther, and Erwann O. Michel-Kerjan
NBER Working Paper No. 23579
July 2017
JEL No. D12,D81,H42,Q54

ABSTRACT

We examine risk preferences using the flood insurance decisions of over 100,000 households. In each contract, households make a small stakes decision, the deductible, and a large stakes one, the coverage limit. Expected utility models predict that households would choose high deductibles and low coverage limits, but households do the opposite. Allowing for probability distortions improves our models. Assessing rank dependent utility models, we find that households follow two tenants of prospect theory: overestimation of small probabilities and diminishing sensitivity to losses. In every tested model, different preferences characterize households' small and large stakes insurance decisions.

Benjamin L. Collier
Department of Risk, Insurance,
and Healthcare Management
Fox School of Business
1801 Liacouras Walk
Philadelphia, PA 19122
collier@temple.edu

Daniel Schwartz
Department of Industrial Engineering
University of Chile
Republica 701
Santiago, Chile
daschwar@dii.uchile.cl

Howard C. Kunreuther
Operations and Information Management
The Wharton School
University of Pennsylvania
3730 Walnut Street, 500 JMHH
Philadelphia, PA 19104-6366
and NBER
kunreuther@wharton.upenn.edu

Erwann O. Michel-Kerjan
The Wharton School
University of Pennsylvania
Center for Risk Management
3730 Walnut Street, 556 JMHH
Philadelphia, PA 19104-6340
erwannmk@wharton.upenn.edu

1 Introduction

Explaining individuals' insurance decisions is a cornerstone of economic theory. Insurance choices helped motivate the concept of expected utility (Bernoulli, 1738) and remain among households' most consequential decisions today. Insurance choices are also fertile ground for understanding risk preferences in the field as contract options mirror the well-defined financial lotteries used in experimental settings.

This paper analyzes risk preferences implied by the small and large stakes flood insurance decisions of over 100,000 households in flood-prone areas of the United States. The insurance contract protects the residence against flood damage. Households make two choices along a continuum of potential losses: the deductible represents a small stakes decision and the coverage limit (i.e., the maximum an insurance claim would pay) a large stakes decision. While empirical research has focused on deductible choices, our data enable us to examine households' preferences in choosing both a deductible and coverage limit. We consider two research questions.

First, we examine whether households' small and large stakes insurance decisions are characterized by similar risk preferences in a context that encourages making deductible and coverage choices in concert. Recent research shows that households' deductible decisions can be difficult to reconcile with expected utility theory. For example, Sydnor (2010) finds that homeowners insurance deductibles indicate triple digit relative risk aversion. Rabin and Thaler (2001) attribute small stakes preferences to decision isolation, the tendency to evaluate outcomes narrowly rather than, as most applications of utility theory posit, in terms of the individual's total wealth.[2] For example, households' risk preferences differ across insurance domains such as auto liability, auto comprehensive, and homeowners (Barseghyan, Prince, and Teitelbaum, 2011). In our setting, households' choices are made in the context of a single premium payment to insure a single asset, their property, against a single hazard, flood damage. Do households' risk preferences differ across their deductible and coverage limit choices in the same insurance contract?

Second, we examine whether households' large stakes insurance preferences align with expected utility theory. Even its critics often concede that while explaining small stakes

[2] They also identify loss aversion as an explanation.

decisions can be difficult for it, “expected utility theory certainly captures some of the intuition for risk aversion over very large stakes” (Rabin and Thaler, 2001, p. 224). Economists may justifiably embrace expected utility theory if it describes households’ most economically important choices.

We model households as constant relative risk aversion (CRRA) expected utility maximizers and, using maximum likelihood estimation (MLE), identify the coefficient of relative risk aversion that best characterizes households’ deductible choices and the coefficient that best characterizes their coverage limit choices.[3] We also allow relative risk aversion to differ across households, modeling it as a function of variables that have been shown to influence individuals’ risky decisions (e.g., recent severe events and proxies for wealth).

We find that households’ small and large stakes flood insurance choices are difficult to explain with standard models of expected utility theory. These models predict that households would choose high deductibles and low coverage limits, but households tend to do neither. Instead, households’ low and high stakes decisions are characterized by different preferences. Households’ deductible decisions reflect a much higher coefficient of relative risk aversion than their coverage limit decisions. Modeling households as constant absolute risk aversion (CARA) expected utility maximizers yields a similar result. We also consider two theories intended to unify households’ preferences across small and large stakes; neither Chetty and Szeidl’s (2007) consumption commitments nor Holt and Laury’s (2002) explanation of expo-power utility seem to explain households’ preferences in our data.

Like Sydnor (2010) and in the spirit of Rabin (2000), we find that households strongly prefer low deductibles, over 94 percent choose one of the two lowest deductible options. Explaining these small stakes preferences requires triple digit relative risk aversion, $\rho = 139$ in our case where ρ is the coefficient of relative risk aversion.

Regarding high stakes insurance preferences, assuming that households are CRRA expected utility maximizers results in relative risk aversion estimates that align with other domains, but do not explain the coverage limit decisions well. Like Gourinchas and Parker (2002), Chetty (2006), Barro and Jin (2011) and others, we find single digit risk aversion when the stakes are high; households’ coverage limit preferences are best explained by $\rho = 2.7$. While this value of

[3] Wakker (2008) reports that CRRA utility is “the most widely used parametric family for fitting utility functions to data” (p.1329).

relative risk aversion is the MLE estimate, this model dramatically under-predicts households' coverage limits: it under-predicts the average household coverage limit ratio (coverage relative to the home's replacement cost) by 34 percentage points. These results hold even when allowing for heterogeneity in risk aversion across households: the heterogeneous model results in a distribution of relative risk aversion around the homogeneous model estimate (e.g., for the coverage limit, a median $\rho = 2.78$ and a mean of 2.81).

The problem with explaining households' high stakes insurance preferences arises from a well-known theoretical result (Mossin, 1968): except under extreme risk aversion, individuals offered an insurance premium that is priced above the actuarially fair rate will partially insure (select a coverage limit that is smaller than their replacement cost in our context). Almost all (97 percent) of the households in our data pay premiums that are above the actuarially fair rate, which is nearly always the case in insurance markets (e.g., because of administrative loads). Thus, theory predicts that *none* of these households would fully insure, but 77 percent do.

We also assess whether households distort probabilities following rank dependent utility (Quiggin, 1982). Substantial evidence indicates that individuals' decisions under risk are marked by misweighting and/or misperceiving the probability of an event (e.g., Preston and Baratta, 1948; Prelec, 1998; Barseghyan et al., 2013). When allowing for probability distortions, we find small and large stakes preferences that are consistent with two of the tenants of prospect theory: households 1) overweight small probabilities, and 2) demonstrate diminishing sensitivity to losses (Kahneman and Tversky, 1979).[4]

In the CRRA rank dependent utility model, households' demand for insurance is explained by their overweighting of small probabilities. This model fits our data best among all the models that we assess (e.g., according to Vuong tests). On average CRRA rank dependent utility over-predicts households' coverage limit ratios by 2 percentage points. Despite its improvements in explaining households' decisions, the CRRA rank dependent utility model continues to indicate that households' choices are characterized by different preferences across small and large stakes. Low stakes preferences show much greater diminishing sensitivity to losses than high stakes

[4] Diminishing sensitivity to losses indicates that households perceive the difference between a $1,000 loss and $2,000 loss as greater than that between a $2,000 and $3,000 loss. Diminishing sensitivity to losses is contrary to the predictions of the standard expected utility of wealth model for a risk averse household.

preferences, and households overweight flood loss probabilities when choosing a deductible at almost twice the level they do when selecting a coverage limit.

Our setting overcomes several constraints that typically challenge examining large stakes preferences with insurance data. First, a household's maximum flood exposure is clearly defined. The potential maximum loss is limited by the value of the home structure, and the flood insurance program provides a specific estimate of that value to households during the application process. In many other cases, households may have difficulty assessing their maximum exposure, such as the healthcare costs associated with unanticipated illnesses. Second, households have greater flexibility in selecting a coverage limit than is typical of property insurance as the flood insurance program does not maintain a coinsurance clause. Third, the insurance covers a single hazard, flood. Other insurance products often cover several hazards, and households' exposures may differ between these hazards, especially constraining the researcher's ability to model rare, consequential outcomes. Finally, almost all (96 percent) of residential flood insurance policies in the U.S. are purchased through this program (Dixon et al., 2006, p.19). All households that we study have the same deductible and contract limit choices and their premiums are priced using the same rules. This structure reduces the endogenous selection problems that can arise from studying the policies of private insurers whose pricing and contracts may differ from their competitors.[5]

The rest of the paper proceeds as follows. Section 2 describes relevant research. Section 3 describes the setting and data, households' contract choices, and our structural approach to estimating households' risk preferences. Section 4 provides our main results. Section 5 explores robustness and considers alternative explanations for these results. Section 6 concludes and discusses implications of our research.

2 Relevant Research

Recent risk preference estimates are typically derived from three broadly defined lines of research: 1) experiments and surveys, 2) major life decisions such as financial investing and

[5] Private property insurance, such as a homeowners' policy normally includes a "coinsurance clause," which requires that policyholders select a coverage limit of at least 80 percent of their property value to be fully reimbursed for damage above the deductible. A large amount of information on flood risk is publicly available, and our data provides detailed information on each insured property (e.g., the approximate likelihood of a claim, the presence of flood obstructing barriers, etc.).

labor supply, and 3) insurance decisions.[6] Experiments are often characterized by strong internal validity and clearly defined choice sets, e.g., asking individuals to select a lottery among several alternatives. Harrison and Rutström (2008) discuss estimating risk aversion in laboratory settings and report many findings with single digit relative risk aversion (or smaller).[7] Experiments also demonstrate that the participants' preferences can be highly influenced by the choice environment as the same individual can demonstrate risk aversion or risk seeking depending on the task (Berg, Dickhaut, and McCabe, 2005; Hey, Morone, and Schmidt, 2009).

A second literature has inferred risk aversion based on expected utility theory from individuals' major decisions, for example, through consideration of the equity premium puzzle (e.g., Barro and Jin, 2011), consumption over the life cycle (Gourinchas and Parket, 2002), and labor supply (Chetty, 2006). These studies often assume a utility function such as CRRA then identify the risk aversion coefficient that best explains their data, typically finding relative risk aversion around 1 or a little higher (e.g., 3 for Barro and Jin, 2011). These choices are undoubtedly consequential, but involve more complex decisions than the clearly defined lotteries posed in experimental settings. For example, a households' labor supply decision might include working more at the same job, taking a second job, working less to seek education, etc., and the choice set differs across households. Without the ability to identify a household's choice set, it is difficult to assess how well derived risk aversion explains a household's decision among its alternatives.

Insurance contract choices represent a third alternative, which Cohen and Einav (2007) call "(almost) an ideal setting for estimating risk aversion" (p. 745). For example, Sydnor (2010) assesses homeowners' insurance deductible choices using CRRA and CARA expected utility models. Over 80 percent of households select a low deductible, and he finds that their deductible decisions imply extremely high risk aversion (e.g., coefficients of relative risk aversion in the hundreds or thousands). He shows that cumulative prospect theory can more plausibly explain households' choices.

---

[6] Barseghyan et al. (2015) provide an excellent overview of field research on risk preferences and discuss research on betting decisions as another important area. Several features of betting markets and preferences in this context (e.g., longshot bias) make them less relevant to our research and so we defer to their discussion on this topic.

[7] While a commonly expressed concern about laboratory settings is that they often examine small stakes decisions, several studies address this concern by increasing the stakes and showing similar preferences for more consequential decisions (e.g., Holt and Laury, 2002; see Harrison and List, 2004 for a discussion).

Our paper is closest methodologically to Barseghyan et al. (2013) who use McFadden's (1974) random utility framework to analyze households' home and auto deductible decisions. They find evidence of risk aversion and probability distortions in households' choices. Our work complements theirs by analyzing both deductible and coverage limit decisions and examining model predictions in more detail.

Baresghyan, Prince, and Teitelbaum (2011) and Einav et al. (2012) examine the extent to which households' risky decisions are consistent across domains. Using households' deductible choices for homeowners, auto liability, and auto comprehensive insurance, Barseghyan, Prince, and Teitelbaum (2011) reject the null hypothesis that households demonstrate consistent risk preferences. Households demonstrate greater risk aversion in home than in auto deductible choices. Einav et al. (2012) examine the insurance coverage (health, prescription drug, dental, and short-term and long-term disability) and 401(k) investment decisions of the employees of a large U.S.-based corporation. These decisions represent both smaller stakes (e.g., health insurance deductible) and larger stakes (e.g., long-term disability coverage) choices. They order individuals by the implied risk of each choice and find that individuals making relatively riskier decisions in one domain are more likely to do so in another (resulting in an average Spearman rank correlation of 0.19 across insurance coverages). We complement their research by examining two decisions along the continuum of a single loss distribution for each household.

Risk preferences research has increasingly focused on heterogeneity in preferences. This research considers systematic differences across individuals (e.g., based on age, gender, nationality) and the distribution of risk aversion in a population. For example, Cohen and Einav (2007) examine individual's auto deductibles selections and find a right-skewed distribution of risk preferences, estimating an average coefficient of absolute risk aversion of 0.0031 (median of 0.000035). In another notable example, Harrison, Humphrey, and Verschoor (2010) estimate risk preferences from lab experiments in Ethiopia, India, and Uganda. Imposing CRRA expected utility, they find that a model assuming homogeneous preferences results in estimated relative risk aversion of around 0.5. Including covariates (e.g., nationality, age, gender) creates a right-skewed distribution of risk preferences, though not to the degree of Cohen and Einav (2007), with both the mean and median between 0.5 and 0.6. We examine heterogeneity in preferences, evaluating how allowing risk aversion to vary across households affects our estimates of risk aversion and model predictions.

3 Methods

3.1 Setting and Data

Homeowners insurance in the U.S. typically does not cover flood risk; protection from flood can be purchased as a standalone contract from the federally-run National Flood Insurance Program (NFIP). We examine contracts that insure the house against flood.[8] The U.S. federal government operates the NFIP through the Federal Emergency Management Agency (FEMA) and underwrites all insured risk. Households buy flood insurance from an authorized insurer or insurance agent who receives a commission intended to cover its origination and administrative costs. In 2013, the NFIP issued 5.6 million policies for a total insured value about $1.3 trillion.

Our data include all policies for single-family dwellings from 2001 to 2009 insured by the NFIP and all its claims from 1982 to 2009, resulting in 16,349,345 policy observations and 635,220 claims observations for these households. We use these data to model flood risk for each policyholder, estimating the probability of a claim and the distribution of losses given a claim.

3.1.1 Baseline Sample

The dataset contains specific contract choices of households' policies with claims between 2003 and 2008.[9] These data include the contract selected by the policyholder and the characteristics of the insured home needed to calculate premiums for any potential policy available to a household.

The data also include the value and replacement cost of the insured property. A home's property value is assessed onsite at the time of a claim and is used in measuring total flood damage to the home. Homeowners are provided an estimate of their replacement cost when they select an insurance policy. Replacement costs provide the amount needed to rebuild the current structure using similar materials and are assessed using insurance industry standards that account for a home's characteristics including its building materials, size, and sales value (NFIP, 2006b, p. 4_175). Policyholders most commonly select a coverage limit that equals the home's

---

[8] Additionally, households can purchase flood coverage for the contents of their homes, which we do not examine. Contents coverage has its own premium, deductible, and coverage limit.

[9] The policy and claim databases each contain some distinct variables that are important to estimating preferences, primarily related to modeling the flood risk, which we discuss in Section 3.1.2. We focus on these years because we can replicate households' premiums using the program's pricing formulas well during this time.

replacement cost. The replacement cost and property value are often similar; the average ratio of replacement cost to property value is 1.06, the median is 0.94.

From 2003 to 2008, households experienced a consistent choice of deductible options ranging from $500, $1,000 up to $5,000, in $1,000 increments, and could purchase property coverage limits up to $250,000 in $100 increments. Because of this maximum coverage limit, we examine only homes with values up to $250,000 so that all households in the analyses could select a coverage limit to protect the full replacement cost of their home.[10] We include only a policyholders' first claim so that each insured household is represented once in the analysis. This restriction intends to reduce potential effects on contract selections of learning, which is not incorporated in standard utility models. The resulting baseline sample includes 103,080 observations.

FEMA manages flood maps that classify the risk of each home in broad terms. Households in the baseline sample reside in areas that FEMA estimates have at least a one percent annual flood risk probability but are not prone to coastal storm surge (FEMA designates this "Zone A"). This is by far the largest risk category in the program, accounting for 47 percent (or 2.1 million) of all policies for single-dwelling homes in 2012. In Section 5, we exploit distinct features of two other zones to examine the robustness of our baseline results: (1) homes with at least a one percent annual flood probability that are also vulnerable to storm surge, and (2) homes with less than a one percent annual flood probability. The flood insurance program also treats homes differently depending on whether they were built before or after flood maps were developed. We discuss specific differences below. The idea behind this distinction is that building standards may not have accounted for the vulnerability of at-risk homes before maps were developed, increasing the home's expected flood losses. The baseline sample includes both homes that were built before and after flood maps were developed.

3.1.2 Flood Risk

We estimate households' claim rates and loss distributions based on the characteristics of the home that the flood insurance program uses to set premium rates. Here, we provide an overview

[10] Ninety-five percent of all insuring households have home replacement costs of $250,000 or less during the studied time period, 2003 to 2008.

of our methodology for modeling flood risk and summary statistics on claim rates and loss distributions. Online Appendix A.1 offers a detailed explanation and additional modeling.

The probability π of flood loss l comprises two elements, the probability of incurring a flood loss $\pi(l > 0)$ and given a flood loss, the probability of a specific loss $\pi(l|l > 0)$ such that

$$\pi(l) = \pi(l > 0)\pi(l|l > 0).$$

We model the claim rate as an approximation of the probability that a policyholder incurs any flood loss. Unlike many other insurance products, flood insurance premiums in this program are not influenced by previous claims experience. Thus, while other forms of insurance create a disincentive to report small flood losses due to the potential that it will increase future premiums (Braun et al., 2006), this flood insurance motivates those who suffer flood losses to report even minor damage so that a professional adjuster can determine if damages exceed the deductible. For example, 2.5 percent of all flood insurance claims in our dataset are for losses less than the minimum deductible of $500.

We estimate the likelihood of having a flood claim using a random effects panel logit model with policies and claims data from 2001 to 2009. We estimate separate models for homes built before and after federal flood maps were developed. We use detailed information regarding the insured home and its vulnerability as explanatory variables. Examples include the number of floors in the home, the presence of obstructions, and an assessment of actions taken by the community to reduce flood risk, and the elevation of the home.[11]

Table 1 provides summary statistics; the estimated average annual flood claim probability is 1.33 percent. We also estimate loss distributions for each household in the program. Using all flood claims from 1982 to 2009, we model losses as a percent of the property value and estimate their probability at each percentile.[12] We find that losses are distributed log-normally and model the two parameters of this distribution, μ and σ, based on a households' observable characteristics using an iterative MLE approach following Aitkin (1987) and Western and Bloome (2009). Table 1 also provides summary statistics for the loss distribution. The expected damage is 20.3 percent of the home's value; however, the median damage is 9 percent.

[11] Home-level elevation information is available for approximately 40 percent of homes in the baseline sample. Our flood models include home-level elevation when it is available.

[12] As described in Section 3.1.1, damages are assessed by the flood insurance program in terms of property value.

Table 1. Policy Summary Statistics

	Mean	St. Dev.	Percentiles				
			1%	10%	50%	90%	99%
Claim Rate	1.33%	0.56%	0.29%	0.70%	1.33%	1.56%	3.24%
Loss Given a Claim	20.3%	26.0%	0.4%	1.5%	9%	59%	100%
Premium ($)	540	326	147	219	490	932	1675
Load	3.22	2.84	0.76	1.37	2.69	5.25	12.17
Property Value ($)	112,387	51,138	16,399	51,325	105,325	186,335	242,196
Replacement Cost ($)	107,177	57,779	13,300	41,000	96,300	195,000	250,000

Note: "Loss given a claim" reports the median loss as a percent of a home's property value for each percentile across policyholders. Load describes a contract's ratio of premiums to expected payouts. Replacement cost provides the amount needed to rebuild the current structure using similar materials. We derive these claim rate and loss estimates from our flood models described in Section 3.1.2. Baseline sample of 103,080 policies.

We compare this parametric approach, which allows for estimating the loss distribution of each household, to several nonparametric strategies in Online Appendix A.1.2. The distribution of potential property losses is influenced by the type of event that occurs; events that affect many households (e.g., a severe storm) result in larger expected losses for each household. We account for this process in our parametric models with fixed effects that group events by the number of claims associated with an event and, in our nonparametric models, by weighing observations using the number of claims for the event. Figure 1 illustrates, showing a parametric, weighted nonparametric, and unweighted nonparametric loss distribution for households in the baseline sample whose homes were built before flood maps were developed. The parametric distribution uses the median parameter values of μ and σ for these households, and tends to fall between the weighted and unweighted nonparametric distributions. We use the probability estimates from the parametric approach in our risk preferences models. Our estimates of risk are quite similar to those of Kousky and Michel-Kerjan (2015) who report on the flood insurance program using a longer time series, policies from 1980 to 2012.[13]

[13] They find an average annual claim rate of 1.45 percent whereas we estimate 1.33. Regarding loss distribution estimates, we include our corresponding estimate in brackets in the following excerpt from their paper: "half of claims [55 percent] are for less than 10 percent of the value of the building, roughly 15 percent [10 percent] of claims exceed 50 percent of the building's value, and approximately 7 percent [6 percent] exceed 75 percent of building value" (p.13).

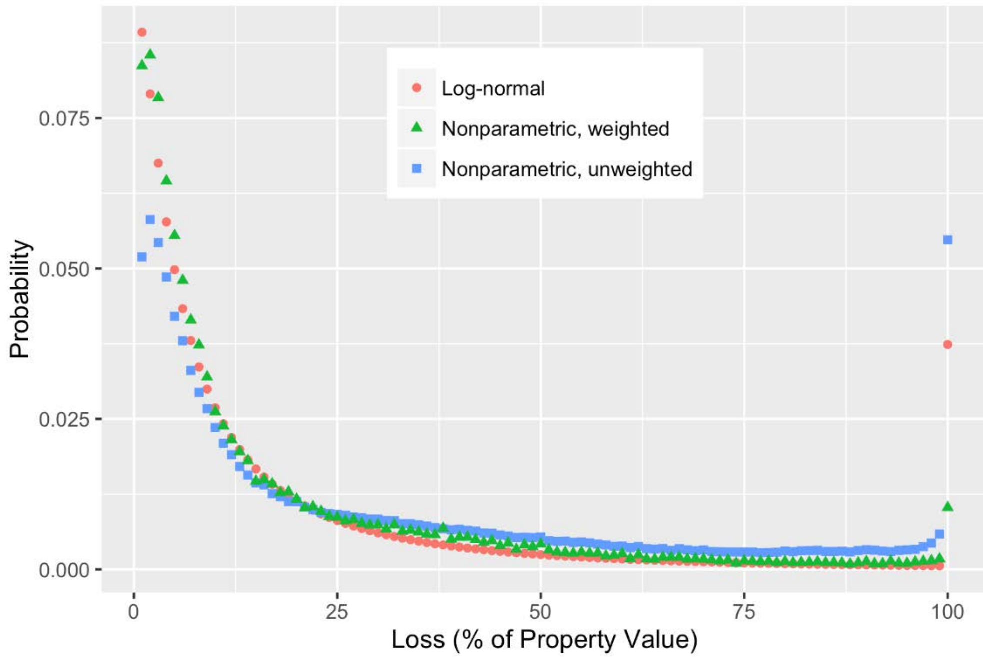

Figure 1 Distribution of Losses Given a Claim

Note: Loss distribution estimates for households in the baseline sample whose homes were built before flood maps were developed. The squares and triangles show flood loss distributions at the aggregate level in our data. The squares weight all observations equally. The triangles weight observations by the severity of the event. Events that affect many households (e.g., a severe storm) result in larger expected losses for each household. We estimate the parameters of the log-normal distribution for each household using the households' characteristics and fixed effects that account for event severity. The circles in the figure show the log-normal distribution using the median parameter estimates, $\mu = -2.49$ and $\sigma = 1.40$, for these households.

3.1.3 Insurance Premiums

The formula used by the flood insurance program that determines premiums p for home i in community k is

$$p_i = \left(\left(\min(c_i, 50{,}000)\, r_{1,i} + \max(c_i - 50{,}000, 0)\, r_{2,i}\right) \delta_i + a_{1,i}\right) CRS_k + a_{2,i}. \quad (1)$$

Households' choices affect premiums through the household's selected coverage limit c_i and a multiplier $\delta_i \in [0.75, 1.10]$ that depends on the selected deductible, with smaller multipliers for larger deductibles (e.g., a household choosing between a \$1,000 and \$5,000 deductible may have $\delta = 1.0$ versus $\delta = 0.785$, respectively). The characteristics of the home (e.g., its elevation relative to the 100-year flood plain) determine the coverage rate: $r_{1,i}$ is the rate per dollar of coverage limit on the first \$50,000 of coverage limit (mean = 0.0072, s.d. = 0.00355); $r_{2,i}$ is the rate for coverage limits in excess of \$50,000 (mean = 0.0030, s.d. = 0.00249). These parameters differ, among other reasons, depending on whether a home was built before federal flood maps were developed for its location. For example, premiums for homes built before the flood maps do not penalize low-lying properties to the extent that they do for homes built afterward. Parameter $a_{1,i}$ addresses the cost of bringing a flood-damaged home up to current building codes (mean = 44.75, s.d. = 29.59). Actions taken by the community to reduce flood vulnerability can reduce households' premiums through a FEMA program called the "Community Rating System," $CRS_k \in [0.65, 1]$.[14] Finally, $a_{2,i} \in \{30, 35\}$ is a policy administration fee. The flood insurance program provides tables to calculate a policyholder's parameters (e.g., r_1 and r_2) for the premium calculation (Equation 1; NFIP, 2006a). We calculated premiums manually from these formulas to confirm that all the policies in our data were priced using this structure.

Table 2 provides an illustrative pricing menu. For convenience, we have shown a subset of coverage limits. Each cell shows the median premium that a household would pay if it chose that deductible and coverage limit combination. For example, the median annual premium is $593 for a contract with a $1,000 deductible and $100,000 coverage limit.

Table 1 includes statistics on contract loads, the ratio of a contract's premium (determined by the NFIP) relative to its expected payout (determined by our flood models). Ninety-seven percent of policyholders in our baseline sample have a load greater than 1 (i.e., above the "actuarially fair" rate), suggesting that if they are expected utility maximizers, these households are risk averse. The median contract load in our baseline sample is 2.7.[15] Loads in our baseline sample tend to be higher than that for the program. We only examine a household's contract choices at the time of its first claim, excluding the policy afterward and so reducing the influence

[14] Community actions include maintaining and disseminating flood maps of the community, preventing building in floodplains, developing flood warning systems, improving community drainage systems, etc. (see FEMA, 2015).

[15] Several common misperceptions about the program (e.g., that the average participating household pays actuarially advantageous rates) motivate a brief discussion of its pricing. Like private insurance markets, the flood insurance program includes expense loads to cover program and insurance agent costs; however, the program differs from private insurance pricing in two significant ways. First, flood premiums do not include capital costs (e.g., reinsurance and reserving due to uncertainty about the underlying risk). Instead, the NFIP sets premiums so that, net of expense loads, collected premiums will cover losses in the "average historical loss year." In years when losses exceed the program's accrued earnings, the NFIP borrows from the U.S. Treasury. Second, premium ratings are less sensitive to risk than those a private insurer would likely charge, e.g., due to grandfathering for homes built before flood maps were developed. The program's rules have resulted in substantial cross-subsidization (discussed below and in Kousky and Shabman, 2014).

of properties with frequent losses on our load statistics.[16] Also, loads in our baseline sample are within the range that previous research identifies for sub-populations in the program. For example, Michel-Kerjan (2010) examines the ratio of NFIP premiums collected to claims paid by state over the period 1978 to 2008 (in 2008 dollars). The ratio varies substantially by state with a maximum of 15.1 in Colorado and minimum of 0.2 in Mississippi. Finally, loads in our baseline sample are within the range of some private catastrophe insurance markets. For example, Kunreuther et al. (1995) survey insurance underwriters regarding loads for a hypothetical catastrophe insurance contract. The geometric mean of their loads ranges from 1.6 to 3.4 depending on the level of uncertainty about the risk (the geometric mean for our baseline sample is 2.7). For uncertain liability risks, they find a geometric mean load over 5.5.

Table 2 Illustrative Pricing Menu

		Coverage Limit ($)						
		25,000	50,000	75,000	100,000	150,000	200,000	250,000
Deductible ($)	500	286	474	558	642	810	979	1,139
	1,000	269	440	516	593	746	899	1,045
	2,000	257	417	489	560	703	846	984
	3,000	249	400	468	536	671	806	937
	4,000	240	383	447	511	639	766	890
	5,000	232	366	426	486	606	726	843

Note: Menu provides median premium values for the households in the baseline sample for different deductible and coverage limit combinations. The table shows the six deductible options. Households may select a coverage limit of up to $250,000 with choices in $100 increments.

3.2 Households' Contract Choices

Households overwhelmingly select low deductibles for their flood insurance, as shown in Table 3. Ninety-four percent of households selected one of the two lowest deductible choices available. The flood insurance program has two default deductible options, depending on whether the insured home was built before federal flood maps were developed for its location. Those built before the maps typically have a default deductible of $1,000 while those built afterward have a

[16] The U.S. Government Accountability Office (GAO, 2006) estimates that one percent of the properties insured "account for 25 to 30 percent of all claims losses." Our baseline sample also differs in that it is limited to homes in Zone A and valued at $250,000 or less (the maximum allowable coverage limit in the program, Section 3.1.1). Both criteria may lead to different contract loads in the baseline sample than those for the entire program.

$500 default deductible. All households in our data are free to choose any of the six available options in the deductible menu, regardless of their default.[17]

Seventy-two percent of households with a $1,000 default keep this deductible, 19 percent select the $500, and about 9 percent choose a higher deductible. Eighty-two percent of households with a $500 default keep this deductible, 15 percent select the $1,000 deductible, and 3 percent choose a higher deductible. Previous research identifies that menu defaults influence consumers' choices (Thaler and Benartzi, 2004); however, the flood insurance program might judiciously select its menu defaults as part of a differentiated marketing strategy. We explore the effects of deductible defaults in two ways. First, in Section 4.1.3, we estimate households' risk aversion and include their menu default as an explanatory variable to assess its effect on our estimates. Second, discussed in Section 5, we re-estimate our main results using only households who selected a deductible that differed from their menu defaults.

Figure 2 shows the coverage limits selected as a proportion of the home's replacement cost separately for households selecting the $500, $1,000, and $5,000 deductibles. Ninety-eight percent of our sample chooses one of these deductible options. Three aspects of this figure are noteworthy. First, households most commonly select a coverage limit equal to their replacement cost; 42 percent of all policyholders do this.

Second, households who select a higher deductible are more likely to select lower coverage limits, with a Spearman rank correlation $\rho = -0.14$, $p < 0.01$.[18] This pattern could be explained by risk aversion: when premiums are above actuarially fair rates, less risk averse households purchase *higher* deductibles and *lower* coverage limits than more risk averse households.

Third, about a third of households select a coverage limit that is more than the replacement cost, which is surprising as they cannot receive a payment greater than this amount. Mossin (1968) predicts this behavior: "there may be some uncertainty as to what will be the actual evaluation of the property in case of damage. We have not taken this kind of randomness into account in our theory, but it seems a perfectly reasonable hypothesis that it will lead to a number

[17] In practice, $\delta = 1$ for the default deductible ($1,000 for homes built before flood maps, and $500 for homes built afterward) in the premium calculation (Equation 1) and so serves as the reference for the pricing tables.

[18] This correlation across choices is similar in magnitude to that found by Einav et al. (2012) who examine individuals' decisions for several types of insurance coverage and 401(k) investments.

of observed cases of full (or over-) insurance" (p. 558).[19] Consequently, we conclude that households who insure at or above the replacement cost of their home intend to fully insure. In Online Appendix A.3.7, we exclude households that over-insure and retest our models, finding that their exclusion does not qualitatively change our main findings.

Table 3. Policyholder Deductible Selection

Deductible	$500	$1,000	$2,000	$3,000	$4,000	$5,000
Policyholders (%)	46.6	47.4	1.4	0.5	0.2	3.8

Note: Baseline sample of 103,080 policies.

[19] In our data, households whose property values are higher than their replacement cost are almost twice as likely to over-insure. Forty three percent of households who did not over-insure have a property value that is greater than the replacement cost while 72 percent of over-insurers do. These households may approach the insurance decision with an estimate of their property value, receive from their insurance agent a lower replacement cost, and buy a coverage limit above the replacement cost to make certain that their coverage is sufficient.

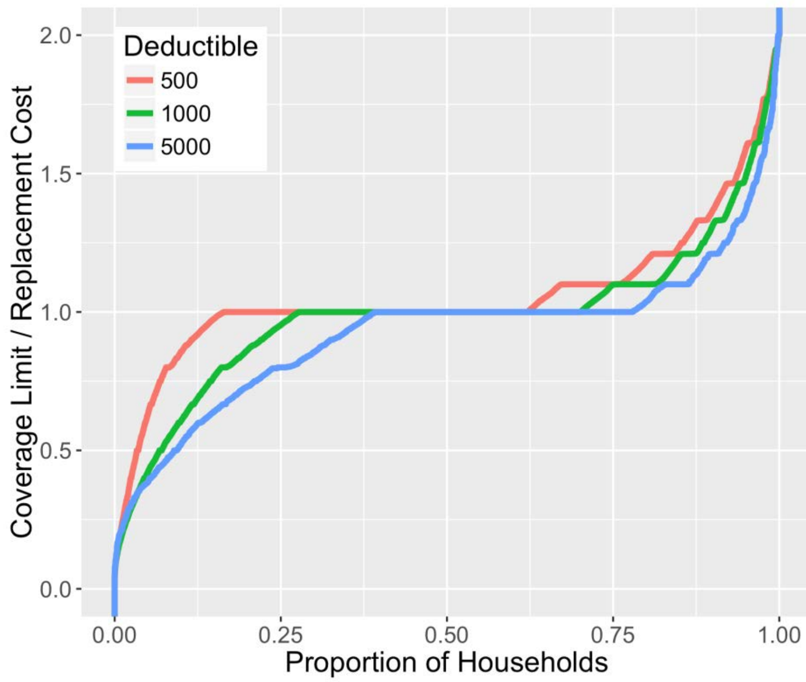

Figure 2 Household Coverage Limit Selections

Note: The figure provides household coverage limit selections as a proportion of the replacement costs of their homes separately for households selecting the $500, $1,000, and $5,000 deductibles. Ninety-eight percent of our sample chooses one of these deductible options. Households cannot receive an insurance payout that exceeds their replacement cost. Baseline sample of 103,080 policies.

Table 4 compares several characteristics of insured homes across the most common deductible choices and divides coverage limits into households that select coverage limits less than their replacement cost ("Partial"), at the replacement cost ("Full"), and above the replacement cost ("Over"). For example, it shows that households choosing the $5,000 deductible tend to have higher claim rates, higher premiums, higher contract loads, homes at lower elevation, older homes, and homes that were built before flood maps were developed. Regarding coverage limits, the table shows that partially insuring households tend to have homes at lower elevation, with lower property values, and that were built before flood maps were developed.

Table 4 Contract Choices and Descriptive Statistics

		Deductible			Coverage Limit		
		$500	$1,000	$5,000	Partial	Full	Over
Claim rate	Median	0.011	0.013	0.014	0.014	0.012	0.013
	Mean	0.012	0.014	0.015	0.015	0.013	0.013
	S.D.	0.006	0.005	0.005	0.0052	0.0057	0.0056
Premium ($)	Median	372	558	611	470	492	503
	Mean	431	625	696	504	543	560
	S.D.	262	334	391	286	329	344
Load	Median	2.43	2.87	2.96	2.46	2.64	2.97
	Mean	3.03	3.35	3.64	2.78	3.07	3.67
	S.D.	2.93	2.55	3.67	1.84	2.25	3.79
Property Value ($)	Median	111,600	99,170	109,400	91,730	106,700	113,700
	Mean	118,000	106,100	115,800	98,280	113,300	120,300
	S.D.	51,559	49,976	50,091	45,593	52,509	50,961
Property value / Replacement cost	Median	1.09	1.06	0.89	0.95	1.03	1.25
	Mean	1.25	1.31	1.02	1.07	1.20	1.48
	S.D.	0.82	1.01	0.69	0.62	0.85	1.10
Age of home (years)	Median	29	40	48	41	34	34
	Mean	30	44	50	45	36	36
	S.D.	18.7	19.5	24.2	22.6	19.8	19.4
Elevation (feet)	Median	1	1	0	1	1	1
	Mean	1.4	1.1	0.6	0.9	1.3	1.5
	S.D.	5.5	3.0	3.1	5.4	3.2	6.5
Default deductible							
$500		82%	15%	1%	13%	44%	43%
$1000		19%	72%	6%	29%	41%	30%

Note: Characteristics of the insured home and households' deductible and coverage limit selections. For the deductible, we show the three most commonly selected deductibles. For coverage limits, we divide households into those that select coverage limits less than their replacement cost ("Partial"), at the replacement cost ("Full"), and above the replacement cost ("Over"). Load describes a contract's ratio of premiums to expected payouts. Elevation is relative to the 100-year flood plain. Baseline sample of 103,080 policies.

3.3 Modeling Households' Preferences

Our analyses of risk preferences begin with expected utility models but expand to include rank dependent probability distortions. Here, we present the more general approach with a probability weighting function, which nests the case in which probabilities are treated objectively (i.e., they are not distorted). We assume that households' deductible and coverage limit decisions are characterized by the weighted utility function

$$\int u(c, d, l, p, w; \rho)\omega(\pi(l); \boldsymbol{\beta}) \quad dl \qquad (2)$$

where u represents the household utility function; c an insurance coverage limit; d deductible; p premiums; w wealth; l flood losses, which occur with probability π and are weighted following function ω. The focus of our analysis is estimating risk aversion coefficient ρ and, for rank dependent utility models (discussed further in Section 4.2), probability distortion parameters $\boldsymbol{\beta}$ using households' deductible and coverage limit decisions.

As we describe in Section 3.1.1, our data include households' deductible and coverage limits choices and the pricing schedule that allows us to determine the premium for any combination of deductible and coverage limit. We estimate flood loss probabilities from the program's claims database. In Equation 2, we apply these probability distributions to a households' replacement cost (i.e., l describes flood losses in terms of a household's cost to address them).[20] Finally, we do not directly observe a household's wealth, but assume that it equals the home's replacement cost. We test our results with other wealth assumptions including using the property value as wealth (Section 5).

3.3.1 Insurance Decisions

We allow for households to treat their deductible and coverage limit choices as separate decisions, but assume that households observe their coverage limit decision when selecting a deductible, and vice versa. This approach is a type of marginal analysis and allows for risk attitudes to differ across deductible and coverage limit decisions; however, if households' preferences are similar for these decisions, it will lead to similar parameter values for each choice. The household deductible selection problem is thus

$$\max_{d \in \{d_1, d_2, .., d_n\}} \int_0^d u(w - p(d) - l)\omega(l) + \int_d^{c^*} u(w - p(d) - d)\omega(l) + \int_{c^*}^{\bar{c}} u(w - p(d) - d - l + c^*)\omega(l) \quad dl \quad (3)$$

where d is the deductible, c^* the coverage limit the household selects, $\bar{c}$ is the total replacement cost, p the premium, l losses with $l \in [0, \bar{c}]$, and $\omega(l)$ the household's transformation of loss probabilities. The first integrand accounts for flood losses less than the deductible, the second

[20] The implicit assumption we make is that if a home is damaged by flood, the household will replace what was damaged. An alternative would be to measure losses in terms of the property value instead of the replacement cost. This modeling choice regarding losses seems to align with households' decisions as households most commonly select a coverage limit that equals their replacement cost (Section 3.2).

accounts for losses above the deductible but less than the coverage limit, and the third accounts for losses greater than the coverage limit selected by the household. Similarly, the households' coverage limit decision is

$$\max_{c\in[\underline{c},\bar{c}]} \int_0^{d^*} u(w - p(c) - l)\omega(l) + \int_{d^*}^{c} u(w - p(c) - d^*)\omega(l) + \int_{c}^{\bar{c}} u(w - p(c) - d^* - l + c)\omega(l) \quad dl \quad (4)$$

where $d^*$is the deductible that the household selects. In Section 4.1.2, we also examine the joint optimization problem in which households select among all possible combinations of the deductible and coverage limit.

$$\max_{d\in\{d_1,d_2,\dots,d_n\},c\in[\underline{c},\bar{c}]} \int_0^{d} u(w - p(c) - l)\omega(l) + \int_{d}^{c} u(w - p(c) - d)\omega(l) + \int_{c}^{\bar{c}} u(w - p(c) - d - l + c)\omega(l) \quad dl \quad (5)$$

In the random utility model (Section 3.3.2), which estimates the model parameters that explain households' decisions, this problem results in a single risk aversion coefficient, accounting for both contract choices.

3.3.2 Estimation

Our estimation approach builds on those pioneered by Camerer and Ho (1994) and Hey and Orme (1994) who adopt a random utility framework (McFadden, 1974) and fit value function parameters using maximum likelihood estimation.[21] For household i, let

$$v_i(k; \boldsymbol{\theta}) \equiv u_i(k; \boldsymbol{\theta}) + \epsilon_{ik}$$

where u represents the household probability-weighted value function, $k \in K$ a specific contract among a set of insurance contracts, θ vector of value function parameters such as relative risk aversion and probability weights for rank dependent utility models, and ϵ_{ik} an i.i.d. error

[21] Barseghyan et al. (2013) provide a set of proofs showing that this framework allows for disentangling risk preferences and probability distortions from insurance purchase decisions (p. 2,510). The intuition behind this identification strategy is that households' selections across a menu of insurance contract options (e.g., more than two deductible options) allows for differentiation between curvature in the value function and probability distortions if the data include both exogenous variation in the risk of loss and premiums. Our data include this variation (Section 3.1).

component, which is assumed to be distributed type 1 extreme value with scale parameter σ. The probability that a household chooses contract k is thus

$$p_{ik} = \frac{\exp\left(v_i(k;\boldsymbol{\theta})/\sigma\right)}{\sum_{k' \in K} \exp\left(v_i(k';\boldsymbol{\theta})/\sigma\right)}. \qquad (6)$$

Our estimation strategy solves the problem

$$\underset{\boldsymbol{\theta},\sigma}{\operatorname{argmax}}\, \mathcal{L} = \sum_{i=1}^{N} \sum_{k=1}^{K} y_{ik} \ln p_{ik}$$

where $\mathcal{L}$ is the log-likelihood function and $y_{ik} = 1$ if household i chooses contract k and 0 otherwise (Cameron and Trivedi, 2005). We model households' deductible choices using the six available options: \$500, \$1,000, \$2,000…\$5,000. We model households' coverage limits from \$10,000 to \$250,000 in \$10,000 increments, rounding households' choices and assumed wealth to the nearest \$10,000.[22]

4 Results

We examine households' small and large stakes decisions in the context of expected utility and rank dependent utility models. We use CRRA utility because of its prevalence (Wakker, 2008) and previous research indicating that risk preferences are approximately proportional in wealth (e.g., Friend and Blume, 1975; Szpiro, 1986), but we also examine CARA and expo-power utility functions (discussed in Section 5). Our primary analyses consider what coefficient of relative risk aversion best characterizes a decision in our data using MLE in the random utility model described in Section 3.3.2 and how well each model explains households' choices.

We report several measures to compare models, the log likelihood, the Akaike information criterion (AIC), and Bayesian information criterion (BIC). The latter measures penalize models with more parameters.[23] We also use Vuong (1989) tests for comparing non-nested models and hypothesis testing.

---

[22] We alternatively examined households' coverage limits in increments of \$1,000 in our baseline sample and derived similar risk preferences to those estimated using the \$10,000 increments. We implement the maximum likelihood estimation in R using differential evolution algorithms for the optimization (Ardia et al., 2016; Price, Storn, and Lampinen, 2006).

[23] Comparisons of model predictions and actual choices (e.g., the percent correctly predicted) are not a suitable goodness-of-fit statistic in this context as random utility models are probabalistic (see Train, 2009, p. 69). As an example, suppose that Model 2 is nested in Model 1 and fit the data better (e.g., it has a lower BIC). These

4.1 Expected Utility Theory

4.1.1 CRRA, Canonical Model

We model households as expected utility maximizers with CRRA utility

$$u(x) = \begin{cases} \frac{1}{1-\rho} x^{1-\rho} & \rho \neq 1 \\ \ln(x) & \rho = 1 \end{cases} \quad (7)$$

$$s.t., x > 0$$

where ρ is the Arrow-Pratt coefficient of relative risk aversion. We assume that households' initial wealth is their home's replacement cost, and impose a minimum such that final wealth never falls below $1. We consider alternative wealth assumptions in Section 5. We begin with a coefficient of relative risk aversion $\rho = 1$ as a canonical value from previous research (Arrow, 1971; Chetty, 2006; Gourinchas and Parker, 2002) and examine different values in the following subsections.

The canonical model, CRRA expected utility with $\rho = 1$, predicts that households would insure much less of their risk than they do. Figure 3 provides households' deductible and coverage limit selections and those predicted by our models. "Actual" is the distribution of households' choices and "EUT CRRA, $\rho = 1$" describes the predictions of the standard model. For the deductible, the model predicts that 92 percent of households would select the highest deductible ($5,000) while more than 90 percent select one of the lowest two deductible options ($500 and $1,000). In total, the model correctly predicts the deductibles selected by 8 percent of households. For the coverage limit, the model predicts that 85 percent of households would select a coverage limit of less than half of their replacement cost while only about 5 percent do.[24]

---

models may predict the same choices. For example, suppose that household i selects the $500 deductible, which the Model 1 estimates has a 5 percent probability and Model 2 estimates has a 10 percent probability. Model 2 would explain this observation better despite the conclusion of both models that the household is unlikely to select this deductible.

[24] The low predicted coverage limits result from the contract loads. Premiums are greater than the expected loss for almost all households (Section 3.1.3). A risk averse household will tend to prefer increasingly larger deductibles as contract loads increase, reducing coverage for moderate stakes but maintaining it for large stakes. However, given the available contract options, the only way that households can reduce their insurance after selecting the $5,000 deductible is by selecting a lower coverage limit, leading to the coverage limit predictions of the canonical model.

In the next subsections, we examine households' decisions using MLE and several utility model specifications. Here, we consider the potential welfare implications of households' deviations from the canonical model. Chetty (2015) contrasts an individual's "decision utility" (factors guiding households' choices) and "experienced utility" (the hedonic realization of its choices), concepts originally introduced by Kahneman, Wakker and Sarin (1997). He posits that households' experienced utility may follow the canonical model even when their decision utility does not.

Table 5 Model Results, Maximum Likelihood Estimation

		ρ	β_0	β_1	σ	Log Likelihood	AIC	BIC
EUT CRRA	Deductible	139			0.377	-168,942	168,946	337,907
		(0.85)			(0.0031)			
	Coverage Limit	2.72			0.191	-299,923	299,927	599,869
		(0.0068)			(0.0012)			
EUT CRRA, heterogeneous	Deductible							
	Mean	113			0.310	-164,935	164,973	330,090
	Median	108			(0.002)			
	Coverage Limit							
	Mean	2.81			0.184	-297.756	297,794	595,730
	Median	2.78			(0.001)			
RDU Yaari	Deductible		0.175	5.80	0.023	-104,378	104,384	208,791
			(0.0007)	(0.041)	(0.00032)			
	Coverage Limit		0.120	1.49	0.00711	-245,684	245,690	491,403
			(0.00048)	(0.024)	(0.000031)			
RDU CRRA	Deductible	-3.05	0.0378	4.77	0.181	-103,880	103,888	207,806
		(0.04)	(0.00031)	(0.022)	(0.00072)			
	Coverage Limit	-0.767	0.00774	3.64	0.125	-244,469	244,477	488,984
		(0.0088)	(0.000034)	(0.027)	(0.00049)			

Note: Standard errors are in parentheses. Table compares utility models describing households' deductible and coverage limit decisions, comparing outcomes using the Akaike information criterion (AIC) and the Bayesian information criterion (BIC). For the AIC and BIC, lower values indicate better fit. The models use maximum likelihood estimation to fit model parameters. The parameter σ describes the scale of model errors, following Equation 6. Parameters $\boldsymbol{\beta}$ show probability distortions $\Omega = \beta_0 + \beta_1 \Pi$ where Π is the cumulative objective probability. The EUT CRRA heterogeneous model estimates risk aversion as a function of the household's characteristics. Yaari (1987) proposes a model in which individuals are risk neutral but distort probabilities. We allow for probability distortions using rank dependent utility (RDU, Quiggin, 1982). Expected utility theory (EUT) assumes that households weight outcomes based on their objective probabilities and so these models fix $\beta_0 = 0$ and $\beta_1 = 1$. Positive values for β_0 and $\beta_1 > 1$ indicate overweighting of the small probabilities in our data. Negative values for the coefficient of relative risk aversion ρ indicate diminishing sensitivity to losses. In the table, if a parameter is missing it is assumed to take the following value, $\rho = 0$, $\beta_0 = 0$, $\beta_1 = 1$. Baseline sample of 103,080 policies.

If so, then households' deviations from the canonical model have welfare implications. Let households' experienced utility be CRRA expected utility with $\rho = 1$, k_i represent the contract choice of household i, and k_i^* the optimum in the canonical model. We identify the compensation a_i required to equate the experienced utility of the household's choice with that of the optimal choice, $Eu(k_i + a_i) = Eu(k_i^*)$. This calculation provides a dollar value for the expected utility loss of deviating from the canonical model. Average compensation a for the deductible model is $62, which is 13 percent on average of a household's optimal premium. Average compensation a is $189 for the coverage limit, which is 95 percent on average of the optimal. The reported percent calculation is $\text{mean}(a_i/p_i^*) = 0.95$ where p^* is the optimal premium. The expected utility cost of households' deviations from the optimal coverage limit is almost as large as the total premium for that contract. Also, the expected utility losses from households' coverage limit choices are three times larger than those of the deductible.

Panel A. Deductible Predictions *Panel B. Coverage Limit Predictions*

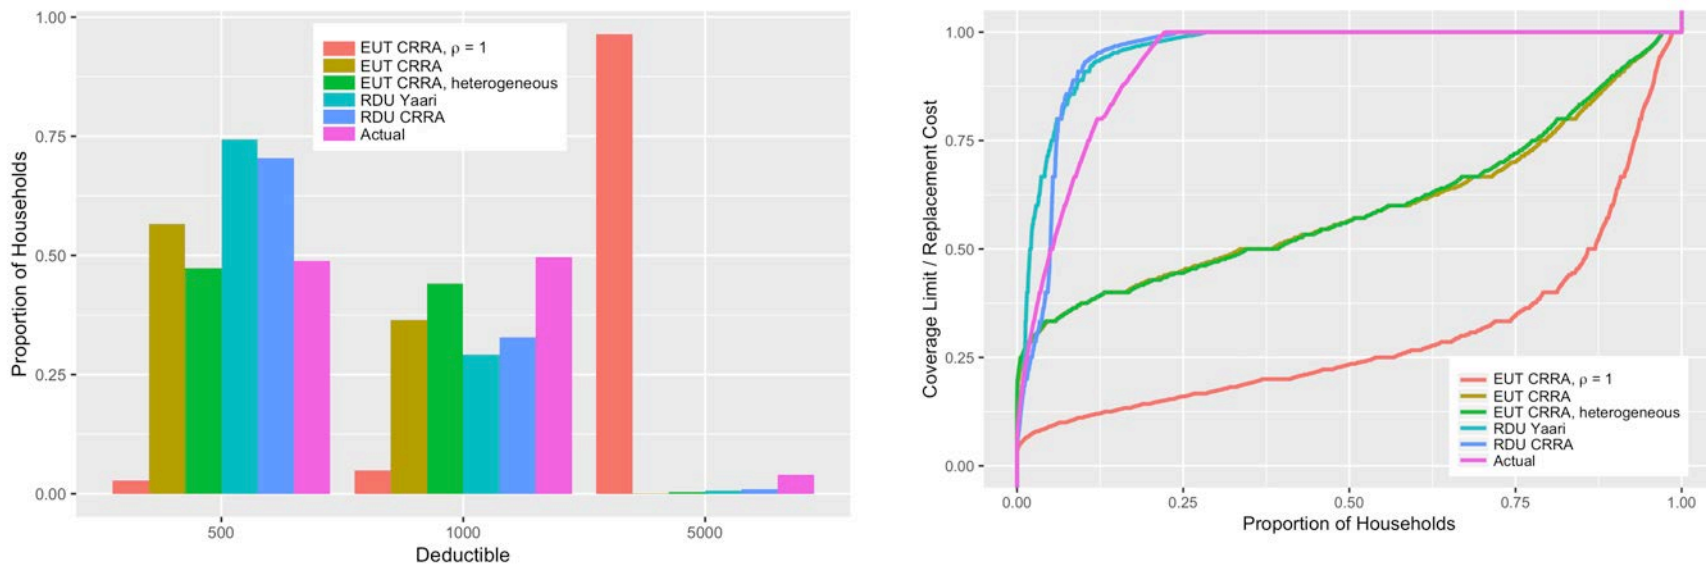

Figure 3 Model Predictions

Note: The graphs compare households' contract decisions ("Actual") with the predictions of several utility models. Panel A shows deductible predictions. For ease of illustration, it only shows the three most frequently chosen deductibles ($500, $1,000, $5,000), which are selected by 98 percent of households. Panel B shows coverage limit ratios, the proportion of the coverage limit to the home's replacement cost, imposing a maximum ratio value of 1. "EUT CRRA, ρ = 1" uses a relative risk aversion coefficient commonly derived in previous research. The remaining models use maximum likelihood estimation to fit model parameters. The parameter values are shown in Table 5. "EUT CRRA" models households as CRRA expected utility maximizers. "EUT CRRA, heterogeneous" models estimates a household risk aversion as a function of the household's characteristics. "RDU Yaari" models households as risk neutral, rank dependent utility maximizers; and "RDU CRRA" as CRRA rank dependent utility maximizers. Baseline sample of 103,080 policies.

4.1.2 CRRA, Maximum Likelihood Estimation

Next, we use MLE to identify the risk aversion coefficients that characterize households' preferences for deductibles and coverage limits, labeled "EUT CRRA" in the figures and Table 5. The EUT CRRA deductible model correctly describes 47 percent of households' deductible selections. It correctly describes the selections of 60 percent of households choosing the $500 deductible, 38 percent of those choosing the $1,000 deductible, but only 0.3 percent of those choosing the $5,000 deductible.

The coefficient of relative risk aversion for the EUT CRRA deductible model is $\rho = 139$. Like Sydnor (2010), we find three-digit relative risk aversion from households' insurance deductible choice. Our risk aversion estimate from the EUT CRRA coverage limit model is much lower, $\rho = 2.7$, and more consistent with estimates found in other domains. For example, Barro and Jin (2011) estimate $\rho = 3$ in their analyses of equity premiums. Unfortunately, the descriptive power of the EUT CRRA coverage limit model is quite poor. For example, the model indicates that only 10 percent of households would insure at least 90 percent of their replacement cost, but 83 percent do. On average, this model under-predicts households' coverage limit ratios by 34 percentage points ($\text{mean}((c_i - \widehat{c_i})/rc_i)$ = 0.34 where c_i and $\widehat{c_i}$ are the selected and predicted coverage limits and rc_i is the replacement cost for household i).

We also estimate risk aversion using a model in which households select a deductible and coverage limit jointly, following Equation 5. While in the above we conduct a marginal analysis on which we hold the coverage limit decision constant and model the households' deductible choice and vice versa, the joint optimization allows for all possible combinations of the deductible and coverage limit and considers what single relative risk aversion coefficient best explains the combined decision. We find a relative risk aversion estimate of 2.9. This suggests that the coverage limit decision largely determines the risk aversion estimates from the joint problem.

In sum, while MLE improves the model fit and descriptive ability of our EUT CRRA model relative to the canonical model, the approach highlights two limitations of the model. First, the deductible and coverage limit decisions are characterized by dramatically different preferences for the population of households in our data. Not only is the high risk aversion of the deductible choice difficult to reconcile with other domains, but also with the coverage limit decision, which occurs in the same domain and in the same insurance contract. Second, even allowing for different risk aversion parameters for low and high stakes decisions, the descriptive ability of these MLE models remains quite low, especially for the coverage limit, suggesting that these models of CRRA expected utility preferences do not characterize households' decisions well in our data.

4.1.3 CRRA with Heterogeneous Risk Aversion

Aggregate estimates of relative risk aversion might mask important heterogeneity in households that explains their contract selections. Our EUT CRRA coverage limit risk aversion estimate of $\rho = 2.7$, discussed in Section 4.1.2, must balance the decisions of both full and partial insurers: larger coefficients of relative risk aversion would better explain the behavior of full insurers, but also lead to greater losses of utility for partial insurers.

To examine this possibility, we estimate risk aversion using characteristics of the household and the home that have been shown to influence individuals' decisions under risk, such as recent severe events (Gallagher, 2014), culture or region (Harrison, Humphrey, and Verschoor, 2010), proxies for income and wealth (Cohen and Einav, 2007), and menu default options (Johnson et al., 1993). We model risk aversion as a function of the home's age and property value (as a proxy for income), whether households have a \$500 or \$1,000 deductible default, and region (using the nine U.S. Census Bureau divisions) and year fixed effects.

The EUT CRRA heterogeneous preferences models fit the data better than the EUT CRRA model with homogeneous preferences according to the AIC and BIC (Table 5). Table 6 reports summary statistics for the estimated distribution of relative risk aversion for the deductible and coverage limit. Online Appendix A.2 provides additional details about the model results.

The average relative risk aversion is 113 for the deductible model and 2.81 for the coverage limit model. In each case, the mean risk aversion of these models assuming heterogeneous preferences is within one standard deviation of the estimates from the models assuming homogeneous preferences. We find that risk aversion is slightly right-skewed for both the deductible and coverage limit estimates.

The deductible model with heterogeneous preferences does not improve the accuracy on households' actual choices compared to the model assuming homogeneous preferences. The heterogeneous preference model predicts 47 percent of households' deductible decisions, and underestimates their coverage limit ratios by 34 percentage points.

While additional explanatory variables might further improve model fit, standard expected utility of wealth models face a fundamental problem in explaining households' coverage limit choices. Except under extreme risk aversion, expected utility maximizing households should partially insure when insurance is priced above the actuarially fair rate (Mossin, 1968). Yet, in our data, most households fully insure. CRRA expected utility requires that households have at least a relative risk aversion of 221 to reach the rate of fully insuring that we observe. A model with $\rho = 221$ does not correctly sort households into full and partial insurers: it correctly predicts whether households fully or partially insure about 66 percent of the time overall.[25]

In sum, standard expected utility theory models with reasonable (i.e., single digit) relative risk aversion predict that households would choose high deductibles and partially insure, but households tend to do neither. Even with triple-digit values of relative risk aversion, our CRRA expected utility models predict a little under half of households' deductible decisions and identify whether households fully versus partially insure about two-thirds of the time.

Table 6 Distributions of Relative Risk Aversion for CRRA Expected Utility Models Assuming Heterogeneous Preferences

			Percentiles				
	Mean	St. Dev.	1%	10%	50%	90%	99%
Deductible	113	44.5	27.5	61.3	108	176	230
Coverage Limit	2.81	0.42	1.99	2.28	2.78	3.38	3.86

Note: Table provides the distribution of relative risk aversion for households' deductible and coverage limit decisions from a model in which households are assumed to be CRRA expected utility maximizers and relative risk aversion is a function of the household's and home's characteristics.

4.2 Rank Dependent Utility

Households' low and high stakes insurance decisions may be characterized by probability distortions. We allow for distortions on the cumulative distribution following Quiggin (1982) by (1) ordering outcomes and their associated probabilities from lowest to highest utility, (2)

[25] We could instead consider households that insure at least 90 percent of their home's replacement cost as "effectively fully insuring." This benchmark might allow our expected utility models to explain households' coverage limit choices without requiring 100 percent coverage of the replacement cost. Eighty-three percent of households insure at least 90 percent of their replacement cost. CRRA expected utility requires that households have at least a relative risk aversion of 21 to reach this rate. Also, 39 percent of households choose a "corner solution," selecting the lowest deductible and a coverage limit to cover their full replacement cost. For these households, we can only identify a lower bound of relative risk aversion: CRRA expected utility requires a relative risk aversion of at least 136 to explain the rate at which households select the $500 deductible and fully insure. A CRRA expected utility model with $\rho = 136$ correctly predicts these corner-solution households 25 percent of the time.

cumulatively adding these probabilities to form the empirical cumulative function, (3) applying distortions on this cumulative with the constraints that the transformed probabilities are restricted to be between zero and one and sum to one. Thus, a net overweighting of flood probabilities results in an underweighting of the outcome in which no flood occurs and vice versa.

While a variety of probability distortion models have been proposed (e.g., Tversky and Kahneman, 1992; Prelec, 1998), the probabilities in our data are small relative to the values typically tested in previous studies. Following Barseghyan et al. (2013), we take a more flexible approach, assessing for probability distortions of households' flood loss probabilities with a polynomial expansion. The models that we discuss below suggest that distortions are approximately linear in the range of probabilities that we examine. The rank-ordered cumulative flood loss probabilities Π_i of household i follow distortion $\Omega_{\mathrm{i}} = \beta_0 + \beta_1 \Pi_{\mathrm{i}}$.[26] If households in these models did not distort probabilities, we would find $\beta_0 = 0$ and $\beta_1 = 1$, providing support for expected utility theory rather than rank dependent utility.

4.2.1 Risk-Neutral, Rank Dependent Utility (Yaari Model)

We begin with a model that assumes a risk neutral, rank dependent utility maximizer who distorts probabilities (based on Yaari, 1987). A Vuong (1989) test indicates that the Yaari models provide a better fit for both the deductible and coverage limit decisions at the 1 percent significance level than the CRRA expected utility models discussed in Section 4.1.2. The Yaari model correctly predicts 61 percent of households' deductible choices. The Yaari model modestly over-predicts coverage limits, estimating that 90 percent of households would choose a coverage limit of at least 90 percent of their replacement cost while 83 percent do. Households tend to overweight small probabilities, as values for $\beta_0 > 0$ and $\beta_1 > 1$ in Table 5. We delay a discussion of the specific probability distortion parameters until the CRRA rank dependent utility model (Section 4.2.2).

[26] We examined up to a cubic polynomial expansion of probability distortions using Chebyshev polynomials and selected the linear distortion model based on the models' AIC and BIC values.

4.2.2 CRRA value function, Rank Dependent Utility

We now fit a rank dependent utility model with a CRRA value function ("RDU CRRA"). This model estimates both probability distortions and the coefficient of relative risk aversion. The CRRA rank dependent utility models provide modest improvements over the Yaari models in terms of AIC and BIC, and models the coverage limit and deductible decisions better at the 1 percent level using Vuong tests.

The CRRA rank dependent utility model correctly identifies 63 percent of households' deductible choices (versus 61 percent for Yaari), and over-predicts households' coverage limits by 2 percentage points (versus 3 for Yaari); CRRA expected utility under-predicts these choices by 34 percentage points).

For both deductible and coverage limit models, we find overweighting of small probabilities and negative values of ρ. These preferences exhibit diminishing sensitivity to losses, typically described as "risk seeking" in utility theory: households perceive the difference between a \$1,000 loss and \$2,000 loss as greater than that between a \$2,000 and \$3,000 loss.[27] Based on Yaari (1987) and Dyer and Sarin (1982), "risk seeking" is perhaps misleading here as the households that we analyze purchased an insurance product that reduces their risk of financial loss. In this context, households' demand for insurance is explained by their overweighting of small probabilities. Thus, our findings are consistent with two of the tenants of prospect theory: households 1) overweight small probabilities, and 2) demonstrate diminishing sensitivity to losses (Kahneman and Tversky, 1979).

The CRRA rank dependent utility results continue to suggest that households' small and large stakes decisions are characterized by different preferences. Regarding the utility function, ρ from their deductible selection is 4 times larger than that from their coverage limit selection. Across the spectrum of small to large stakes, these preferences cannot be universally convex in losses as any line connecting the two must be at some point concave. Figure 4 illustrates this

[27] Decreasing marginal utility of wealth is equivalent to increasing sensitivity to losses for typical expected utility models. Let the value x in the CRRA model in Equation 7 represent some initial wealth w_0 minus losses, $x = w_0 - l$. For $\rho \neq 1$, the first and second derivative with respect to l are

$$\frac{\partial u}{\partial l} = -(w_0 - l)^{-\rho}$$

$$\frac{\partial^2 u}{\partial l^2} = -\rho(w_0 - l)^{-(1+\rho)}.$$

The first derivative is negative: utility decreases in losses. The second derivative is positive if $\rho < 0$ and negative if $\rho > 0$.

point. The solid green line provides the shape of the value function derived from the deductible decision in the range of potential losses related to the deductible (losses up to $5,000). Similarly, the solid blue line provides the shape of the value function derived from the coverage limit decision in the range of potential losses related to selecting a low coverage limit (set at losses greater than 20 percent of the property value for this example). The dotted lines help illustrate the magnitude of this difference. They show how these preferences would extend toward more severe losses for the deductible and less severe ones for the coverage limit. A similar result applies to each of the expected utility models that we test as our low stakes risk aversion estimates are always larger in magnitude than those for larger stakes.

Regarding probability distortions, households overweight probabilities more when choosing a deductible than they do when selecting a coverage limit. For example, their estimates of the likelihood of a flood are 80 percent larger for the deductible decision as they are for the coverage limit decision. When making coverage limit decisions, households overweight the median annual claim rate of 1.33 percent by a factor of 4.2, acting as if it occurs with a 5.6 percent annual probability. When making deductible decisions, they overweight by a factor of 7.6, acting as if the median annual claim rate occurs with a 10 percent annual probability.[28]

[28] For example, for the deductible model and the median claim rate, $\Omega = \beta_0 + \beta_1 \Pi = 0.0378 + 4.77 \times 0.0133 = 0.101$. The ratio of the transformed probability to the objective probability is $\Omega / \Pi = 0.101/0.0133 = 7.59$.

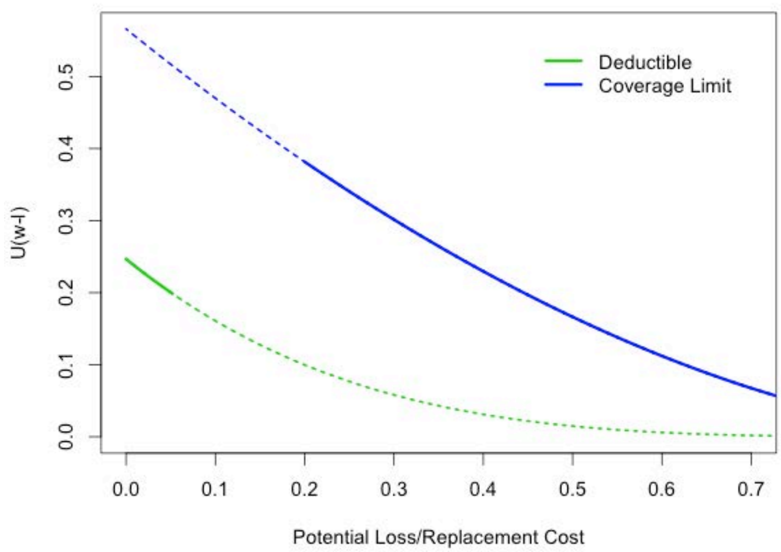

Figure 4 Risk Preferences from Deductible and Coverage Limit Selections

Note: Figure shows the CRRA rank dependent utility model with the MLE of model parameters fitted to deductible and coverage limit selections. The solid green line provides the shape of the value function derived from the deductible selection in the range of potential losses related to the deductible (losses up to $5,000). Similarly, the solid blue line provides the shape of the value function derived from the coverage limit selection in the range of potential losses related to selecting a low coverage limit (set at losses greater than 20 percent of the property value for this example). The dotted lines show how these preferences would extend toward more severe losses for the deductible and less severe ones for the coverage limit.

5 Robustness Tests and Alternative Explanations

In this section, we consider aspects of the setting and our modeling approach that might interfere with estimating households' preferences from their observed deductible and coverage limits. We discuss each here and provide additional details in the online appendix.

Other utility functions. Households' small and large stakes decisions may be characterized by expected utility but rely on some utility function other than CRRA. We consider CARA and expo-power preferences. The expo-power utility function (Saha, 1993) can accommodate a variety of combinations of relative and absolute risk preferences (e.g., decreasing absolute risk aversion, decreasing relative risk aversion, etc.). Holt and Laury (2002) posit that expo-power utility might reconcile households' small and large stakes preferences and find supporting evidence in their lab experiments. While CARA is nested in expo-power utility, we include a specific examination of it because of its prevalence.

The limitations of CRRA expected utility also apply to these models. CARA and expo-power predict that households would choose much smaller coverage limits than what they do.

For example, CARA and expo-power each underestimate households' coverage limit ratios choices by an average of 27 percentage points (versus 34 for CRRA). They also result in different model parameters for households' deductible and coverage limit decisions. For CARA, absolute risk aversion is 17 times larger for the deductible decision than the coverage limit. While these tests of expo-power and CARA do not exhaust the many possible utility functions that might be employed in expected utility models, they further illustrate the difficulty of explaining households' small and large stakes insurance preferences with common applications of expected utility theory (Online Appendix A.3.1).

Consumption commitments. Households' decisions may reflect budget or liquidity constraints. For example, Chetty and Szeidl's (2007) theory of consumption commitments posits that large shocks lead to adjustments in fixed liabilities (e.g., housing costs); however, such adjustments can be costly so small shocks are fully born by discretionary consumption. These conditions lead to a kinked value function in which households demonstrate higher levels of risk aversion when the stakes are small than when they are large, which might explain the differences in low and high stakes preferences that we observe.

Chetty and Szeidl (2007) find that fixed liabilities account for over 50 percent of the average household's budget. As a household's income grows, the utility loss of paying an insurance deductible (of say \$500) would be expected to decrease because it represents a smaller proportion of the household's discretionary consumption. This argument predicts that (risk averse) households with higher incomes would select higher deductibles in this program. Instead, we find that households with higher property values (a proxy for income) are slightly more likely to select a low deductible. We estimate an ordered logit, regressing the deductible choice on property values (in \$10,000) for household i

$$Deductible_i = \alpha_{jt} + \beta\ PropertyValue_i + \varepsilon_i$$

where α_{jt} are $County \times Year$ fixed effects. We find a statistically significant, negative relationship between property values and deductible choices, an odds ratio of $\beta = 0.986$ ($s.e. = 0.0013,\ p < 0.01$). We conclude that such commitments may influence the decisions of some households in our data, but our results do not seem to be explained by consumption commitments (Online Appendix A.3.2).

Mortgages and strategic non-repayment. The differences in households' preferences across low and high stakes might be explained by their use of mortgages: homeowners may be unconcerned about losses beyond their equity. By modeling households' potential loss as the total value of the home, we would overestimate their exposure and conclude that households selecting low coverage limits are less risk aversion than they actually are. If households are insuring based on their limited liability for damage to the home, we would expect them to select lower coverage limits relative to their replacement costs in "non-recourse" states, where lenders' ability to collect on defaulted mortgages is limited.[29] We find a small difference in recourse and non-recourse states. For example, 75 percent of households in non-recourse states fully insure compared to 78 percent of households in recourse states. The magnitude of this difference on our preference estimates is negligible and cannot explain the large disparity between low and high stakes (Online Appendix A.3.3).

Program requirements. Households in our baseline sample are required to insure against flood if they have a mortgage from a federally regulated lender. This mortgage requirement does not seem to be consistently enforced (Dixon et al., 2006); however, contracts may not reflect household preferences if, say, lenders determined coverage limits (e.g., required that households fully insure their homes). We compare coverage limit selections for our baseline sample with households that participate in the program but live in a zone not subject to the federal requirement. Their coverage limit selections are remarkably similar, and so would seem to support the assumption that contract coverage limits reflect household preferences. For example, 78 percent of households fully insure in the baseline sample versus 80 percent fully insure in the zone where flood insurance is not required (Online Appendix A.3.4).

Household wealth. Perhaps households are considering a much lower or higher level of wealth than what we assume in Section 4, affecting our risk preference estimates**.** We set wealth equal to varying proportions of the replacement cost ranging from 5 percent to 10 times its value. We also set wealth equal to the home's property value instead of the replacement cost. Across wealth assumptions, our findings are qualitatively consistent with the main results (Section 4.1.2): households' insurance decisions are difficult to explain with CRRA expected utility. For example, assuming that household wealth is 10 times its replacement cost (on average about $1

---

[29] Non-recourse laws prevent lenders from seeking a deficiency judgment when the sale from a foreclosure is insufficient to cover the homeowner's liability on a defaulted mortgage. Orlando (2011) identifies Alaska, Arizona, California, Hawaii, Minnesota, Montana, North Dakota, Oklahoma, Oregon, and Washington as non-recourse states during the same time period as our data.

million in absolute terms) leads to risk aversion estimates of $\rho_{Deductible} = 1{,}378$ and $\rho_{CoverageLimit} = 56$ (Online Appendix A.3.5).

Policy tenure. Previous research identifies that households often maintain the same financial contract terms despite changes over time that would be expected to lead them to update their contracts. For example, Sydnor (2010) finds that while new customers frequently select one of the lowest deductible options for their homeowners insurance in his data, households with longstanding policies (e.g., 15 years) are even more likely to have a low deductible. In this context, households' insurance contract decisions may result from inertia rather than current risk preferences. Like Michel-Kerjan, de Forges, and Kunreuther (2012), we find that the mean flood insurance policy tenure during this time is four years (the median is two years), reducing concerns that the households' contract choices reflect severely outdated preferences. We find a small effect of contract duration on households' choices. For example, 44 percent of households with policy tenures in the first quartile (1 year or less) selected the lowest deductible while 46 percent of households with policy tenures in the fourth quartile (5 years or greater) did so. Households with first quartile policy durations fully insured their homes 73 percent of the time while those with fourth quartile policy durations did so 78 percent of the time.

Deductible menu defaults. Households are influenced by program default options in many contexts (e.g., Madrian and Shea, 2001). Most households (77 percent) select the default deductible option in our data, which might indicate inattention or that the defaults are judiciously chosen to match households' risk preferences. Our baseline sample includes two deductible defaults: one group has a $500 default and another a $1,000 default. Households with the $1,000 default deductible have riskier homes with higher premiums and so this default may result from this group's tendency to select higher deductibles. For example, about 6 percent of these households select the highest deductible available $5,000 while only 1 percent of households with the $500 deductible do so. We find a modest effect of deductible menu defaults on our estimates. We re-estimate our CRRA expected utility models excluding households who adopted the default option and find risk aversion estimates of $\rho_{Deductible}$ = 108 and $\rho_{CoverageLimit}$ = 2.6 (Online Appendix A.3.6). We also include a household's assigned default deductible in the model of heterogeneous preferences (Online Appendix A.2). While standard models predict that a household's assigned default deductible would not influence the preferences inferred from its contract choices, we find that being assigned a default of $1,000 (rather than the $500 default) reduces a households' relative risk aversion by 3.6 after controlling for other factors such as the age and property value of the home.

Misunderstandings about the insurance coverage limit. A large group of households (35 percent) purchased coverage limits that were higher than their replacement costs. We tend to agree with Mossin's (1968) explanation that households may buy too much insurance to address the risk that their replacement cost is higher than anticipated; however, perhaps instead these households did not understand the contract that they were purchasing, thinking that they could receive larger payments than the replacement cost. Our models, which account for the contract structure, may not correctly estimate the risk preferences of households that misunderstood it. Given this possibility, we re-estimate our CRRA expected utility models excluding these households and find qualitatively similar results to our estimates with the baseline sample (e.g., $\rho_{Deductible}$ = 147 and $\rho_{CoverageLimit}$ = 2.5, Online Appendix A.3.7).

Events with small probabilities. A great deal of research on risk preferences, especially lab research, considers events with likelihoods greater than the average annual claim rate of 1.3 percent in our data. Perhaps our preference estimates are an anomalous byproduct of examining relatively rare events. We also examine a population that is vulnerable to inundation as well as wave damage and their average annual claim rate is 4.1 percent. This claim rate is similar to Sydnor (2010) and Barseghyan, Prince, and Teitelbaum (2011) who examine homeowners insurance deductibles. Premiums are much higher in this zone, with a mean of $1,497 (and median of $1,201), almost 3 times the average in the baseline sample. Households in this zone choose higher deductibles: 78 percent choose one of the lowest two deductibles (versus 94 percent in the baseline sample) and 12 percent choose the highest deductible of $5,000 (versus 3.8 percent in the baseline sample). Eighty percent of these households fully insure. Our risk preference results are qualitatively similar for this higher risk sample to our baseline. For example, CRRA expected utility models result in relative risk aversion of 139 for the deductible and 3.6 for the coverage limit (Online Appendix A.3.8).

6 Conclusion

We assess whether households' risk preferences for small and large stakes losses are well characterized by commonly applied expected utility models. Previous research on insurance decisions evaluates deductible choices to examine risk preferences. We build on this research using a unique dataset that also contains information about large stakes exposures. We examine the deductible and coverage limit decisions from the flood insurance contract choices of over 100,000 households.

We find that both households' small and large stakes choices are difficult to explain with standard expected utility of wealth models. Our results confirm previous research that households' tendency to select a low deductible suggests extremely high relative risk aversion (e.g., Sydnor, 2010). Regarding coverage limits, standard models predict that households would partially insure (i.e., select low coverage limits relative to their home's replacement cost, Mossin, 1968). For example, our CRRA expected utility with $\rho = 1$ predicts that households should, on average, select a coverage limit of 29 percent of their replacement cost. Instead, over three-quarters of households fully insure. The coverage limits of the quarter of our households that partially insure generally suggest single digit risk aversion; however, only relative risk aversion of at least $\rho = 221$ would motivate CRRA expected utility maximizers to fully insure at the rates that we observe.

Allowing for probability distortions greatly improves model descriptions of households' low and high stakes decisions. Using CRRA rank dependent utility, we find that households' behavior follows two of the tenants of prospect theory: households demonstrate diminishing sensitivity to losses and overweight small probabilities. However, in rank dependent utility models, households' low and high stakes decisions continue to be characterized by different preferences, as indicated by different model parameters. We conclude that while our context stands out as one that facilitates making low and high stakes choices in concert as they are part of a single contract covering a continuum of losses, households seem to treat them as separable decisions toward which they have differing attitudes.

We find Mossin's (1968) discussion of his theoretical result on partially insuring remarkably prescient, though he may not have anticipated the degree to which expected utility theory misrepresents households' insurance choices in this setting: "Casual empirical evidence seems to contradict the conclusion, however; some of our best friends take full coverage. Several explanations for such behavior can of course be offered, among them: (a) they simply behave irrationally, for example, by not bothering to determine optimal coverage; (b) there may be some uncertainty as to what will be the actual evaluation of the property in case of damage. We have not taken this kind of randomness into account in our theory, but it seems a perfectly reasonable hypothesis that it will lead to a number of observed cases of full (or over-) insurance; (c) they may uniformly overestimate the probability distribution for damage…" (p. 558). We find that households overestimate flood probabilities and over-insure (buying coverage limits greater than their replacement costs) in the way Mossin describes.

A substantial literature has emerged since Mossin's paper that examines the systematic biases and heuristics utilized by individuals for making choices under risk and uncertainty (Kahneman, 2011) and their relationship to insurance purchase and coverage decisions (Kunreuther, Pauly, and McMorrow, 2012). These behavioral features may also explain the differences in low and high stakes insurance preferences that we observe in the rank dependent utility model.

Our findings present an interesting public policy dilemma. A stated objective of the NFIP is increasing take-up of flood insurance as a means to reduce federal spending on flood events (National Research Council, 2015). Recent policy interventions emphasize choice architecture that induces households to make choices that better align with normative models such as expected utility (e.g., Thaler and Sunstein, 2008); however, convincing households in our data to act as expected utility maximizers would run counter to this program's objective. At least under the program's current pricing, such interventions would lower coverage limits, leaving households more exposed to severe events.

Finally, we highlight two limitations of our analyses. First, we cannot rule out the possibility that households may maximize some unconsidered utility function or that myriad value functions may be represented in our data (e.g., List, 2004, and Harrison and Rutström, 2009, find a mixture of expected utility and prospect theory preferences in their data). We have instead examined the most frequently used utility functions and considered whether these individually characterize households' decisions well. Second, how our results regarding large stakes flood insurance preferences generalize to other domains is unclear. While our findings of single-digit relative risk aversion for households' large stakes decisions are consistent with estimates from other domains (e.g., life cycle models, Gourinchas and Parker, 2002, and equity premiums, Barro and Jin, 2011), individuals' decisions may more closely align with expected utility theory in those settings than ours. These points motivate additional research.

References

Aitkin, M. (1987). Modelling variance heterogeneity in normal regression using GLIM. *Applied Statistics*, *35*(3), 332-339.

David A., Mullen, K. M., Peterson, B. G., Ulrich, J. (2016). 'DEoptim': Differential Evolution in 'R'. version 2.2-4.

Arrow, K. J. (1971). The theory of risk aversion. *Essays in the theory of risk-bearing*, 90-120. Markham Publishing Co., Chicago.

Barro, R. J., & Jin, T. (2011). On the size distribution of macroeconomic disasters. *Econometrica*, *79*(5), 1567-1589.

Barseghyan, L., Molinari, F., O'Donoghue, T., & Teitelbaum, J. C. (2013). The nature of risk preferences: Evidence from insurance choices. *American Economic Review*, *103*(6), 2499-2529.

Barseghyan, L., Molinari, F., O'Donoghue, T., & Teitelbaum, J. C. (2015). Estimating risk preferences and risk perceptions in the field. Forthcoming in *Journal of Economic Literature*.

Barseghyan, L., Prince, J., & Teitelbaum, J. C. (2011). Are risk preferences stable across contexts? Evidence from insurance data. *American Economic Review*, *101*(2), 591-631.

Berg, J., Dickhaut, J., & McCabe, K. (2005). Risk preference instability across institutions: A dilemma. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(11), 4209-4214.

Bernoulli, D. (1738). Specimen Theoriae Novae de Mensura Sortis [Exposition of a new theory on the measurement of risk], Commentarii Academiae Scientiarum Imperialis Petropolitanae, Tomus V [Papers of the Imperial Academy of Sciences in Petersburg, Vol. V], 175-192. Translated and reprinted in *Econometrica*, *22*(1), 23-36, 1954.

Braun, M., Fader, P. S., Bradlow, E. T., & Kunreuther, H. (2006). Modeling the "pseudodeductible" in insurance claims decisions. *Management Science*, *52*(8), 1258-1272.

Camerer, C. F., & Ho, T. H. (1994). Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty*, *8*(2), 167-196.

Cameron, A. C., & Trivedi, P. K. (2005). *Microeconometrics: Methods and Applications*. Cambridge University Press.

Chetty, R. (2015). Behavioral economics and public policy: A pragmatic perspective. *The American Economic Review*, *105*(5), 1-33.

Chetty, R. (2006). A new method of estimating risk aversion. *The American Economic Review*, *96*(5) 1821-1834.

Chetty, R., & Szeidl, A. (2007). Consumption Commitments and Risk Preferences. *The Quarterly Journal of Economics*, *122*(2), 831-877.

Cohen, A., & Einav, L. (2007). Estimating risk preferences from deductible choice. *The American Economic Review*, *97*(3), 745-788.

Dixon, L., Clancy, N., Seabury, S. A., & Overton, A. (2006). *The National Flood Insurance Program's Market Penetration Rate*. RAND Corporation.

Dyer, J. S., & Sarin, R. K. (1982). Relative risk aversion. *Management Science*, *28*(8), 875-886.

Einav, L., Finkelstein, A., Pascu, I., & Cullen, M. R. (2012). How general are risk preferences? Choices under uncertainty in different domains. *The American Economic Review*, *102*(6), 2606-2638.

FEMA [U.S. Federal Emergency Management Agency] (2015). National Flood Insurance Program Community Rating System Coordinator's Manual. URL: http://www.fema.gov/media-library/assets/documents/8768. Last accessed: February 26, 2016.

Friend, I., & Blume, M. E. (1975). The demand for risky assets. *The American Economic Review*, 900-922.

Gallagher, J. (2014). Learning about an infrequent event: evidence from flood insurance take-up in the United States. *American Economic Journal: Applied Economics*, *6*(3), 206-233.

GAO [U.S. Government Accountability Office] (2006). Federal Emergency Management Agency: Challenges for the National Flood Insurance Program (25-JAN-06, GAO-06-335T).

Gourinchas, P. O., & Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, *70*(1), 47-89.

Harrison, G. W., Humphrey, S. J., & Verschoor, A. (2010). Choice under uncertainty: evidence from Ethiopia, India and Uganda. *The Economic Journal*, *120*(543), 80-104.

Harrison, G. W., & List, J. A. (2004). Field experiments. *Journal of Economic literature*, *42*(4), 1009-1055.

Harrison, G. W., & Rutström, E. E. (2008). Risk aversion in the laboratory. In *Research in Experimental Economics*, Vol. 12, ed. James C. Cox and Glenn W. Harrison, 41-196. Bingley, UK: Emerald Group Publishing Limited.

Harrison, G. W., & Rutström, E. E. (2009). Expected utility theory and prospect theory: One wedding and a decent funeral. *Experimental Economics*, *12*(2), 133-158.

Hey, J. D., Morone, A., & Schmidt, U. (2009). Noise and bias in eliciting preferences. *Journal of Risk and Uncertainty*, *39*(3), 213-235.

Hey, J. D., & Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica*, *62*(6), 1291-1326.

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, *92*(5), 1644-1655.

Johnson, E. J., Hershey, J., Meszaros, J., & Kunreuther, H. Framing, probability distortions, and insurance decisions. *Journal of Risk and Uncertainty* 7, no. 1 (1993): 35-51.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263-291.

Kahneman, D., Wakker, P. P., & Sarin, R. (1997). Back to Bentham? Explorations of experienced utility. *The Quarterly Journal of Economics*, *112*(2), 375-406.

Kousky, C., & Michel-Kerjan, E. (2015). Examining flood insurance claims in the united states: six key findings. *The Journal of Risk and Insurance,* DOI: 10.1111/jori.12106

Kousky, C., & Shabman, L. (2014). Pricing flood insurance: How and why the NFIP differs from a private insurance company. Resources for the Future Discussion Paper 14-37.

Kunreuther, H., Meszaros, J., Hogarth, R. M., & Spranca, M. (1995). Ambiguity and underwriter decision processes. *Journal of Economic Behavior & Organization*, *26*(3), 337-352.

Kunreuther, H. C., Pauly, M. V., & McMorrow, S. (2013). *Insurance and behavioral economics: Improving decisions in the most misunderstood industry*. Cambridge University Press.

List, J. A. (2004). Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica*, *72*(2), 615-625.

Madrian, B. C., & Shea, D. F. (2001). The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly Journal of Economics*, *116*(4), 1149-1187.

McFadden, D. (1974). The measurement of urban travel demand. *Journal of Public Economics*, *3*(4), 303-328.

Michel-Kerjan, E. O. (2010). Catastrophe economics: the national flood insurance program. *The Journal of Economic Perspectives*, *24*(4), 165-186.

Michel-Kerjan, E., Lemoyne de Forges, S., & Kunreuther, H. (2012). Policy tenure under the U.S. National Flood Insurance Program (NFIP). *Risk Analysis*, *32*(4), 644-658.

Mossin, J. (1968). Aspects of rational insurance purchasing. *Journal of Political Economy*, *76*(4), 553-568.

National Research Council (2015). National Flood Insurance Program history and objectives in *Affordability of National Flood Insurance Program Premiums, Report 1*. pp. 23-34. The National Academies Press, Washington, D.C.

NFIP [National Flood Insurance Program] (2006a). Flood insurance manual, effective May 1, 2006. URL: https://www.fema.gov/flood-insurance-manual-effective-may-1-2006. Last accessed: February 2, 2017.

NFIP [National Flood Insurance Program] (2006b). Transaction record reporting and processing plan for the Write Your Own Program. URL: http://bsa.nfipstat.fema.gov/manuals/manuals.html. Last accessed: March 1, 2016.

Orlando, James (2011). Comparison of state laws on mortgage deficiencies and redemption periods. Office of Legislative Research, Research Report. URL: https://www.cga.ct.gov/2010/rpt/2010-R-0327.htm

Prelec, D. (1998). The probability weighting function. *Econometrica, 66*(3), 497-527.

Preston, M. G., & Baratta, P. (1948). An experimental study of the auction-value of an uncertain outcome. *The American Journal of Psychology*, *61*(2) 183-193.

Price, K. V., Storn, R. M., and Lampinen, J. A. (2006). *Differential Evolution - A Practical Approach to Global Optimization.* Berlin Heidelberg: Springer-Verlag. ISBN 3540209506.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization*, *3*(4), 323-343.

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

Rabin, M. & Thaler, R. H. (2001). Anomalies: risk aversion. *Journal of Economic Perspectives*, *15*(1), 219-232.

Saha, A. (1993). Expo-power utility: A 'flexible'form for absolute and relative risk aversion. *American Journal of Agricultural Economics*, *75*(4), 905-913.

Sydnor, J. (2010). (Over) insuring modest risks. *American Economic Journal: Applied Economics*, *2*(4), 177-199.

Szpiro, G. G. (1986). Measuring risk aversion: an alternative approach. *The Review of Economics and Statistics*, *68(1)*, 156-159.

Thaler, R. H., & Benartzi, S. (2004). Save more tomorrow™: Using behavioral economics to increase employee saving. *Journal of Political Economy*, *112*(S1), S164-S187.

Thaler, R. H., & Sunstein, C. R. (2008) *Nudge: Improving Decisions About Health, Wealth, and Happiness.* Yale University Press.

Train, K. E. (2009). Logit. *Discrete choice methods with simulation*. Chapter 3, 34-75, Cambridge university press.

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

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, *57*(2), 307-333.

Wakker, P. P. (2008). Explaining the characteristics of the power (CRRA) utility family. *Health Economics*, *17*(12), 1329-1344.

Western, B., & Bloome, D. (2009). Variance function regressions for studying inequality. *Sociological Methodology*, *39*(1), 293-326.

Yaari, M. (1987). The dual theory of choice under risk. *Econometrica*, *55*(1), 95-115.

FOR ONLINE PUBLICATION

A Online Appendix

A.1 Estimating Policyholder Risk

This appendix describes our claim rate and loss distribution models.

A.1.1 Claim Rates

We assume that policyholders rely on the same information as the flood insurance program in estimating their flood risk. The included observables are those used in determining flood insurance premiums, which Table 7 lists and defines.

The flood insurance program considers whether a home was built before its flood insurance rate maps (FIRM) were developed. Zoning regulations and building codes intend to reduce vulnerability in designated flood hazard areas and so flood claims and losses may substantially differ depending on whether a home was built pre- or post-FIRM. We model flood risk separately for pre- and post-FIRM dwellings. We estimate the likelihood of having a claim using a random effects panel logit for household i in community j

$$p_i = \beta_0 + \beta_1 CRSclass_j + \beta_2 I(Basement_i) + \beta_3 D(Obstruction_i) + \beta_5 D(Floors_i) + \beta_6 I(Mobile_i) + \beta_7 D(Elevation_i) \times D(Source_i) + \varepsilon_i \quad (A1)$$

where I is an indicator variable and D describes a dummy set. For example, $D(Floors)$ indicates whether the home has one, two, or three or more floors. Only a policyholder's first claim in each year is included in the analysis.[30] Table 8 provides the results of the panel logit model of household claim rates with coefficients reported as log odds ratios. These results are consistent with previous research. For example, Kousky and Michel-Kerjan (2015) find in their assessment of NFIP claims that a community's Community Rating System score, the number of floors in a home, its elevation, and whether it has a basement all significantly influence flood claims and in qualitatively similar ways to our findings.

[30] Focusing on the first claim in a year great facilitates modeling the household's decision in Section 3.3.1 Allowing for more than one claim is computationally prohibitive as it requires accounting for all possible combinations of losses from each flood event. Two percent of policies with claims have a second claim in the same year. If the likelihood of having more than one claim in a year is equal across households, the probability of having at least two claims in a year for the household with the median claim rate is 0.00027 ($0.0133 \times 0.02 = 0.00027$).

.

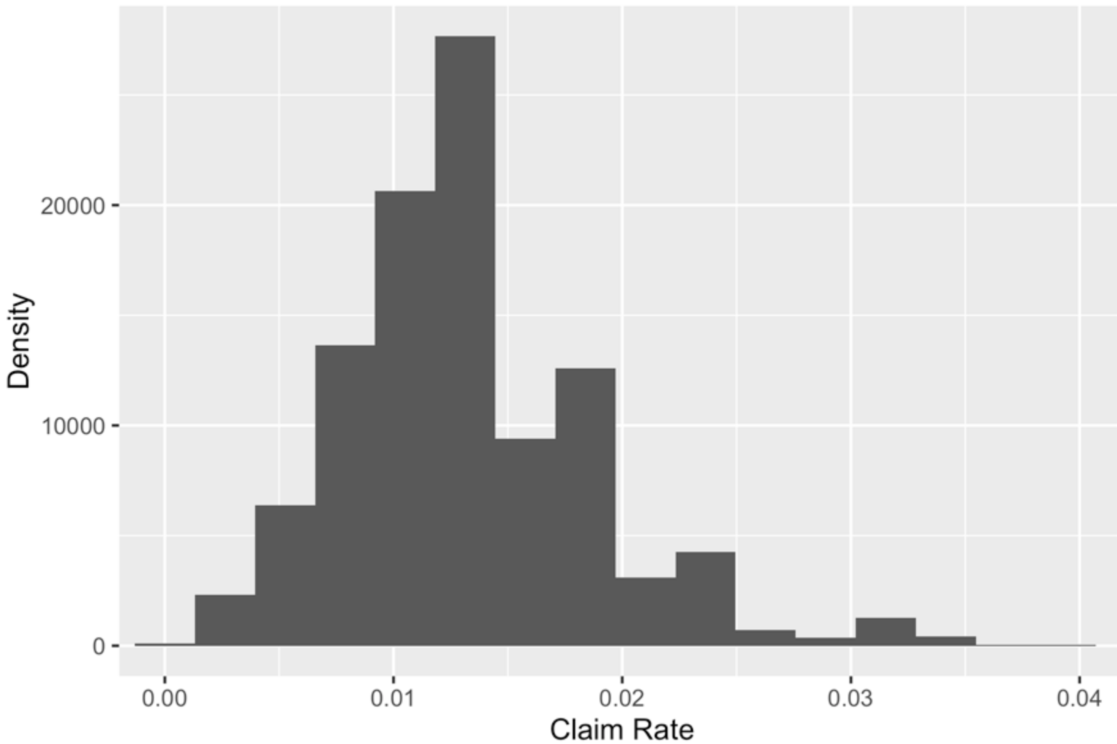

Figure 5 Claim Rate Estimates

Note: Claim rates using random effects panel logit. Baseline sample of 103,080 policies.

Figure 5 shows the histogram of claim rates for our baseline sample. The claim rates bunch around elevation relative to the 100-year flood plain, which FEMA rounds to the nearest foot, taking values from 2 or more feet below the flood plain to 5 or more feet above the flood plain. For example, while the median home in our data has an elevation of 1 foot above the flood plain, none of the homes with estimated claims rates above 3 percent are elevated above the flood plain.

Table 7. Explanatory Variables Used in Regressions

Explanatory variables	*Description*
CRS class	The community's score on the Community Rating System (CRS). The CRS is a voluntary program that rewards communities for taking actions to mitigate flood risk beyond minimum NFIP requirements. Community actions reduce policyholder premiums by up to 45%. CRS class is the associated premium reduction, ranging from 0 to 45.
Basement	Indicates that the property has a basement
Obstruction	Indicates that an elevated building has an enclosed area and/or machinery attached to the building below the lowest floor. This variable is not applicable to un-elevated buildings. Thus, it takes three values: 0, 1, not applicable.
Source	Indicates that elevation estimates were collected by a contracted engineer. Otherwise, the elevation data are provided by the NFIP and may not have been assessed at the level of the individual home. The source of elevation data is not always reported in the policy database, and in those cases, we record the observation's source as "Not Indicated."
Elevation	An estimate of the elevation in feet of a policyholder's home relative to the 100-year floodplain. Values are rounded to the nearest foot and bound with the lowest and highest elevation at -2 and 4 feet, respectively.
Mobile	Indicates that the structure is a manufactured or mobile home.
Floors	Number of floors in the home, taking three possible values: 1, 2, or 3 or more

Note: More information can be found on these variables from NFIP (2006).

Table 8. Claim Rate Regressions

		Pre-FIRM		Post-FIRM	
		Coefficient	Std. Error	Coefficient	Std. Error
Intercept		-2.81	1.158	-3.74	0.43
CRS Class		-0.03	0.0004	-0.06	0.001
Basement		0.00	0.009	-0.42	0.019
Obstruction		0.19	0.010	0.05	0.015
Obstruction, Not Applicable		-0.30	0.005	-0.48	0.009
Mobile		0.02	0.021	-0.02	0.027
Floors (Reference Group: Floors = 1)					
2		0.06	0.007	-0.04	0.009
3 or more		0.14	0.011	-0.29	0.018
Elevation x Source (Reference Group: Elevation = -2 x Source = NFIP)					
Elevation	*Source*				
-2	*Contractor*	-0.52	1.158	0.42	0.433
-2	*Not Indicated*	-0.50	1.158	0.75	0.432
-1	*NFIP*	-	-	-0.32	0.866
-1	*Contractor*	-1.04	1.158	0.07	0.434
-1	*Not Indicated*	-0.43	1.158	0.29	0.433
0	*NFIP*	-1.16	1.316	0.35	0.528
0	*Contractor*	-1.29	1.158	-0.33	0.433
0	*Not Indicated*	-0.61	1.158	0.39	0.432
1	*NFIP*	-0.88	1.255	-0.75	0.633
1	*Contractor*	-1.65	1.158	-0.66	0.433
1	*Not Indicated*	-1.12	1.158	-0.12	0.432
2	*NFIP*	-1.75	1.307	-0.94	0.603
2	*Contractor*	-2.08	1.158	-0.96	0.433
2	*Not Indicated*	-1.47	1.158	-0.50	0.433
3	*NFIP*	-1.59	1.377	-1.12	0.733
3	*Contractor*	-2.16	1.158	-1.09	0.433
3	*Not Indicated*	-1.61	1.158	-0.66	0.433
4	*NFIP*	-2.19	1.549	-0.44	0.506
4	*Contractor*	-2.05	1.158	-0.85	0.433
4	*Not Indicated*	-1.62	1.158	-0.35	0.432
Not Indicated	*NFIP*	-0.94	1.158	-0.97	0.434
Not Indicated	*Contractor*	-1.23	1.159	-2.03	0.470
Not Indicated	*Not Indicated*	-1.19	1.158	-1.11	0.434
Log Likelihood		-894,641		-381,598	
Observations		9,108,126		6,171,368	

Note: Random effects panel logit panel regression. Coefficients are reported as log odds ratios.

A.1.2 Loss Distributions

We model household losses beginning with an examination of the loss distribution, considering losses for all households in the same flood risk category concurrently. From this approach, we find that our flood loss distributions are approximately log-normal. Finally, we fit the parameters of the log-normal distribution for each household based on its observables.

A.1.2.1 Parametric distributions

This section compares parametric specifications for modeling flood losses. Throughout, we use A zone, pre-FIRM as an example. Given a claim event, the distribution of potential property losses is influenced by the type of event that occurs: events that affect many policyholders result in larger expected losses for each policyholder. We model observations as a mixture of two loss generating processes: (1) isolated loss events, floods affecting an individual or small group of policyholders (e.g., due to unusually heavy, localized rain), or (2) correlated loss events, losses due to a storm, hurricane, etc., affecting many policyholders. We define a correlated loss event as an observation of at least 30 claims with a date of loss on the same day in the same state. That event continues for each consecutive day with at least 30 claims in the same state. The 30 claims threshold is located at the 95th percentile of the distribution of claims by date of loss by state. These criteria create 1,319 correlated loss events out of a total of 635,220 flood insurance claims occurring between 1982 and 2009. All observations that are not associated with a correlated loss event are considered an isolated loss event.

Observations from isolated loss events are weighted by

$$\phi_h = \frac{\overline{\pi_h}}{n_h}$$

where $\overline{\pi_h}$ describes the average across years of the percent of claims generated by isolated loss events and n_h indicates the total number of claims generated by isolated loss events ($n_h =$ 85,075). On average, 22 percent of claims are generated by correlated loss events. Observations for correlated loss events are weighted by

$$\phi_{m,j} = \frac{(1 - \overline{\pi_h})}{n_m * n_{m,j}}$$

where n_m indicates the total number of correlated loss events and $n_{m,j}$ indicates the total number of claims generated by correlated loss event j.

These weights provide empirical loss distributions for the isolated and correlated loss events. No time trends are present in the means and variances of the empirical loss distributions across years: regressing respectively the mean and variance of losses as a percent of the property value by year on time results in non-significant F-statistics of 0.89 and 0.60.

Using the weighted empirical loss distribution, we fit parametric distributions using maximum likelihood estimation (MLE) and compare models using the Anderson-Darling test, Kolmogorov-Smirnov test, and the Akaike Information Criterion (AIC). Comparisons consistently indicate that the log-normal is the best fitting parametric model; this result is consistent across flood zones. Table 9 provides the results for the A zone, Pre-FIRM, correlated loss events distribution for a subset of the tested parametric distributions.

As a sensitivity test, we generate several alternative distributions: (1) defining correlated loss events using a minimum of 10 claims (rather than 30) with the same date of loss in a state, (2) using a minimum of 100 claims, (3) treating all claims as independent and equally weighted. Table 10 shows deciles from the cumulative distributions using these alternative assumptions in Columns 1 through 5; we explain Column 6 in the next section. Excluding the case in which all losses are equally weighted, the distributions are quite similar across approaches. The log-normal results in a slightly thicker right tail than the non-parametric. The equally weighted approach results in the greatest probability of large losses due to severe events.

Table 9. Fit Comparisons Across Parametric Distributions

	A-D	K-S	AIC
Log-normal	53.41	0.060	-14949.5
Pareto III	65.61	0.064	-14286.2
Pareto IV	68.79	0.066	-14364.4
Pareto II	72.06	0.071	-14322.7
Generalized Extreme Value	85.10	0.063	-14175.4
Gamma	143.86	0.100	-13966.1

Note: The results are for the A zone, Pre-FIRM, correlated loss event data. The Anderson-Darling (A-D) and Kolmogorov-Smirnov (K-S) tests and the Akaike Information Criterion suggest that the log-normal is the parametric distribution that best fits the loss data.

Table 10. Cumulative Distributions for A Zone, pre-FIRM

	Homogeneous					Heterogeneous
	(1)	(2)	(3)	(4)	(5)	(6)
Loss	Log-normal	Nonparametric (30 claims)	Nonparametric (10 claims)	Nonparametric (100 claims)	Nonparametric, equally weighted	Log-normal
10.0	53.1	55.3	58.0	51.9	40.6	50.6
20.0	73.8	71.7	73.9	68.5	55.4	69.0
30.0	83.3	80.3	82.1	77.6	64.5	78.0
40.0	88.4	86.5	87.9	84.3	71.8	83.5
50.0	91.6	90.8	91.6	89.0	77.6	87.1
60.0	93.6	93.5	94.0	92.2	81.9	89.6
70.0	95.0	95.4	95.7	94.3	85.3	91.5
80.0	96.0	96.7	97.0	96.0	88.2	92.8
90.0	96.8	97.9	98.0	97.4	91.2	93.9

Note: Alternative specifications of the cumulative distribution for losses. Losses are measured as a percent of the structure's value. The parametric distributions use the correlated loss event definition of 30 claims. Columns 1 through 5 report deciles assuming a homogeneous loss distribution across all policyholders in the A zone, pre-FIRM. Column 1 reports deciles for the log-normal distribution. Columns 2 through 4 report for nonparametric distributions for which weights are based on definitions of the correlated loss events. Columns 2, 3, and 4 respectively use definitions of a correlated loss event as a minimum of 30, 10, and 100 claims in the same state on the same day. Column 5 reports deciles for the nonparametric distribution using equal weights for all observations. Column 6 reports deciles using the median log-normal parameter μ and σ estimates fit for each household based on observables.

A.1.2.2 Household-level loss distributions

We develop loss distribution estimates for each policyholder using the log-normal distribution and our assumption that the loss generating process is a mixture of isolated and correlated loss events. Our approach is similar to that of Aitkin (1987) and Western and Bloome (2009). Aitkin proposes modeling variance heterogeneity to address heteroscedasticity. Western and Bloome note that the estimations of these variance models may themselves be of interest for research related to within-group differences and so adapt Aitkin's approach to study income inequality. They propose an iterative MLE. We use this approach to fit parameters μ and σ of the log-normal distribution. Consider the two-equation model

$$\mathrm{E}[\log l_i] = \boldsymbol{x}_i' \boldsymbol{\beta}$$
$$\log \sigma_i^2 = \boldsymbol{x}_i' \boldsymbol{\lambda}$$

where $l_i \in (0,1]$ is the flood loss as a proportion of the home value for policyholder $\boldsymbol{i}$, $\boldsymbol{x_i}$ is a vector of policyholder observables, and σ_i^2 is the estimated variance in losses for each policyholder. We fit the model using an iterative MLE approach:

1) Estimate $\widehat{\boldsymbol{\beta}}$ using a Tobit model, which is right censored at 0, as $\log l_i \in (-\infty, 0]$. Save the residuals, $e_i = \log l_i - \boldsymbol{x}_i' \widehat{\boldsymbol{\beta}}$.
2) Estimate $\hat{\boldsymbol{\lambda}}$ with a gamma regression of the squared residuals e^2, using a log link function. Save the fitted values, $\widehat{\sigma^2} = \exp(\boldsymbol{x}_i' \hat{\boldsymbol{\lambda}})$.
3) Estimate $\widehat{\boldsymbol{\beta}}$ using a Tobit model with weights $1/\hat{\sigma}^2$ and update the residuals.
4) Repeat steps 2 and 3 until the log likelihood converges.

This iterative approach (1) addresses heteroscedasticity in the mean model by weighting observations based on the fitted variance, and (2) corrects the standard errors in the variance model by increasing the precision of the coefficient estimates from the mean model (Western and Bloome, 2009).

Fixed effects in these regressions account for correlated loss events. We order correlated loss events by the number of claims for each event and bin the events every 5 percentiles, creating twenty fixed effects across the distribution of correlated loss events (i.e., vigintiles). For example, a fixed effect is included for all correlated loss events for which the number of claims is below the fifth percentile, another one for events with claims from the fifth and to tenth percentiles, etc.

We follow the estimation equation A1 that we use in our claim rate regressions, though some minor differences exist in our claims and policy data. For example, these regressions do not include a households Community Rating System score. Table 11 shows the output of the mean and variance models for our baseline sample (Zone A). Its results are also qualitatively consistent with the findings of Kousky and Michel-Kerjan (2015) with respect to explanatory variables such as elevation, whether the home has a basement, and whether the home is permanent or a mobile home.

Loss observations are associated with a specific event; however, we are attempting to estimate each policyholder's loss distribution across all possible events. That is, for each household, we would like the expected loss given the explanatory variables and the occurrence of a loss event $\mathrm{E}[l_i|l_i > 0 \cap \boldsymbol{x_i}]$ but the fixed effects model provides an observation given specific event m_j, $\mathrm{E}[l_i \big| l_i > 0 \cap \boldsymbol{x_i} \cap m_j\big]$. To address this in our probability estimates, we weight each event by its probability in these data. For example, the model of log losses can be written as

$$\log l_i = \alpha + \boldsymbol{x_i'\beta} + \boldsymbol{FE'\gamma} + \epsilon_{\boldsymbol{i}}$$

where $\boldsymbol{\gamma}$ is the vector of coefficients on the fixed effects $\boldsymbol{FE}$. The probability from this equation can then be written as

$$\mathrm{E}[\log l_i|l_i > 0 \cap \boldsymbol{x_i}] = \overline{\pi_h}\alpha + (1 - \overline{\pi_h})(\alpha + \boldsymbol{\gamma})'\boldsymbol{\pi_{FE}} + \boldsymbol{x_i'\beta}$$

where $\overline{\pi_h}$ is the probability that an observed loss is generated from an isolated loss event (rather than a correlated loss event), and $\boldsymbol{\pi_{FE}}$ are the probabilities for each fixed effect event. In this case, each quantile has a 5 percent probability, given that a correlated loss event occurs. The term $\overline{\pi_h}\alpha + (1 - \overline{\pi_h})(\alpha + \boldsymbol{\gamma})'\boldsymbol{\pi_{FE}}$ is a constant and provides the intercept for the predictive model and so can be used to estimate fitted values for each household. The same approach is taken for the variance model. The mean and variance estimates for each household are then used to fit parameters of the log-normal distribution.

Table 10 provides the cumulative loss distribution for the median policyholder at each decile following this approach in Column 6. The median of these loss distributions results in a slightly higher probability of a total loss than the group-level log-normal (Column 1) and the weighted nonparametric distributions (Columns 2–4) but a lower probability than the distribution in which all observations are equally weighted (Column 5). Figure 1 illustrates a similar result using the entire baseline sample (the results in Table 10 are only for pre-FIRM homes in the baseline sample). The triangles in Figure 1 correspond with Column 2, the squares with Column 5, and the circles with Column 6.

A.2 Modeling Heterogeneous Preferences

We estimate a household's coefficient of risk aversion in our structural model using the regression equation

$$\rho_i = \gamma_j + \delta_t + \beta_1 PropertyValue_i + \beta_2 Age + \beta_3 I(DefaultDeductible_i = \$1{,}000) + \varepsilon_i \quad (A2)$$

where γ_j and δ_t are respectively region, using the nine U.S. Census Bureau divisions, and year fixed effects. Age is the age of the home. The property value serves as a proxy for wealth as well as other aspects of the household such as educational attainment may affect its risk preferences but are not included in the utility function. Property value and age are standardized (i.e., values are subtracted from the mean and divided by the standard deviation). Their coefficients show how a one standard deviation change in these variables affects risk aversion. *Table 12* provides the regression results.

Table 11. Models of Flood Loss as a Percent of Home Value

	Zone A, pre-FIRM				Zone A, post-FIRM			
	Mean Model (Tobit)		Variance Model (Gamma, log link)		Mean Model (Tobit)		Variance Model (Gamma, log link)	
	Coefficient	Std. Error	Coefficient	Std. Error	Coefficient	Std. Error	Coefficient	Std. Error
Intercept	-2.621	0.024	0.730	0.049	-3.032	0.037	1.036	0.050
Basement	-0.111	0.009	-0.037	0.016	-0.221	0.025	0.002	0.035
Mobile	0.492	0.023	0.077	0.039	0.901	0.042	-0.097	0.059
Floors (Reference group: Floors = 1)								
2	-0.297	0.007	-0.086	0.013	-0.289	0.013	-0.026	0.019
3 or more	-0.443	0.010	-0.143	0.019	-0.574	0.025	0.037	0.035
Elevation × Source (Reference group: Elevation = 0 × Source = Contractor)								
Elevation Source								
-2 NFIP	0.147	0.038	0.387	0.069	0.710	0.045	-0.311	0.068
-2 Contractor	0.220	0.046	0.235	0.089	0.061	0.038	0.029	0.055
-1 NFIP	0.036	0.044	0.379	0.080	0.431	0.049	-0.052	0.069
-1 Contractor	0.241	0.070	0.152	0.137	0.076	0.064	-0.010	0.093
0 NFIP	-0.179	0.024	0.239	0.051	0.169	0.031	-0.088	0.045
1 NFIP	-0.400	0.026	0.563	0.052	0.111	0.031	-0.014	0.045
1 Contractor	-0.152	0.032	0.171	0.065	-0.178	0.036	0.032	0.052
2 NFIP	-0.284	0.029	0.329	0.058	0.149	0.033	-0.028	0.049
2 Contractor	-0.284	0.036	0.372	0.070	-0.300	0.037	0.022	0.055
3 NFIP	-0.375	0.038	0.457	0.070	-0.117	0.040	0.033	0.057
3 Contractor	-0.528	0.045	0.481	0.083	-0.365	0.041	0.059	0.058
4 NFIP	-0.386	0.049	0.385	0.087	-0.181	0.049	0.165	0.066
4 Contractor	-0.567	0.058	0.337	0.106	-0.457	0.043	0.126	0.062
5 NFIP	-0.315	0.037	0.271	0.071	-0.301	0.043	0.188	0.059
5 Contractor	-0.479	0.046	0.445	0.080	-0.650	0.033	0.066	0.048
Not Indicated NFIP	0.145	0.022	0.022	0.047	0.672	0.049	-0.155	0.068
Not Indicated Contractor	0.312	0.023	0.039	0.049	0.738	0.093	-0.173	0.129
Log(scale)	0.257	0.001			0.360	0.003		
Quantile Fixed Effects	Yes		Yes		Yes		Yes	
Log Likelihood	-431,805		-484,037		-117,922		-164,652	
Deviance	495,371		1,262,186		156,414		411,915	
Num. obs.	409,179		409,179		126,014		126,014	
Right-censored	17,712				8,954			

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$, Standard errors in parentheses. Models predict the mean and variance of flood loss for A zone, pre-FIRM and A zone, post-FIRM. We model flood losses between 1982 and 2009. The mean and variance models are fit using an iterative maximum likelihood approach proposed by Aitkin (1987). The mean model is a Tobit with a dependent variable of log(loss) where loss is measured as a proportion of the structure's value. The variance model uses the squared residuals of the mean model as a dependent variable. In turn, the inverse of the predicted value of the variance model is used to weight observations in the mean model. Quantile fixed effects describe a set of dummies we include to account for major storms.

Table 12 Coefficient Estimates for CRRA Expected Utility Models Assuming Heterogeneous Preferences

Dependent Variable: Coefficient of Relative Risk Aversion

	Deductible		Coverage Limit	
	Coefficient	Standard errors	Coefficient	Standard Errors
Intercept	112.84	0.037	3.42	0.029
Property Value	39.23	0.022	0.25	0.015
Age of Home	-3.41	0.044	-0.14	0.013
I(Default Deductible = $1,000)	-4.38	0.019	-0.32	0.029
Region (Reference Group: East North Central)				
I(East South Central)	-21.88	0.254	-0.38	0.021
I(Middle Atlantic)	36.92	1.572	0.10	0.020
I(Mountain)	19.16	9.301	-0.08	0.099
I(New England)	24.72	0.783	0.68	0.070
I(Other)	-30.51	0.022	-0.39	0.052
I(Pacific)	60.49	3.342	0.17	0.040
I(South Atlantic)	31.75	0.950	-0.04	0.016
I(West North Central)	-20.94	0.734	-0.31	0.032
I(West South Central)	-1.17	0.500	-0.44	0.012
Year (Reference Group: 2008)				
I(2003)	-17.47	0.226	-0.44	0.019
I(2004)	-0.34	0.057	-0.37	0.036
I(2005)	-7.90	0.696	-0.19	0.011
I(2006)	2.86	0.494	-0.20	0.027
I(2007)	8.96	1.608	0.10	0.023
Scale	0.31	0.002	0.18	0.001
Log Likelihood	-164,935		-297,756	
AIC	164,973		297,794	
BIC	330,090		595,730	
Observations	183,080		183,080	

Note: *** p < 0.01, ** p < 0.05, * p < 0.1. Table shows the effects of a household's characteristics on its relative risk aversion based on its deductible and coverage limit decisions. Regressions follow Equation (A2). Property value and the age of the home are standardized variables (i.e., values are subtracted from the mean and divided by the standard deviation). Default deductible indicates whether households have a $500 or $1,000 default deductible. Region uses U.S. Census Bureau Divisions where "other" applies to U.S. territories such as Puerto Rico.

A.3 Robustness Tests and Alternative Explanations

This section provides technical details for the discussion in Section 5 and follows its structure. We discuss some of the most relevant alternative explanations for our results. Our analysis also builds on previous research regarding insurance decisions and risk preferences from Cohen and Einav (2007), Sydnor (2010), Barseghyan, Prince, and Teitelbaum (2011), Barseghyan et al. (2013), and others, and each of those papers provides a thoughtful discussion of alternative explanations that often also apply to our setting.

A.3.1 Other Utility Functions

Households' low and high stakes decisions may be characterized by some utility function other than CRRA. We consider constant absolute risk aversion (CARA) and expo-power preferences. CARA is perhaps the second most commonly employed utility function and assumes that risk aversion is independent of households' wealth. It takes the form

$$u(x) = -\exp(-rx)$$

where r is the coefficient of absolute risk aversion, which we fit using MLE. The expo-power utility function (Saha, 1993) can accommodate a variety of combinations of relative and absolute risk preferences (e.g., DARA, DRRA, etc.). It takes the form

$$u(x) = -\exp(-\tau x^{\psi})$$

with restrictions $\tau \neq 0$, $\psi \neq 0$, and $\tau\psi > 0$. Holt and Laury (2002) suggest that households' low and high stakes preferences might be reconciled by such a combination and use expo-power utility in their research. While CARA is nested in expo-power utility (when $\psi = 1$), we include a specific examination of it because of its prevalence.

Table 13 provides the results. Both CARA and expo-power preferences provide an improvement over CRRA in explaining both households' deductible and coverage limit preferences according to the AIC and BIC. A Vuong test leads to a similar conclusion, resulting in a fit that is significantly better at the 1 percent level in each case. The expo-power model parameters indicate a combination of increasing absolute risk aversion (IARA) and increasing relative risk aversion (IRRA) for the deductible model. For the expo-power coverage limit model, we find a combination of IRRA and constant absolute risk aversion. That is, expo-power cannot improve on the CARA model in explaining households' coverage limit decisions.

CARA and expo-power preferences correctly predict households' deductible choices at rates similar to CRRA. CRRA utility correctly predicts 47 percent of households' deductible, slightly more than CARA and expo-power utility, respectively 45 and 46 percent. CARA and expo-power predict that households would choose much smaller coverage limits than what they do: expo-power underestimates households' coverage limit ratios choices by an average of 27 percentage points (versus 34 for CRRA).

Like CRRA preferences, CARA and expo-power preferences result in substantially different model parameters for households' deductible and coverage limit decisions. For example for CARA, absolute risk aversion is 17 times larger for the deductible decision.

Thus, despite its flexibility, expo-power preferences do not overcome the problems identified in Section 4.1.2 for CRRA preferences: 1) even using maximum likelihood estimate in our data, these models do not explain the average households' decision well, and 2) that households' preferences do not seem to be similar across low and high stakes.

Table 13 CARA and Expo-Power Model Results

	Parameters			Log Likelihood	AIC	BIC
CARA	r		σ			
Deductible	0.00093		0.204	-153,529	153,533	307,081
	(2.4e-06)		(0.0014)			
Coverage Limit	0.000055		0.208	-294,245	294,249	588,513
	(2.8e-07)		(0.0012)			
Expo-power	τ	ψ				
Deductible	0.00018	1.14	0.202	-152,867	152,873	305,769
	(3.3e-06)	(0.0015)	(0.0014)			
Coverage Limit	0.000055	1.00	0.208	-294,245	294,251	588,525
	(7.3e-07)	(0.00079)	(0.0012)			

Note: Standard errors are in parentheses. Table compares utility models describing households' deductible and coverage limit decisions, comparing outcomes using the Akaike information criterion (AIC) and the Bayesian information criterion (BIC). For the AIC and BIC, lower values indicate better fit. The remaining models use maximum likelihood estimation to fit model parameters. Baseline sample of 103,080 policies.

A.3.2 Consumption Commitments

We exploit variation in property values in our data to assess for empirical evidence that households' preferences for low deductibles are explained by consumption commitments. While Chetty and Szeidl's (2007) describe consumption commitments as a percent of income, all households in our data face the same deductible menu. As a household's income grows, the utility loss of paying an insurance deductible (of say $500) would be expected to decrease because it represents a smaller proportion of the household's discretionary consumption. This argument predicts that (risk averse) households with higher incomes would select higher deductibles in this program.

Figure 6 shows households' property values (as a proxy for income) and their deductible choices. We do not observe a consistent relationship; households with high and low property values prefer low deductibles. We also estimate an ordered logit model using the deductible choice and property value (in $10,000) for household i

$$Deductible_i = \alpha_{jt} + \beta\ PropertyValue_i + \varepsilon_i$$

where α_{jt} are $County \times Year$ fixed effects. We find a statistically significant, negative (but perhaps not economically meaningful) relationship between property values and deductible choices, an odds ratio of $\beta = 0.986$ ($s.e. = 0.0013,\ z = -10.22,\ p < 0.01$).

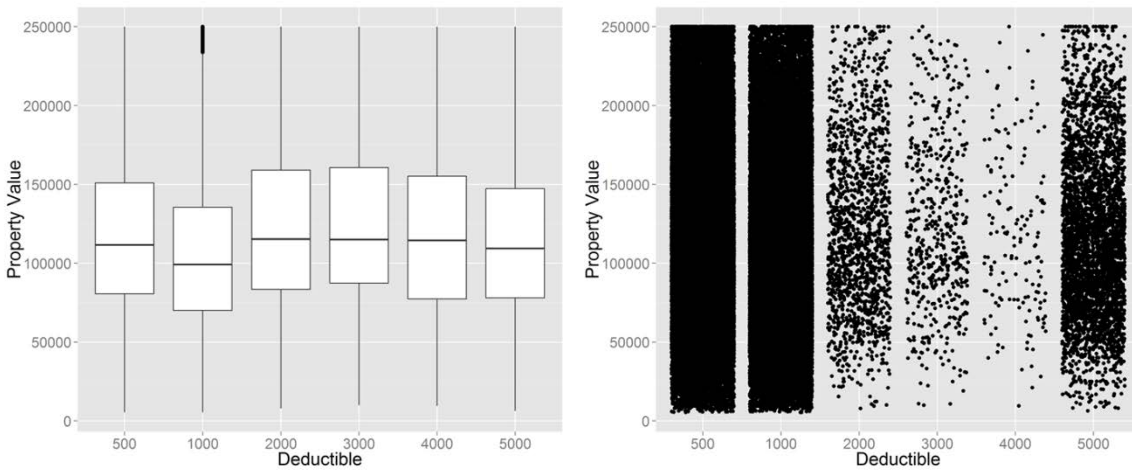

Figure 6 Boxplot and Jitter Plot of Households' Deductible Choices and Property Values

Note: The left figure provides a boxplot of household deductible choices and property values. The solid line in the center of each box shows the median property value, the upper and lower edges of the box provide 25th and 75th percentiles, and the solid lines extend to 1.5 times the interquartile range, the distance between the 25th and 75th percentiles. Dots outside of these lines identify individual outliers. The right figure shows individual observations of deductible choices and property values. Baseline sample of 103,080 policies.

A.3.3 Mortgages and Strategic Non-repayment

The differences in households' preferences across low and high stakes might be explained by their use of mortgages: homeowners may be unconcerned about losses beyond their equity. We explore this possibility by exploiting differences in lenders' recourse rights for defaulted mortgages across states. If households are insuring strategically based on their limited liability for damage to the home, we would expect them to select lower coverage limits relative to their replacement costs in "non-recourse" states.[31] We find a small difference in recourse and non-recourse states, as shown in Figure 7. For example, 75 percent of households in non-recourse states fully insure compared to 78 percent of households in recourse states.

[31] Non-recourse laws prevent lenders from seeking a deficiency judgment when the sale from a foreclosure is insufficient to cover the homeowner's liability on a defaulted mortgage. Orlando (2011) identifies Alaska, Arizona, California, Hawaii, Minnesota, Montana, North Dakota, Oklahoma, Oregon, and Washington as non-recourse states during the same time period as our data.

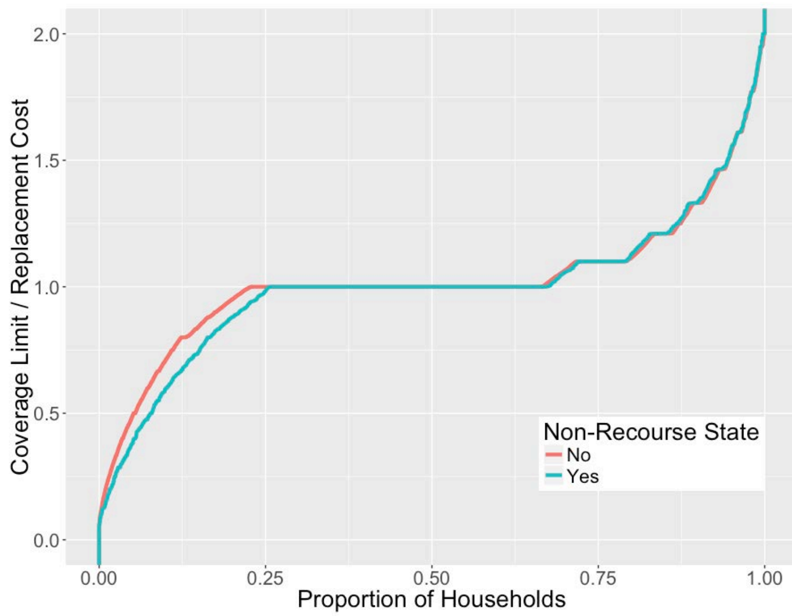

Figure 7 Coverage Limits Selected in Recourse and Non-Recourse States

Note: The figure compares coverage limit selections for households that live "non-recourse" states versus other states. Orlando (2011) identifies Alaska, Arizona, California, Hawaii, Minnesota, Montana, North Dakota, Oklahoma, Oregon, and Washington as non-recourse states during the same time period as our data.

A.3.4 Program Requirements

Households are required to insure against flood if they have a federally backed mortgage and live in a zone in which FEMA estimates of annual flood risk probabilities exceed one percent. While this mortgage requirement is not consistently enforced (Dixon et al., 2006), contracts would not reflect household preferences if, say, lenders determined coverage limits (e.g., required that households fully insure their homes).

We compare coverage limit selections for our baseline sample, which are subject to the federal requirement to insure, with households that participate in the program but live in a zone not subject to the federal requirement. Figure 8 compares the distribution of coverage limits for our core sample (labeled "Federal Requirement") and the less vulnerable group (labeled "No Requirement"). The similar coverage limit distribution for these two groups suggests that the federal requirement has little effect on the coverage limit choices that insured households select.

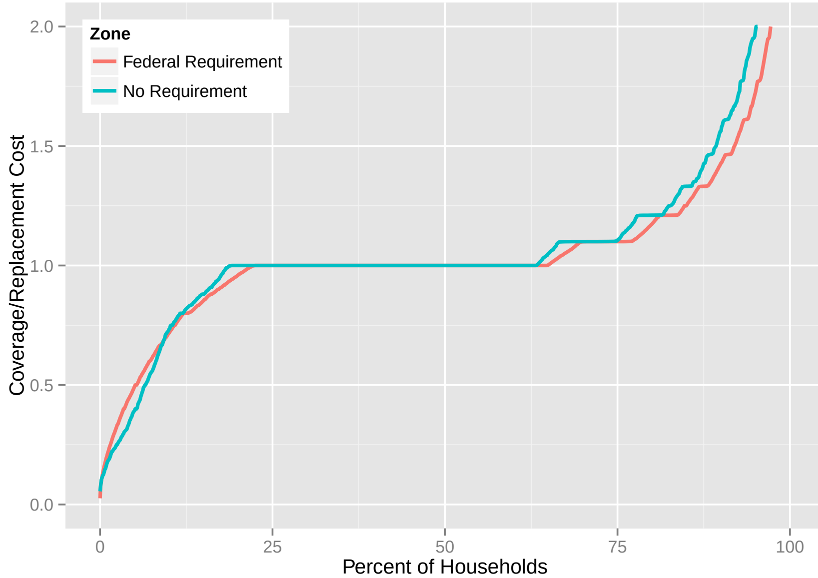

Figure 8 Coverage Limits Selected in Zones Where Insurance is Sometimes Required vs. Voluntary

Note: The figure compares coverage limit selections for households that live in a zone in which the federal government requires that households with federally backed mortgages insure with a zone in which households are not required to insure.

A.3.5 Household Wealth

We analyze the sensitivity of our findings to our wealth assumptions using a random sample of 10,000 policies from the baseline sample. Our results in Table 5 assume that households' wealth equals their replacement cost. Perhaps households are considering a much lower or higher level of wealth than what we assume, which explains the divergence in risk preferences across low and high stakes decisions that we observe. Also, while not typical of classical utility theory, the reference that households use for wealth might change across contexts, and so this analysis shows how reference wealth would need to differ to reconcile households' deductible and coverage limit preferences.

We set wealth equal to varying proportions of the replacement cost RC, ranging from 5 percent to 10 times its value, using scalar α. To illustrate, we rewrite the coverage limit problem (Equation 4), replacing wealth w with αRC

$$\max_{c\in[\underline{c},RC]} \int_0^{d^*} u(\alpha RC - p(c) - l)\omega(l) + \int_{d^*}^{c} V(\alpha RC - p(c) - d^*)\omega(l) + \int_{c}^{RC} V(\alpha RC - p(c) - d^* - l + c)\omega(l) \quad dl$$

where d^* is the deductible that the household selects, p the premium, l losses with $l \in [0, RC]$, and $\omega(l)$ the household's transformation of loss probabilities. As before, we set a lower bound so that wealth never falls below \$1. The average replacement cost is \$107,177 in our data, and so similar to Sydnor (2010), these wealth assumptions range from an average of about \$5,000 to \$1 million across values of α. Rather than the replacement cost, the final column uses the home's property as reference wealth (average value: \$112,387). Property value and replacement cost have a Pearson correlation of 0.58.

Table 14 shows the results. For the deductible, relative risk aversion follows the expected pattern: for a given insurance contract, estimates of risk aversion increase in household wealth. For the coverage limit, we observe that risk aversion decreases in wealth until around 50 percent of the replacement cost, then increases. The increase in risk aversion for low values of α results from households selecting large coverage limits relative to their assumed wealth.

For each wealth assumption, we continue to find different risk attitudes for the deductible and coverage limit decisions. The deductible and coverage limit risk aversion estimates diverge as reference wealth grows. The table also shows that deductible and coverage limit preferences could be reconciled by allowing households to select different reference wealth for each decision. For example, a reference wealth of about 5 percent of the replacement cost in the deductible decision and around 200 percent of the replacement cost in the coverage limit decision would seem to result in similar relative risk aversion coefficients. While different reference wealth values could help overcome the differences in small and large stakes preferences that we observe, the expected utility models for coverage limits continue to under-predict households' coverage limit selections. For example, the CRRA expected utility model with $\rho = 9.33$, $\sigma = 0.149$, and wealth as 200 percent of the home's replacement cost predict an average coverage limit ratio of 70 percent while the actual average is 93 percent.

Finally, using property value as reference wealth results in similar relative risk aversion estimates as using replacement cost (with $\alpha = 1$ in both cases), though the relative risk aversion from the coverage limit is slightly higher (3.7 versus 2.6).

Table 14. Relative Risk Aversion Estimates Across Wealth Assumptions

		Deductible					Coverage Limit				
	α	ρ	Std. Error	σ	Std. Error	Log Likelihood	ρ	Std. Error	σ	Std. Error	Log Likelihood
Replacement Cost	0.05	7.56	0.229	0.399	0.007	-15,199	1.42	0.002	0.121	0.002	-26,578
	0.10	15.4	0.303	0.355	0.008	-15,719	1.29	0.002	0.11	0.001	-25,358
	0.25	38.1	0.564	0.334	0.008	-15,951	1.16	0.002	0.145	0.002	-25,654
	0.50	62.5	0.206	0.372	0.009	-16,242	1.12	0.003	0.206	0.003	-27,546
	1	145	2.193	0.368	0.010	-16,368	2.62	0.035	0.230	0.007	-29,095
	2	278	4.134	0.355	0.009	-16,280	9.33	0.057	0.149	0.003	-28,276
	5	689	10.347	0.346	0.009	-16,222	26.8	0.163	0.147	0.003	-28,103
	10	1,378	19.835	0.343	0.009	-16,204	55.7	0.322	0.146	0.003	-28,057
Property Value	1	141	0.77	0.31-	0.008	-16,120	3.71	0.057	0.312	0.006	-30,194

Note: Table shows estimated relative risk aversion for households' deductible and coverage limit decisions using CRRA expected utility models. In the rows labeled "Replacement Cost" (the first eight rows), α is multiplied by the replacement cost so that, for example, the first row shows a risk aversion of $\rho = 7.56$ from the deductible decision when household wealth is assumed to be 5 percent of the replacement cost (around $5,000 in these data). The final row assumes that household wealth is the home's property value rather than the replacement cost. Random sample of 10,000 policies from the baseline sample.

A.3.6 Deductible Menu Defaults

Households are influenced by program default options in many contexts (e.g., Madrian and Shea, 2001). We re-estimate our models excluding households who adopted the default option. Table 15 provides the results, labeled "Active Choosers." As we are omitting households based on their deductible choices, we observe larger differences in the parameter estimates from households' deductible decisions. For example, we estimate relative risk aversion of 108 for these households versus 139 for our baseline sample.

A.3.7 Misunderstandings About the Insurance Coverage Limit

A large group of households (35 percent) purchased coverage limits that were higher than their replacement costs. Perhaps these households did not understand the contract that they were purchasing, thinking that they could receive larger payments than the replacement cost. Given this possibility, we re-estimate our model excluding these households. Table 15 provides the results, labeled "No Over-Insurers."

Table 15 Estimates for Households That Did Not Select Their Default Deductible and Households That Did Not Over-Insure

		ρ	σ	Log Likelihood	AIC	BIC
Active Choosers						
EUT CRRA	Deductible	108	2.78	-43,169	43,173	86,360
		(3.46)	(0.153)			
	Coverage Limit	2.56	0.188	-69,907	69,911	139,836
		(0.0115)	(0.00231)			
No Over-Insurers						
EUT CRRA	Deductible	147	0.44	-110,901	110,905	221,824
		(0.807)	(0.00461)			
	Coverage Limit	2.48	0.159	-184,929	184,933	369,880
		(0.00638)	(0.0011)			

Note: See Section A.3.6 and Section A.3.7 for details. Number of observations: 24,193 (“Active Choosers”) and 66,732 (“No Over-Insurers”).

A.3.8 Events with Small Probabilities

We also test our main results in an area with a higher probability of flooding. Our data also include a population that is vulnerable to inundation as well as wave damage (dwellings located directly on the coast, in FEMA-defined V zones) and their average annual claim rate is 4.1 percent, over three times higher than that of our core sample. We analyze the coverage limit selections for 7,113 households in this higher risk zone. Premiums are much larger for these households than those in our baseline sample, with a mean of $1,497 (and median of $1,201) compared to $540 for the baseline (Table 1). Households in this zone choose higher deductibles: 78 percent choose one of the lowest two deductibles (versus 94 percent in the baseline sample) and 12 percent choose the highest deductible of $5,000 (versus 3.8 percent in the baseline sample). Eighty percent of these households fully insure, selecting coverage limit of at least their replacement cost. Table 16 provides the results. We find qualitatively similar risk aversion estimates in the CRRA expected utility models to those in our baseline sample. Rank dependent utility models continue to show overweighting of small probabilities and diminishing sensitivity to losses.

Table 16 Parameter Estimates Across Models for High Risk Zone

		ρ	β_0	β_1	σ	Log Likelihood	AIC	BIC
EUT CRRA	Deductible	139			1.24	-12,525	12,529	25,068
		(3.3)			(0.064)			
	Coverage Limit	3.58			0.135	-21,084	21,088	42,186
		(0.035)			(0.0047)			
RDU CRRA	Deductible	-10.2	0.087	4.2	0.401	-10,318	10,326	20,671
		(0.24)	(0.0015)	(0.042)	(0.0065)			
	Coverage Limit	-0.681	0.0382	0.778	0.127	-15,486	15,494	31,007
		(0.026)	(0.0004)	(0.12)	(0.0017)			

Note: Standard errors are in parentheses. Table compares utility models describing households' deductible and coverage limit decisions, comparing outcomes using the Akaike information criterion (AIC) and the Bayesian information criterion (BIC). For the AIC and BIC, lower values indicate better fit. The models use maximum likelihood estimation to fit model parameters. The parameter σ describes the scale of model errors, following Equation 6. Parameters $\boldsymbol{\beta}$ show probability distortions $\Omega = \beta_0 + \beta_1 \Pi$ where Π is the cumulative objective probability. Expected utility theory (EUT) assumes that households weight outcomes based on their objective probabilities and so these models fix $\beta_0 = 0$ and $\beta_1 = 1$. We allow for probability distortions using rank dependent utility (RDU, Quiggin, 1982). Positive values for β_0 and $\beta_1 > 1$ indicate overweighting of the small probabilities in our data. Negative values for the coefficient of relative risk aversion ρ indicate diminishing sensitivity to losses. Number of observations: 7,113.