Comparing the Effects of Behaviorally-Informed Interventions on Flood Insurance Demand: An Experimental Analysis of 'Boosts' and 'Nudges'*

Jacob Bradt†

Draft accepted at *Behavioural Public Policy*

July 9, 2019

Abstract

This paper compares the effects of two types of behaviorally-informed policy, nudges and boosts, that are designed to increase consumer demand for insurance against low-probability, high-consequence (LPHC) events. Using previous findings in the behavioral sciences literature, this paper constructs and implements two nudges (an "informational" and an "affective" nudge) and a statistical numeracy boost and then elicits individual risk beliefs and demand for flood insurance using a contingent valuation survey of 331 participants recruited from an online labor pool. Using a two-limit Tobit model to estimate willingness-to-pay (WTP) for flood insurance, this paper finds that the affective and informational nudges result in increases in WTP for flood insurance of roughly \$21/month and \$11/month relative to the boost, respectively. Taken together, the findings of this paper suggest that nudges are the more effective behaviorally-informed policy in this setting, particularly when the nudge design targets the affect and availability heuristics; however, additional research is necessary to establish sufficient conditions for this conclusion.

Keywords: Nudge, Boost, Flood Risk, Insurance, Contingent Valuation

JEL Codes: D81, D91, G22, Q54

*I would like to thank Cass Sunstein and an anonymous reviewer for their valuable feedback on earlier drafts of this paper. I am grateful to Arevik Avedian for her helpful comments on the survey instrument as well as the students in the Behavioral Economics, Law, and Public Policy Seminar at Harvard Law School for their engaging discussion on this topic. The author is solely responsible for any errors contained herein. Data and code to replicate the analysis in this paper are available from the author on request.

†John F. Kennedy School of Government, Harvard University; 79 John F. Kennedy St., Cambridge, MA, 02138. Email: `jbradt@g.harvard.edu`. Tel.: (301) 655-5262.

Introduction

Risk is pervasive: individuals are often required to make decisions without knowing which possible outcome will occur.[1] These decisions range from the mundane to the extraordinary. From the decision to pack an umbrella based on a weather forecast to the choice of occupation—countless decisions that individuals make are done without knowing which of the possible outcomes will actually occur. A particular class of decisions in which risk plays an important role is consumer demand for insurance against low-probability, high-consequence (LPHC) events, such as natural disasters.

Numerous studies in both the cognitive psychology and behavioral science literature examine how individuals interact with risk information in LPHC settings. In general, the literature finds that when making decisions associated with risk, individuals tend to either neglect or overweigh low-probability risks (Slovic et al., 1977; Lichtenstein et al., 1978). Moreover, individuals' tendency to do so can be correlated with various factors such as their emotional state; the framing of the risk and outcomes; and past experiences (Johnson et al., 1993; Kunreuther et al., 2001; Thaler et al., 1997; Tom et al., 2007; Browne et al., 2015).

A particular policy setting in which this behavior is observed is the market for flood insurance: experiments have confirmed that many individuals neglect low-probability flood risks and do not purchase insurance, while others reveal a willingness-to-pay (WTP) for flood insurance that exceeds the loss' expected value. This observed behavior violates common rational agent-based theories of decision-making under risk that assume individuals make decisions—such as the amount they are willing to pay for insurance—by maximizing their expected utility or payoff. Empirical evidence validates this behavior in the larger population (Botzen and van den Bergh, 2012; Botzen et al., 2013). Kunreuther (2018) describes cognitive biases that may explain the neglect of low probability flood risk in the context of the United

[1]This paper adopts the distinction between risk and uncertainty commonly found in the decision theory literature: uncertainty (i.e., "Knightian uncertainty") implies the inability to assign a probability distribution to the set of possible outcomes of a decision, whereas risk implies a well-defined probability distribution for the set of possible outcomes of a decision (Knight, 1921).

States (U.S.) flood insurance market: myopia, amnesia, optimism, inertia, simplification, and herding. The literature suggests that the use of the availability and affect heuristics—mental shortcuts that rely on immediate examples or emotional salience, respectively, to evaluate in this case the likelihood of flooding—may also lead certain individuals to inflate their perceptions of the risk of these events (Keller et al., 2006).

In light of these findings relating to consumer behavior, this paper seeks to examine the relative effect of different behaviorally-informed interventions on the demand for insurance against LPHC events. Interest in behavioral policy interventions has increased in recent years (Oliver, 2013; Shafir, 2013; Chetty, 2015). These policies assume many different forms; however, a helpful taxonomy within the set of non-incentivizing, choice-preserving behavioral policies distinguishes between "nudges" and "boosts" (Grüne-Yanoff et al., 2018). Generally, nudges are defined as changes to a decision frame—the manner in which a decision is presented—that alter individuals' behavior in predictable ways without excluding options or altering the incentive structure (Thaler and Sunstein, 2008). Boosts are interventions characterized by their goal of expanding the decision-maker's set of competences in order to enable them to accomplish their objectives (Grüne-Yanoff and Hertwig, 2016).

Focusing in particular on the context of demand for flood insurance, this paper seeks to compare the effects of nudges and boosts designed to increase consumer demand for insurance against LPHC events. To do so, this paper elicits WTP for flood insurance in a hypothetical scenario involving coastal flood risk using a web-based survey in which 331 participants recruited from an online labor pool are randomly assigned to receive either a nudge, a boost, or neither. Given the large literature suggesting that affect and availability heuristics impact agent decision-making in these settings, two nudges are included in the study: one purely informational (henceforth, "informational nudge"), the other adding affective and availability cues to the informational nudge (henceforth, "affective nudge"). Whereas in the control condition the risk of flooding is presented as an annual probability, the informational nudge alters the presentation of the flood risk by translating this annual probability into the probability of experiencing a flood over a 30-year period. The affective nudge adds information about coastal flooding in the U.S. to the informational nudge's decision frame, to include information on current and future exposure of at-risk population and assets and a reference to a particularly damaging U.S. hurricane season. Lastly, based on a large literature on statistical training, the boost condition is designed to provide individuals with an intuitive heuristic applying the Law of Large Numbers, a theorem that states that under general conditions, the sample average will be close to the population average with very high probability when the sample size is large, as is the case when considering the probability of flooding over a large sample of homes.

Ultimately, the study finds that nudges are the more effective behaviorally-informed policy tool in this setting. In particular, using a two-limit Tobit model to estimate WTP for flood insurance, this paper finds that the affective and informational nudges result in increases in WTP for flood insurance of roughly \$21/month and \$11/month relative to the boost, respectively. Moreover, this paper finds that the efficacy of a nudge relative to the baseline condition of no behavioral intervention is contingent on targeting the affect and availability heuristics. These findings suggest that policymakers may prefer nudges over boosts in the context of flood insurance or insurance against other LPHC events. Ultimately, this study provides two key takeaways for policymakers interested in increasing take-up of insurance against LPHC events: behavioural policies can play a large role in accomplishing this objective and particular attention should be given to the framing of the risk in question. The next section discusses the differences between nudges and boosts and introduces relevant literature.

Background

Nudges and boosts are associated with two distinct programs in the behavioral sciences: the former is viewed as a product of the Heuristics and Biases program (Tversky and Kahneman, 1986; Kahneman and Tversky, 1996); the latter the Fast and Frugal Heuristics program (Gigerenzer and Todd, 1999). While these two forms of behavioral policy are the result of different strains of research, they are similar in that they both seek to alter agents' behavior without substantially changing incentives or restricting agents' autonomy through legal mandates. Moreover, nudges and boosts both assume that individuals use a finite set of heuristics—i.e., mental processes or "shortcuts"—to make decisions and that the result of the use of these heuristics depends on properties of the decision frame (Thaler and Sunstein, 2008; Gigerenzer and Todd, 1999; Grüne-Yanoff and Hertwig, 2016; Grüne-Yanoff et al., 2018).

Much of the literature discussing the differences between nudges and boosts focuses on the ethical implications of each policy type (Grüne-Yanoff et al., 2018). Some contend that any effort to draw a normative distinction between these two forms of policy is ill-founded (Sims and Müller, 2018). This paper brackets the ethical implications of nudging and boosting and focuses instead on evaluating the different effects of these two policy types in a specific setting.

While there are those who argue that the distinction between nudges and educative boosts is perhaps tenuous, particularly in the case of informational nudges (Sunstein, 2015), this paper adopts the view that this distinction is well-founded on mechanical grounds. In particular, nudges and boosts differ in both (1) the point in the cognitive decision-making process which they target and (2) the mechanism that they employ to accomplish their desired outcome. Assuming that agents' set of heuristics is stable, nudges target the decision frame as a means of using cognitive biases to accomplish an outcome. In comparison, while boosts assume that agents make use of cognitive heuristics, they do not assume that the set of heuristics is fixed. Boosts therefore target agents' supply of cognitive heuristics directly, seeking to either improve existing heuristics or provide decision-makers with heuristics to apply to the decision in question (Grüne-Yanoff et al., 2018).

While this is possibly an imperfect classification, the purpose of this paper is not to further clarify the distinction between nudges and boosts. Given that the literature has largely adopted a distinction between nudges and boosts, this paper seeks to provide evidence of the relative effect of each form of policy in a specific setting. Though there are numerous areas in which nudges and boosts can be and are currently employed, the analysis herein focuses on the effect of a nudge and a boost in the context of decision-making with respect to LPHC events given that there exists a robust behavioral science literature on this topic.

LPHC Events: Behavioral Biases and Nudges

The behavioral science literature has established numerous facts regarding agents' decision-making under risk.[2] Several studies have found that individuals tend to overestimate small probabilities and underestimate large probabilities (Lichtenstein et al., 1978; Hertwig et al., 2004). Moreover, individuals' perceptions of the likelihood of an event appear to be driven in part by the exposure, memorability, or imaginability of the event for that individual (Slovic, 2000; Hertwig et al., 2004), which suggests that individuals make use of the availability and affect heuristics when making decisions regarding low-probability outcomes (Tversky and Kahneman, 1973; Finucane et al., 2000). Overall, when asked to make decisions under conditions of risk, the evidence suggests that individuals depart from conventional models of expectational reasoning (Johnson et al., 1993; Kunreuther et al., 2002).

The field of regulatory focus theory builds on findings in the cognitive psychology literature, emphasizing the role that motives play in individual decision-making. In particular, regulatory focus theory contends that two motivational systems shape decision-making: the promotion and prevention systems (Higgins, 1998). In general, promotion motivations are those that seek to improve on the status quo, whereas prevention motivations are those that seek to maintain the status quo by avoiding loss. Krantz and Kunreuther (1998) emphasize the role of motivational reasoning in decision-making under risk, particularly in the context of LPHC events such as the decision to purchase flood insurance. Botzen et al. (2013) demon-

[2]For a particularly helpful and rather comprehensive survey of probabilistic reasoning and judgement biases, see Benjamin (2018).

strate that both situational motivation—motivation directly related to a specific context—and global motivation—motivation that applies to all contexts—play an important role in driving individual decision-making in the context of flood insurance.

It appears as though when making decisions regarding low-probability events, agents have a threshold probability value, below which they neglect a given outcome (Slovic et al., 1977; McClelland et al., 1993; Schade et al., 2012). Examining motorists' seatbelt use behavior, Slovic et al. (1978) find that reporting the probability of serious injury over the course of a lifetime of automobile trips rather than over a single trip leads more people to wear seat-belts. Browne et al. (2015) find evidence to suggest that individuals neglect low-probability outcomes and overweigh high-probability outcomes in the market for insurance. A possible reason for this behavior is that individuals are unable to internalize probabilities below a certain threshold as they do not believe that such unlikely events could ever happen to them, a phenomenon known as probability neglect (Tversky and Shafir, 1992; Sunstein, 2002). In short, nudges have been shown to affect individual decision-making with respect to LPHC events.

Boosting Statistical Reasoning

Numerous studies examine the effects of training, or boosting, on biased decision-making. In particular, there exists a robust literature that seeks to use the cognitive model of heuristics to address well-defined biases by either providing agents with new heuristics or enhancing the application of existing heuristics. Several studies find that brief training in formal, inferential rules may enhance agents' use of statistical reasoning—their understanding of and ability to apply statistical concepts—in everyday life (Nisbett et al., 1987). Evidence suggests that providing individuals with an intuitive heuristic applying the Law of Large Numbers improves their statistical reasoning (Fong et al., 1986; Fong and Nisbett, 1991).

Other efforts to measure the effects of training on statistical reasoning have sought to examine the impact of interventions on agents' Bayesian reasoning, which refers to individuals' ability to use new information to update their prior beliefs about the likelihood of an event. A key result of this literature is that providing individuals with probabilities in frequency format as opposed to probability format improves their ability to conduct Bayesian reasoning (Gigerenzer and Hoffrage, 1995; Gigerenzer, 1996, 2014). The reason for this improvement is argued to lie in the mechanics of solving these types of problems: it is argued that Bayesian computations are simpler to perform with natural frequencies than with probabilities (Sedlmeier and Gigerenzer, 2001).

These findings have led some to contend that training or boosting decision-making can produce persistent reductions in cognitive biases (Gigerenzer and Brighton, 2009; Morewedge et al., 2015). In the context of agents' behavior with respect to LPHC events, this literature suggests that boosting individuals' expectational reasoning enhances the biased decision-making discussed in the previous subsection (Slovic et al., 2002; Kunreuther et al., 2002). This paper is not aware of previous studies examining the effect of boosts on decisions to insure against LPHC events, not to mention comparing its effect relative to nudges in this setting.

Flood Insurance Demand

The market for insurance against flooding represents a particularly important and policy-relevant area in which individuals' departure from rational agent-based models of behavior in the context of LPHC events is observed. In the U.S., the take-up rate (i.e., the proportion of households purchasing insurance) for flood insurance in high risk areas remains around 49%, despite mandates requiring the purchase of coverage (Kousky, 2018; Kunreuther, 2018). Moreover, several revealed preference studies of insurance demand in the U.S. indicate that many homeowners do not internalize flood risk and fail to purchase flood insurance, even when it is partly subsidized (Kunreuther and Slovic, 1978; Atreya et al., 2015).

While it is possible that the low take-up rates for flood insurance observed in practice represent unbiased preferences in the population, this is unlikely. Several studies have found that WTP for flood insurance is conditional on a number of factors which are suggestive of biased behavior in these markets. Examining the market for flood insurance in the state of Georgia, Atreya et al. (2015) find that experience with recent flood events temporarily increases purchases of policies. Surveys of coastal residents in the Netherlands suggest that the framing of risk as well past experience with flooding alter WTP for insurance (Botzen and van den Bergh, 2012; Botzen et al., 2013; de Boer et al., 2015).

Hypotheses and Survey Methodology

Despite the presence of financial incentives as well as a—albeit, poorly-enforced—mandate requiring the purchase of flood insurance, relatively few at-risk households in the U.S. do so. While there are a number of non-behaviorally-informed policy remedies that are likely necessary to address this issue (see e.g., Kunreuther (2018)), the poor performance of incentives and mandates to increase the purchase of flood insurance coupled with the many applicable insights available from the behavioral sciences literature begs the consideration of possible behavioral policies in this area. With this in mind, this paper seeks to provide early evidence of the relative performance of two such types of policy, nudges and boosts.

Hypotheses

To compare the relative effects of nudges and boosts in this setting, a survey is constructed in which respondents are randomly assigned to one of a Control condition, a boost (Treatment 1), an informational nudge (Treatment 2a), or an affective nudge (Treatment 2b) and then asked to consider their WTP for insuring a hypothetical endowment against a given annual probability of flooding. The differences between the two forms of a nudge (Treatments 2a and 2b) as well as the boost (Treatment 1) and Control conditions are outlined in the subsections that follow and in Table 1.

Below are the set of hypotheses (H) based on the relevant literature which the survey instrument seeks to test.

- **H1**: There will be a large number of respondents who will not purchase insurance independent of treatment status (i.e., in all treatment and control groups).
- **H2**: Providing respondents with a brief statistical reasoning boost intended to give them an intuitive heuristic applying the Law of Large Numbers will have a relatively minor effect on their WTP for insurance.
- **H3**: Obtaining a non-zero treatment effect of a behaviorally-informed nudge on WTP for flood insurance is conditional on targeting the availability/affect heuristics.

The above hypotheses are informed by previous findings in the literature. In particular, **H1** is informed by the finding that, all else equal, individuals often fail to purchase insurance against LPHC events in both experimental and empirical settings. **H2** is based on the numerous findings in the literature that suggest that certain well-documented biases affect individuals' decision to forgo insurance in LPHC settings. This paper contends that the cognitive biases that affect individuals' probabilistic reasoning in this setting are sufficiently strong to ensure that any attempt to increase take-up and WTP without making explicit use of these biases will be unsuccessful. Lastly, **H3** is informed by the finding that affect and availability play a significant role in determining individuals' insurance decisions with respect to LPHC events. The subsections that follow describe the survey instrument used to test the above hypotheses.

[**Table 1 about here.**]

Eliciting WTP for Flood Insurance

Following assignment of treatment status, respondents are presented with a hypothetical scenario in which they are informed that they own a single-family, detached home located in the coastal U.S. with a total value of \$300,000.[3] Respondents are told that their home faces a 1% probability in any given year of experiencing flooding resulting in approximately \$75,000 worth of damages.[4] Thus, the expected damages to the endowed home in the hypothetical scenario are (0.01) x \$75,000 = \$750/year or \$62.50/month.

After being asked to imagine this hypothetical scenario, respondents are informed that there exists an annual insurance policy that will cover the cost of damages associated with the flooding risk described. Respondents are then asked to consider this policy and indicate their prior belief of the cost of the policy on a sliding scale ranging from \$0-\$125/month. This question is included to prime the respondents to consider the actual cost of such a policy before eliciting their willingness-to-pay for coverage. Interest in purchasing an annual flood insurance policy for the home described in the hypothetical scenario is then elicited. If the respondents answer "Yes" or "Maybe" then they are asked to indicate the "highest amount" that they are willing to pay to purchase the annual insurance policy described in the hypothetical scenario on a sliding scale ranging from \$0-\$125/month. If the respondents answer "No," they are assigned a WTP value of \$0/month. The use of this elicitation format introduces methodological challenges, which are discussed and addressed below.

The different treatment and control conditions vary in their presentation of the hypothetical scenario (see Appendix A). In particular, the descriptions of the hypothetical scenario across treatment and control conditions vary in their presentation of the risk of flooding (holding the description of the loss associated with flooding constant). In the Control condition, respondents are presented with the flooding risk in probabilistic terms (i.e., "there is a 1% chance, in any given year, that you will experience flooding…"). This risk presentation is chosen as the baseline as this aligns closely with the format in which risk is presented to

[3]This figure is within the 10-year range of the average quarterly seasonally-adjusted sale price time series for homes sold in the U.S. from 2008 to 2018 (sources: U.S. Census Bureau and U.S. Department of Housing and Urban Development, 2019).

[4]These numbers are selected due to their relative ease of use in mental arithmetic as well as the similarity to the policy environment of flood insurance in the U.S.: communication of incidence within a 100-year floodplain— defined as those areas with a 1% annual probability of inundation—is one of the primary informational mechanisms available to homeowners when assessing their risk exposure.

homeowners in practice (see Footnote 4).

Nudge and Boost Design

Prior to being presented with the hypothetical scenario, respondents assigned to Treatment 1 are provided with a brief description of how to interpret probabilities in frequency terms (see Appendix A). This includes a description of a simple example as well as an easy to use inferential rule to interpret probabilities in frequency terms. This boost is designed based on the relevant literature on statistical training: the objective of the boost is to provide individuals with an intuitive heuristic applying the Law of Large Numbers (Fong et al., 1986; Nisbett et al., 1987; Fong and Nisbett, 1991). Respondents in this treatment group are then shown the hypothetical scenario in which flooding risk is presented in both probabilistic and relative frequency terms (i.e., "there is a 1% (or 1 in 100) chance, in any given year, that you will experience flooding…"). This risk presentation is intended to improve respondents' Bayesian reasoning (Gigerenzer and Hoffrage, 1995; Gigerenzer, 1996; Sedlmeier and Gigerenzer, 2001).

Similar to the Control condition, respondents assigned to Treatment 2a are shown the hypothetical scenario prior to eliciting their WTP for insurance and are presented the flooding risk in probability terms. However, in addition to presenting the risk of flooding on an annual basis, respondents in Treatment 2a observe the probability of inundation over a 30-year period (i.e., "…over the course of 30 years of home ownership, the probability of experiencing a severe flood of this type is approximately 26%"), which corresponds to the modal amortization period for fixed-rate mortgages in the U.S.[5]

Prior to being presented with the hypothetical scenario, respondents in Treatment 2b are provided with information on coastal flooding in the U.S., to include: a description of at-risk populations and assets; a summary of coastal flooding projections resulting from climate change; and a description of damages resulting from salient tropical cyclones that

[5]Assuming a static, 1% annual probability of flooding over the course of 30 years, the probability of experiencing a flood in 30 years is $1 - (1 - 0.01)^{30} \approx 0.26$.

made landfall in the U.S. during the 2017 hurricane season. Respondents in this group are then shown a set of images of flooding in coastal areas (see Appendix A). This intervention is intended to elicit an affective response and prime respondents to consider salient, recent examples of coastal flooding. Respondents in this treatment group are then presented with the same hypothetical scenario and risk presentation as in Treatment 2a.

Motivation and Prior Beliefs

Based on the regulatory focus theory literature (Higgins, 1998), respondents are asked to consider a set of questions designed to measure the situational and global motivations behind their stated WTP value similar to Botzen et al. (2013) and de Boer et al. (2014). Two questions are specifically designed to measure the degree to which respondents' maximum WTP for flood insurance is motivated by the risk described in the frame ("situational" motivation) and another two questions are designed to measure respondents' general motivations to insure against risk ("global" motivation). These questions are designed to target both the prevention and promotion systems in both the situational and global cases. For a full description of the motivation elicitation questions, see Appendix B.

Respondents' prior beliefs regarding climate change and sea-level rise are also elicited in order to control for strong or weak priors on future flooding in coastal areas. Given the possible role of previous exposure to flooding or similar natural disasters on insurance demand, respondents prior experience with these events is also elicited. Additional demographic variables are elicited, including age, education, income, home state, and home ownership status. A complete description of variables elicited via the survey instrument is given in Table 2.

[Table 2 about here.]

Empirical Methods

A single model is estimated for WTP elicited via the survey instrument. Due to the method employed to elicit respondents' WTP, a two-limit Tobit model is used to estimate the marginal effects of treatment on individuals' WTP for flood insurance.

Linear WTP Model

It is assumed that, for each respondent $i \in \{1, \ldots, N\}$, WTP is a linear function of the form

$$WTP_i = f(X_i, D_i, P_i) + \epsilon_i \tag{1}$$

where $f(\cdot)$ is a linear function; X_i is a $1 \times j$ vector of j individual attributes; D_i is the value of property exposed to flood damage; P_i is the probability of inundation; and $\epsilon_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is a normally-distributed error term with mean zero and variance σ_ϵ^2. In particular, WTP is parameterized as the following:

$$WTP_i = \beta_0 + X_i\beta_1 + D_i\beta_2 + P_i\beta_3 + \epsilon_i \tag{2}$$

The literature finds that individuals may follow a belief updating pathway suggestive of Bayesian reasoning when they are provided with novel information regarding the likelihood of adverse outcomes (Viscusi, 1985; Botzen and van den Bergh, 2012). Assuming that respondents possess a prior belief regarding the likelihood of flooding in the hypothetical scenario, then WTP should be a function not of the flood probability stated in the survey, P_i, but rather of the posterior belief regarding flooding in the hypothetical scenario. This posterior belief is modelled as a function of respondents' prior beliefs, π_i, and the probability information presented in the hypothetical scenario:

$$\tilde{P}_i = \frac{\eta_1 \pi_i + \eta_2 P_i}{\eta_1 + \eta_2} \tag{3}$$

where η_1 and η_2 are the weights assigned to the prior belief about the probability of flooding in the hypothetical scenario, π_i, and the probability of the hypothetical endowment experiencing flooding that is given in the survey, P_i (Viscusi and O'Connor, 1984; Viscusi, 1985).

Given that numerous studies show that individuals often do not follow a Bayesian model of belief updating, largely due to the existence of certain behavioral biases (Kahneman et al., 1982; Cameron, 2005), Bayesian updating is not imposed on agents in this model. Similar to Botzen and van den Bergh (2012), the framework of Bayesian updating is introduced herein to incorporate the possibility of a meaningful relationship between WTP for flood insurance and individual flood risk perceptions. In fact, much of the behavioral sciences literature suggests that including factors that influence individual priors on flood risk (e.g., past experience of a flood) is important. Incorporating Eq. (3) into Eq. (2) gives

$$WTP_i = \beta_0 + X_i\beta_1 + D_i\beta_2 + \left(\frac{\eta_1 \pi_i + \eta_2 P_i}{\eta_1 + \eta_2}\right)\beta_3 + \epsilon_i \tag{4}$$

Respondents' prior beliefs regarding the probability of flooding in the hypothetical scenario presented in the survey are likely informed by their prior beliefs regarding flooding in coastal areas. These prior beliefs are not observed; however, similar to Botzen and van den Bergh (2012), a vector of variables, C_i, are used as proxy for π_i in Eq. (4):

$$WTP_i = \beta_0 + X_i\beta_1 + D_i\beta_2 + C_i\tilde{\beta}_3 + P_i\tilde{\beta}_4 + \epsilon_i \tag{5}$$

where

$$\tilde{\beta_3} = \frac{\eta_1}{\eta_1 + \eta_2}\beta_3 \qquad\qquad \tilde{\beta_4} = \frac{\eta_2}{\eta_1 + \phi_2}\beta_3 \tag{6}$$

Lastly, adding treatment assignment, T_i, to Eq. (5) gives the following model of WTP that is used to measure the effect of the different treatments:

$$WTP_i = \beta_0 + X_i\beta_1 + D_i\beta_2 + C_i\tilde{\beta_3} + P_i\tilde{\beta_4} + T_i\gamma + \epsilon_i \tag{7}$$

where T_i assumes the values $T_i \in \{0, 1, 2, 3\}$ corresponding to the Control condition, Treatment 1, Treatment 2a, and Treatment 2b. For a discussion of the two-limit Tobit model used to estimate WTP, see Appendix C.

Results

To test the above hypotheses, 331 participants are recruited to take the survey instrument using Amazon's Mechanical Turk in exchange for a modest payment. Of the 331 respondents recruited via Mechanical Turk, 10 are excluded from the final sample due to incomplete responses. An additional respondent is flagged and removed from the final sample due to invalid and inconsistent answers. The final sample used in the analysis therefore includes 320 total respondents.

Sample Characteristics

Summary statistics for the main explanatory variables elicited from the pool of Mechanical Turk respondents are given in Table 3.[6] Overall, the sample characteristics match those of

[6]Though there are several studies that suggest online surveys of the Mechanical Turk labor pool are externally valid when testing certain phenomena, this paper does not view the final sample as representative of the broader U.S. population (Berinsky et al., 2012; Clifford et al., 2015). However, this paper does view the findings of the analysis herein as suggestive of behavior in the broader population of interest, so comparisons are made to the U.S. population.

the broader U.S., up to a point.[7] The sample has slightly more male (57.81%) than female (41.56%) respondents. Respondents in the sample have more children and are more educated than the U.S. population: approximately 42.13% of the sample indicated that they have at least one child and 63.44% of the sample have a Bachelor's degree or higher. Overall, the sample skews towards younger individuals, with only 4.52% of the final sample over the age of 60. The median annual pre-tax household income in the sample is $62,500, which corresponds to the category "$50,001-$75,000" and is similar to the median household income in the U.S. in 2017.

[Table 3 about here.]

Descriptive Analysis of WTP

The results of the questions eliciting WTP for insurance against flooding of a hypothetical endowment are reported in Table 4. The proportion of respondents who indicate that they are willing-to-insure (WTI), the mean WTP of all respondents, and the mean WTP conditional on indicating positive interest (i.e., answering "yes" or "maybe") in purchasing insurance (CWTP) are shown for the full sample, Control condition, Treatment 1, Treatment 2a, and Treatment 2b.

Each of these values measures different components of respondents' insurance demand. The WTI measure captures the percentage of respondents who, when presented with information about the hypothetical flood risk, are willing to pay a non-zero amount for flood insurance. The mean WTP value for each group measures the utility respondents receive from the flood insurance policy described in the hypothetical scenario under each treatment status. Lastly, the CWTP value is computed as the mean WTP of respondents who are willing to pay a positive amount for flood insurance. This value is of interest as it captures the premium that respondents are willing to pay for the insurance product above or below the expected value of the loss described in the hypothetical scenario. In particular, for individual

[7] Sources for U.S. population data: U.S. Census Bureau (USCB), Current Population Survey Annual Social and Economic Suplements, 2018; USCB, American Community Survey, 5-year Estimates (2013-2017).

i:

$$\begin{aligned}\text{Risk Premium}_i &= CWTP_i - (\text{Flood Probability} \times \text{Flood Damages}) \\ &= CWTP_i - 62.50 \end{aligned} \tag{8}$$

The CWTP values reported in Table 4 average $CWTP_i$ values across all individuals in each group.

[Table 4 about here.]

WTI across treatment and control groups ranges from 81.40% to 98.55%. A cursory comparison of WTI values across the treatment and control groups reveals that, relative to the control subsample, the Treatment 1 subsample has a considerably lower (approximately 8 percentage points) WTI and the Treatment 2b subsample has a considerably higher (approximately 10 percentage points) WTI, whereas the Treatment 2a subsample WTI is roughly equal to that for the control subsample.

A similar pattern across treatment status emerges along the WTP and CWTP variables (see Figure 1). Ranging from $49.63/month to $71.46/month, mean WTP values for each of the treatment and control subsamples show that, relative to the Control group, individuals assigned to Treatment 1 have on average lower WTP values; individuals assigned to Treatment 2b have on average higher WTP values; and individuals assigned to Treatment 2a have on average roughly similar WTP values. Table 5 summarizes the results of a pairwise-comparison of the difference in mean WTP across treatment and control groups using Tukey's HSD.

[Figure 1 about here.]

Mean CWTP ranges from $60.97/month to $72.51/month across treatment and control groups. Using the CWTP results for each treatment and control group reported in Table 4 and Eq. (8), the risk premium across treatment and control subsamples ranges from -

\$1.53/month to \$10.01/month. In particular, whereas the mean risk premium for individuals in the Control group is estimated to be positive (\$2.89/month), the mean risk premium for individuals in Treatment 1 is estimated to be negative (-\$1.53/month), suggesting that these individuals either derive negative utility from the acquisition of the insurance policy described or are systematically underestimating their risk exposure. Individuals in the nudge treatments, Treatments 2a and 2b, have mean risk premiums greater than two and three times that of the Control group, respectively. This suggests that individuals assigned to these treatment groups either derive positive utility from the acquisition of insurance or are systematically overestimating their risk exposure.

[Table 5 about here.]

Estimation Results of Two-Limit Tobit

Inspection of the distribution of WTP results for the full sample suggests that the use of a two-limit Tobit is appropriate (see Figure 2 and Appendix C for a discussion). Table 6 reports the coefficients and standard errors of the two-limit Tobit model estimated by Newton-Raphson maximization as well as the resulting estimated marginal effects of a unit change in the explanatory variables on WTP.

[Figure 2 about here.]

Overall, point estimates for the parameter values have the expected valence, with several exceptions. The test for joint significance of all the covariates is the log-likelihood ratio, which is generated as the statistic $-2\log(L_R/L_U)$. In this case, the log-likelihood ratio is sufficient to reject the null hypothesis that all of the coefficients on the covariates in the model are equal to zero at the 1%-significance level ($p = 0.0003$). The log-likelihood ratio is also sufficient to reject the null hypothesis that all of the coefficients on the treatment status indicator variables are equal to zero at the 1%-significance level ($p = 0.001$).

[Table 6 about here.]

Assignment to several of the treatment conditions appear to have a significant effect on

WTP for flood insurance. In line with the findings of the descriptive analysis, the effect of assignment to Treatment 1 is estimated to result in a large decrease in WTP that is statistically-significant at the 90%-confidence level and the effect of assignment to Treatment 2b is estimated to result in a large increase in WTP that is statistically-significant at the 95%-confidence level. The estimated marginal effects for assignment to Treatments 1 and 2b are -\$10.74/month and \$12.50/month, respectively. The point estimate on the coefficient of the binary variable indicating assignment to Treatment 2a is positive; however, the size of the standard error on this term suggests that it is not possible to distinguish this parameter from zero at any common level of statistical significance. This is also in line with the findings of the descriptive analysis.

Past exposure to flooding and prior beliefs regarding climate change and sea-level rise also appear to have positive effects on WTP. In particular, past experience with a severe flood is estimated to have a large, statistically-significant effect on WTP for flood insurance. Though the point estimate for the coefficient on prior beliefs regarding climate change and sea-level rise is positive (i.e., stronger beliefs that climate change and resulting sea-level rise are occurring/will occur result in a higher WTP value), the estimated effect is small and not guaranteed to be non-zero. Of the remaining explanatory variables, few appear to have a large effect on WTP.

The Role of Situational and Global Motivation

Figure 3 shows the group means for the two situational and two global variables across the treatment and control groups. Overall, there does not appear to be significant variation in each of the motivation variables across treatment status, particularly in the case of the two global motivation variables. The two nudge treatments (Treatments 2a and 2b) do have slightly higher mean levels of worry and prevention motivation; however, it is unclear whether these differences across treatment groups are non-trivial.

[Figure 3 about here.]

To further examine the role of situational and global motivation, the four motivation variables are added to the two-limit Tobit model. Table 7 reports the coefficients and standard errors of the two-limit Tobit model including transformations of the motivation variables as z-scores as well as the resulting estimated marginal effects of a unit change in the covariates on WTP.

[Table 7 about here.]

Overall, the estimated model parameters match those estimated in the model in which motivation is not observed. All of the motivation variables are statistically significant at the 95%-confidence level or greater. The marginal effect of the financial security global motivation variable is large and negative, a somewhat counterintuitive result. As expected, the estimated marginal effect of the global safety variable is large in magnitude and positive. Moreover, the estimated marginal effect of the situational variables on WTP for insurance are both large in magnitude and positive. This strong relationship between situational motivation and insurance demand is shown in Figure 4.

[Figure 4 about here.]

Interestingly, the effect of assignment to Treatment 2b is somewhat smaller than the specification in which motivation is not observed and is no longer statistically significant, whereas the coefficient on the variable indicating assignment to Treatment 1 is larger in magnitude and now statistically significant at the 95%-confidence level. This, coupled with Treatment 2b's higher group mean for both situational motivation variables shown in Figure 3 suggests that a significant portion of the effect of assignment to Treatment 2b found in the model in which motivation is not observed is due to higher levels of situational motivation.

Discussion

The above results contribute to five key findings, which are now discussed in the context of the set of hypotheses (H) enumerated previously as well as the broader literature on this topic.

First, WTI in the final sample is far higher than anticipated, ranging from around 81% to 98%, which contradicts hypothesis **H1**. This is a relatively high range compared to other findings in the literature; however, these studies examined individuals' insurance preferences using either monetary payoffs (e.g.,McClelland et al. (1993)) or real-world endowments (e.g., Botzen and van den Bergh (2012)), so it is natural that WTI would be lower in these cases. Moreover, the inclusion of the option "maybe" in response to the question eliciting interest in purchasing an insurance policy rather than a simple binary response format may contribute to higher WTI values.

Second, respondents demonstrate relatively effective expectational reasoning skills in the contingent valuation experiment employed herein. As previously noted, expected damages to the endowed home in the hypothetical scenario are held constant at $62.50/month. The fact that the mean WTP and mean CWTP for the Control group are $58.30/month and $65.39/month, respectively, suggests that individuals in this sample are relatively well-equipped ex ante to calculate the actuarially fair premium (i.e., the premium equal to the expected damage amount) in this experimental setting. Again, this is likely an artifact of the manufactured scenario in which agents' demand is elicited: the relatively easy to manipulate numbers used in constructing the scenario, the lack of real-money consequences, and numerous other factors likely mitigate the behavioral biases that are evident in other experimental (e.g., Slovic et al. (1977); McClelland et al. (1993)) and real-world settings (e.g., Kunreuther and Slovic (1978); Atreya et al. (2015)). Moreover, Botzen and van den Bergh (2012) similarly find that respondents risk premium for hypothetical flood insurance policies are relatively small. Given that the primary objective of this research is to compare the relative effect of different behaviorally-informed policies, this finding does not invalidate the core findings regarding the nudge and boost interventions.

Third, the statistical numeracy boost as designed and implemented herein reduces respondents' WTI and WTP for insurance in the hypothetical scenario. In both specifications of the two-limit Tobit model of WTP, assignment to the boost is shown to have a statistically-significant, negative effect on WTP for flood insurance. This is a partial validation of **H2**: though a null effect of the boost is hypothesized, a robust negative effect is not. There are several possible explanations for this finding. First, it is possible that the boost employed herein is not designed effectively given its objective of improving statistical reasoning. This would possibly explain a null effect; however, it is unlikely that this alone would produce a robust negative effect. Second—and more likely— it is possible that the boost design inadvertently targets an existing behavioral bias that effectively nudges insurance demand downward. Slovic et al. (1978) may prove useful here: with respect to LPHC events, individuals appear to have a threshold probability value below which they effectively ignore the risk. Perhaps equipping individuals with a prescriptive heuristic to interpret probabilities in frequency terms in the case of flood insurance simply underscores the fact that the risk in question is below the threshold probability below which they ignore the risk.

Fourth, relatively simple, inexpensive nudges are more effective in increasing take-up and WTP for flood insurance than boosts. As shown in Table 5, the mean differences in WTP between Treatment 2a and Treatment 1 and Treatment 2b and Treatment 1 are both positive, though the former is not statistically significant. This is in line with previous results suggesting that altering the risk framing (Johnson et al., 1993; Botzen et al., 2013) is an effective means of enhancing risk perceptions in the insurance context. In particular, this matches Slovic et al. (1978)'s finding that extending the scale over which an LPHC risk is presented results in greater attention paid to the given risk. Overall, this paper finds sufficiently strong evidence to justify a weak preference for nudges over boosts in the context of flood insurance based on a comparison of their effects on WTP for insurance. While other factors may enter into the policymaker's objective function, these are outside of the scope of this paper and should be weighed against the results presented herein.

Fifth, targeting the affect and availability heuristics in addition to altering the risk presentation results in a significant increase in WTP for flood insurance, thereby validating **H3**.

In fact, it appears as though in order for a nudge to effectively alter risk perceptions in the context of flood insurance relative to the baseline of no behaviorally-informed intervention, it is necessary to explicitly target the affect and availability heuristics. As shown in Table 6, assignment to Treatment 2b (the affective nudge) is estimated to have a statistically-significant marginal effect on WTP for flood insurance of $12.50/month. While including the motivation variables in the estimation of WTP diminishes this effect (see Table 7), this is likely due to the fact that individuals assigned to Treatment 2b report, on average, higher levels of situational motivation (see Figure 3), which is associated with higher WTP (see Table 7 and Figure 4). While this suggests that the affect and availability heuristics are responsible for the large, positive effect of Treatment 2b, additional research is necessary to causally identify the role of these mechanisms.

Though great care is taken to ensure the robustness of these findings, there are several limitations to the analysis conducted herein that are worth noting explicitly. First, as discussed above, the lack of monetary consequences to insurance decisions made in the hypothetical scenario likely influences individuals' stated insurance preferences. Future work comparing these two policy types in this setting or similar settings should structure choice experiments in which insurance decisions have real monetary consequences. Second, the use of an open-ended contingent valuation method does raise some concern regarding the validity of the results; however, certain estimation techniques are employed to address this issue. Additional work in this area should make use of discrete choice experiments or other closed-form referendum formats in eliciting WTP for insurance against LPHC events.

Conclusion

That individuals in both experimental and empirical settings fail to purchase insurance against LPHC events—even in settings where premiums are partially subsidized—is a well-established fact. This finding challenges many conventional microeconomic models of consumer behavior in insurance markets (see, e.g. Rothschild and Stiglitz (1976)). In light of the many findings in the cognitive psychology and behavioral sciences literature suggesting that this phenomenon is likely the result of behavioral biases, behaviorally-informed policy tools are appropriate candidate solutions for increasing low take-up of insurance. Two categories of behaviorally-informed policies that are often used to address behavioral biases are nudges and boosts.

This study provides experimental evidence comparing the effectiveness of nudges and boosts in increasing low take-up rates in the context of flood insurance. Overall, this study finds that nudges are more effective than boosts in increasing take-up of and WTP for flood insurance. Moreover, the effectiveness of a nudge in increasing WTP for flood insurance relative to the baseline case in which no behaviorally-informed intervention is implemented appears to be conditional on the provision of information intended to elicit an affective response and prime respondents to consider salient, recent examples of coastal flooding.

These findings provide novel insight to policymakers seeking to increase take-up rates for insurance against flooding or other LPHC events. In particular, this study suggests two main takeaways for policymakers and practitioners. First, behavioral policy matters: the differences between individuals' WTP for flood insurance in two of the three behaviorally-informed interventions used in this study and the control condition in which no behaviorally-informed intervention is used are significant. Thus, any effort to increase take-up of insurance against flooding or other LPHC events should make use of behaviourally-informed interventions.

Second, policymakers should pay particular attention to the framing of risk when providing information in LPHC settings. Extending the time horizon over which risk probabilities are provided and presenting risk information alongside salient examples increases the attention individuals pay to the risk in question. In the case of flood insurance, this finding suggests that the common practice of presenting flood risk in terms of annual probability thresholds (i.e., a 1% probability of flooding in a given year) is ineffective and that this risk should be reframed if policymakers are to increase the salience of this risk communication.

Though additional work is necessary to provide an empirically-based framework for understanding those conditions under which nudges or boosts are most effective in this setting, this paper suggests that policymakers should consider using relatively inexpensive, affective nudges to increase take up of insurance against LPHC events.

References

Amemiya, T. (1984). Tobit models: A survey. *Journal of Econometrics*, 24(1-2):3–61.

Atreya, A., Ferreira, S., and Michel-Kerjan, E. (2015). What drives households to buy flood insurance? New evidence from Georgia. *Ecological Economics*, 117:153–161.

Benjamin, D. J. (2018). Errors in probabilistic reasoning and judgment biases. Working Paper 25200, National Bureau of Economic Research.

Berinsky, A. J., Huber, G. A., and Lenz, G. S. (2012). Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk. *Political Analysis*, 20(3):351–368.

Botzen, W., de Boer, J., and Terpstra, T. (2013). Framing of risk and preferences for annual and multi-year flood insurance. *Journal of Economic Psychology*, 39:357–375.

Botzen, W. and van den Bergh, J. (2012). Risk attitudes to low-probability climate change risks: WTP for flood insurance. *Journal of Economic Behavior & Organization*, 82(1):151–166.

Browne, M. J., Knoller, C., and Richter, A. (2015). Behavioral bias and the demand for bicycle and flood insurance. *Journal of Risk and Uncertainty*, 50(2):141–160.

Cameron, T. (2005). Updating Subjective Risks in the Presence of Conflicting Information: An Application to Climate Change. *Journal of Risk and Uncertainty*, 30(1):63–97.

Chetty, R. (2015). Behavioral Economics and Public Policy: A Pragmatic Perspective. *American Economic Review*, 105(5):1–33.

Clifford, S., Jewell, R. M., and Waggoner, P. D. (2015). Are samples drawn from Mechanical Turk valid for research on political ideology? *Research & Politics*, 2(4):1–9.

de Boer, J., Botzen, W. J., and Terpstra, T. (2015). More Than Fear Induction: Toward an Understanding of People's Motivation to Be Well-Prepared for Emergencies in Flood-Prone Areas: More Than Fear Induction. *Risk Analysis*, 35(3):518–535.

de Boer, J., Botzen, W. J. W., and Terpstra, T. (2014). Improving flood risk communication by focusing on prevention-focused motivation. *Risk Analysis*, 34(2):309–322.

Finucane, M. L., Alhakami, A., Slovic, P., and Johnson, S. M. (2000). The Affect Heuristic in Judgements of Risks and Benefits. *Journal of Behavioral Decision Making*, 13(1):1–17.

Fong, G., H Krantz, D., and Nisbett, R. (1986). The Effects of Statistical Training on Thinking About Everyday Problems. *Cognitive Psychology*, 18:253–292.

Fong, G. T. and Nisbett, R. E. (1991). Immediate and delayed transfer of training effects in statistical reasoning. *Journal of Experimental Psychology: General*, 120(1):34–45.

Gigerenzer, G. (1996). Why do frequency formats improve Bayesian reasoning? Cognitive algorithms work on information, which needs representation. *Behavioral and Brain Sciences*, 19(01):23.

Gigerenzer, G. (2014). Breast cancer screening pamphlets mislead women. *BMJ (Clinical research ed.)*, 348:g2636.

Gigerenzer, G. and Brighton, H. (2009). Homo Heuristicus: Why Biased Minds Make Better Inferences. *Topics in Cognitive Science*, 1(1):107–143.

Gigerenzer, G. and Hoffrage, U. (1995). How to Improve Bayesian Reasoning Without Instruction: Frequency Formats. *Psychological Review*, 102(4):684–704.

Gigerenzer, G. and Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart*, Evolution and cognition, pages 3–34. Oxford University Press, New York, NY, US.

Gneezy, U. and Potters, J. (1997). An Experiment on Risk Taking and Evaluation Periods. *The Quarterly Journal of Economics*, 112(2):631–645.

Grüne-Yanoff, T. and Hertwig, R. (2016). Nudge versus boost: How coherent are policy and theory? *Minds and Machines: Journal for Artificial Intelligence, Philosophy and Cognitive Science*, 26(1-2):149–183.

Grüne-Yanoff, T., Marchionni, C., and Feufel, M. A. (2018). Toward a Framework for Selecting Behavioral Policies: How to Choose between Boosts and Nudges. *Economics and Philosophy*, 34(2):243–266.

Hertwig, R., Barron, G., Weber, E. U., and Erev, I. (2004). Decisions from Experience and the Effect of Rare Events in Risky Choice. *Psychological Science*, 15(8):534–539.

Higgins, E. (1998). Promotion and Prevention: Regulatory Focus as a Motivational Principle. *Advances in Experimental Social Psychology*, 30:1–46.

Hsee, C. K. (1996). The Evaluability Hypothesis: An Explanation for Preference Reversals between Joint and Separate Evaluations of Alternatives. *Organizational Behavior and Human Decision Processes*, 67(3):247–257.

Johnson, E. J., Hershey, J., and Meszaros, J. (1993). Framing, Probability Distortions, and Insurance Decisions. *Journal of Risk & Uncertainty*, 7(1):35–51.

Kahneman, D., Slovic, P., and Tversky, A. (1982). *Judgment under uncertainty: heuristics and biases*. Cambridge University Press, Cambridge.

Kahneman, D. and Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, 103(3):582–591.

Keller, C., Siegrist, M., and Gutscher, H. (2006). The Role of the Affect and Availability Heuristics in Risk Communication. *Risk Analysis*, 26(3):631–639.

Knight, F. (1921). *Risk, Uncertainty, and Profit.* Number No. 31 in Hart, Schaffner, and Marx Prize Essays. Houghton Mifflin, Boston and New York.

Kousky, C. (2018). Financing Flood Losses: A Discussion of the National Flood Insurance Program. *Risk Management and Insurance Review*, 21(1):11–32.

Krantz, D. H. and Kunreuther, H. C. (1998). Goals and Plands in Decision Making. *Judgement and Decision Making*, 2(3):137–168.

Kunreuther, H. (2018). Improving the National Flood Insurance Program. *Behavioural Public Policy*, https://doi.org/10.1017/bpp.2018.26.

Kunreuther, H., Meyer, R., Zeckhauser, R., Slovic, P., Schwartz, B., Schade, C., Luce, M. F., Lippman, S., Krantz, D., Kahn, B., and Hogarth, R. (2002). High Stakes Decision Making: Normative, Descriptive and Prescriptive Considerations. *Marketing Letters*, 13(3):259–268.

Kunreuther, H., Novemsky, N., and Kahneman, D. (2001). Making Low Probabilities Useful. *The Journal of Risk and Uncertainty*, 23(2):103–120.

Kunreuther, H. and Pauly, M. (2018). Behavioral Economics of Multiperiod Insurance Purchasing Behavior: The Role of Emotions. *Foundations and Trends in Microeconomics*, 12(2):109–199.

Kunreuther, H. and Slovic, P. (1978). Economics, Psychology, and Protective Behavior. *American Economic Review*, 68(2):64–69.

Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., and Combs, B. (1978). Judged Frequency of Lethal Events. *Journal of Experimental Psychology: Human Learning and Memory*, 4(6):551–578.

Logar, I. and Brouwer, R. (2017). The effect of risk communication on choice behavior, welfare estimates and choice certainty. *Water Resources and Economics*, 18:34–50.

McClelland, G. H., Schulze, W. D., and Coursey, D. L. (1993). Insurance for Low-Probability Hazards: A Bimodal Response to Unlikely Events. *Journal of Risk and Uncertainty*, 7:95–116.

Morewedge, C. K., Yoon, H., Scopelliti, I., Symborski, C. W., Korris, J. H., and Kassam, K. S. (2015). Debiasing Decisions: Improved Decision Making With a Single Training Intervention. *Policy Insights from the Behavioral and Brain Sciences*, 2(1):129–140.

Nisbett, R. E., Fong, G. T., Lehman, D. R., and Cheng, P. W. (1987). Teaching reasoning. *Science*, 238(4827):625–631.

Oliver, A. E. (2013). *Behavioural Public Policy.* Cambridge University Press, Cambridge.

Peters, E., Kunreuther, H., Sagara, N., Slovic, P., and Schley, D. R. (2012). Protective Measures, Personal Experience, and the Affective Psychology of Time. *Risk Analysis*, 32(12):2084–2097.

Rothschild, M. and Stiglitz, J. (1976). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *The Quarterly Journal of Economics*, 90(4):629–649.

Schade, C., Kunreuther, H., and Koellinger, P. (2012). Protecting Against Low-Probability Disasters: The Role of Worry: Protecting Against Low-Probability Disasters. *Journal of Behavioral Decision Making*, 25(5):534–543.

Sedlmeier, P. and Gigerenzer, G. (2001). Teaching Bayesian Reasoning in Less Than Two Hours. *Journal of Experimental Psychology: General*, 130(3):380–400.

Shafir, E. (2013). *The Behavioral Foundations of Public Policy*. Princeton University Press, Princeton.

Sims, A. and Müller, T. M. (2018). Nudge Versus Boost: A Distinction without a Difference. *Economics and Philosophy*, pages 1–28.

Slovic, P., editor (2000). *The perception of risk*. Earthscan Publications, London, England.

Slovic, P., Finucane, M., Peters, E., and MacGregor, D. G. (2002). Rational actors or rational fools: implications of the affect heuristic for behavioral economics. *The Journal of Socio-Economics*, 31(4):329–342.

Slovic, P., Fischhoff, B., and Lichtenstein, S. (1978). Accident probabilities and seat belt usage: A psychological perspective. *Accident Analysis & Prevention*, 10(4):281–285.

Slovic, P., Fischhoff, B., Lichtenstein, S., Corrigan, B., and Combs, B. (1977). Preference for Insuring Against Probable Small Losses: Insurance Implications. *Journal of Risk and Insurance (pre-1986); Malvern*, 44(2):237–258.

Sunstein, C. (2015). Nudges, Agency, and Abstraction: A Reply to Critics. *Review of Philosophy and Psychology*, 6(3):511–529.

Sunstein, C. R. (2002). Probability Neglect: Emotions, Worst Cases, and Law. *The Yale Law Journal*, 112(1):61–107.

Thaler, R., Tversky, A., Kahneman, D., and Schwartz, A. (1997). The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test. *The Quarterly Journal of Economics*, 112(2):647–661.

Thaler, R. H. and Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press, New Haven.

Tobin, J. (1958). Estimation of Relationships for Limited Dependent Variables. *Econometrica*, 26(1):24–36.

Tom, S. M., Fox, C. R., Trepel, C., and Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, 315(5811):515–518.

Tukey, J. W. (1949). Comparing Individual Means in the Analysis of Variance. *Biometrics*, 5(2):99–114.

Tversky, A. and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2):207–232.

Tversky, A. and Kahneman, D. (1986). Rational Choice and the Framing of Decisions. *Journal of Business*, 59(4):S251–S278.

Tversky, A. and Shafir, E. (1992). The Disjunction Effect in Choice under Uncertainty. *Psychological Science*, 3(5):305–310.

Viscusi, W. K. (1985). Are Individuals Bayesian Decision Makers? *The American Economic Review*, 75(2):381–385.

Viscusi, W. K. and O'Connor, C. J. (1984). Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers? *The American Economic Review*, 74(5):942–956.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*, volume 1. The MIT Press, Cambridge, 2 edition.

Tables

Table 1: Summary of different treatment and control conditions. Each condition is characterized by (1) its intervention type, (2) the format in which the inundation risk is presented, and (3) the structure of its decision frame. "Simple hypothetical" refers to the hypothetical scenario used to elicit WTP.

	Control	**Treatment 1**	**Treatment 2a**	**Treatment 2b**
Intervention:	–	Statistical numeracy boost	Altered decision environment	Altered decision environment
Risk presentation:	Annual risk probability	Annual risk frequency	Multi-year risk probability	Multi-year risk probability
Decision frame:	Simple hypothetical	Statistical boost and simple hypothetical	Simple hypothetical	Affective / availability cue and simple hypothetical
Pr{Selection} =	0.25	0.25	0.25	0.25

Table 2: Description of primary explanatory variables elicited via survey instrument.

Variable	Description
WTP motivation:	
$Situational\_Worry_i$	Categorical variable (1-5); indicates agreement with the statement "I would worry about the possibility of experiencing flooding" in reference to the hypothetical scenario; 1 = strongly disagree, 5 = strongly agree.
$Situational\_Prepare_i$	Categorical variable (1-5); indicates agreement with the statement "I would make sure that I am prepared for possible flooding" in reference to the hypothetical scenario; 1 = strongly disagree, 5 = strongly agree.
$Global\_Financial_i$	Categorical variable (1-5); indicates agreement with the statement "financial security is important to me;" 1 = strongly disagree, 5 = strongly agree.
$Global\_Safe_i$	Categorical variable (1-5); indicates agreement with the statement "a safe environment is important to me;" 1 = strongly disagree, 5 = strongly agree.
Prior beliefs:	
$FloodExp_i$	Binary variable; 1 = respondent has experienced a flood in the past, 0 = respondent has not experienced a flood in the past.
$DisasExp_i$	Binary variable; 1 = respondent has experienced a disaster in the past, 0 = respondent has not experienced a disaster in the past.
$ClimateChange_i$	Categorical variable (1-5); indicates agreement with the statement "man-made climate change is occurring;" 1 = strongly disagree, 5 = strongly agree.
SLR_i	Categorical variable (1-5); indicates agreement with the statement "sea-level rise is occurring as a result of climate change;" 1 = strongly disagree, 5 = strongly agree.
Political views:	
$Political\_Party_i$	Categorical variable (1-7); 1 = strongly Republican, 7 = strongly Democratic.
$Political\_Ideology_i$	Categorical variable (1-7); 1 = extremely conservative, 7 = extremely liberal.
Home ownership:	
$DetachedHome_i$	Binary variable. 1 = respondent owns a single-family, detached home; 0 = respondent does not own a singly-family, detached home.
$CoastalState_i$	Binary variable; 1 = respondent lives in a coastal state, 0 = respondent does not live in a coastal state.
Demographics:	
Age_i	Categorical variable (1-7); "under 20 years old" to "over 69 years old." Modeled as a continuous variable assigning each respondent the midpoint of their range.
$Education_i$	Categorical variable (1-9); "no schooling completed" to "doctoral degree."
$Children_i$	Number of children; possible responses: 0, 1, 2, 3, 4+ (assigned 4).
$Income_i$	Categorical variable (1-6); "less than \$25,000/year" to "greater than \$125,000/year." Modeled as a continuous variable assigning each respondent the midpoint of their range.

Table 3: Summary statistics of the explanatory variables.

Variable	Mean	St. Dev.
WTP motivation:		
$Situational\_Worry_i$	3.819	1.038
$Situational\_Prepare_i$	4.088	0.870
$Global\_Financial_i$	4.350	0.761
$Global\_Safe_i$	4.166	0.800
Prior beliefs:		
$FloodExp_i$	0.291	0.455
$DisasterExp_i$	0.469	0.500
$ClimateChange_i$	4.103	1.062
SLR_i	4.109	0.919
Political views:		
$Political\_Party_i$	3.603	2.113
$Political\_Ideology_i$	3.653	1.782
Home ownership:		
$DetachedHome_i$	0.344	0.476
$CoastalState_i$	0.609	0.489
Demographics:		
Age_i	35.203	11.332
$Female_i$	0.416	0.494
$Education_i$	5.322	1.434
$Children_i$	0.697	0.953
$Income_i$	56,875.00	34,621.78

Table 4: Willingness-to-insure (WTI), mean willingness-to-pay (WTP) and mean conditional willingness-to-pay (CWTP) for the full sample, Control condition, Treatment 1, Treatment 2a, and Treatment 2b.

Panel	WTI (% of respondents)	WTP ($/month)	CWTP ($/month)
Full sample ($n = 320$)	89.06	59.54	66.85
		(2.03)	(1.87)
Control ($n = 83$)	89.16	58.30	65.39
		(3.90)	(3.57)
Treatment 1 ($n = 86$)	81.40	49.63	60.97
		(3.91)	(3.63)
Treatment 2a ($n = 82$)	89.02	61.15	68.68
		(4.32)	(4.04)
Treatment 2b ($n = 69$)	98.55	71.46	72.51
		(3.70)	(3.60)

Note: standard errors reported in parentheses.

Table 5: Pairwise treatment/control group comparison of mean WTP (Tukey's test).

	Difference ($/month)	95% CI ($/month)	Adjusted p-value
(Treatment 1−Control)	-8.673	[-22.869, 5.523]	0.393
(Treatment 2a−Control)	2.845	[-11.520, 17.210]	0.956
(Treatment 2b−Control)	13.163	[-1.868, 28.193]	0.109
(Treatment 2a−Treatment 1)	11.518	[-2.722, 25.758]	0.159
(Treatment 2b−Treatment 1)	21.836	[6.925, 36.747]	0.001
(Treatment 2b−Treatment 2a)	10.317	[-4.754, 25.389]	0.291

Note: Tukey's test (Tukey, 1949) assumes i.i.d. random sampling; observations are normally-distributed within group; and homoscedasticity. Examination of the within sample distributions and consideration of the data generating process suggest that these assumptions are reasonable (see Appendix D). One way ANOVA testing H_0 : no difference between the mean WTP across treatment and control groups rejects the null hypothesis with F-value of 4.859 ($p = 0.003$).

Table 6: Estimation results of two-limit Tobit model of WTP.

	Coefficient	Standard Error	Marginal Effect
(Treatment 1)$_i$	-12.176^{*}	6.369	-10.740
(Treatment 2a)$_i$	2.621	6.394	2.312
(Treatment 2b)$_i$	14.169^{**}	6.643	12.498
$FloodExp_i$	20.412^{***}	5.285	18.005
$ClimatePrior_i$	0.241	2.606	0.212
$Political_i$	-0.620	0.748	-0.547
Age_i	-0.150	0.222	-0.133
$Children_i$	4.577^{*}	2.734	4.037
$Income_i$	0.00002	0.0001	0.00002
$University_i$	-3.354	5.110	-2.959
$DetachedHome_i$	-3.140	5.574	-2.770
$CoastalState_i$	-0.411	4.737	-0.362
Constant	59.216^{***}	11.162	–
σ	39.741^{***}	1.046	–
Observations			320
Log Likelihood			$-1{,}436.801$
$-2\log(L_{R_1}/L_U)$			36.097^{***}
$-2\log(L_{R_2}/L_U)$			15.856^{***}

Note: *(Treatment 1)*$_i$, *(Treatment 2a)*$_i$, and *(Treatment 2b)*$_i$ are binary variables indicating individual i's treatment assignment. $ClimatePrior_i$ is generated as the sum of $ClimateChange_i$ and SLR_i for individual i ($\alpha = 0.84$) and is normalized to have mean of zero and variance of one. $Political_i$ is generated as the sum of $Political\_Party_i$ and $Political\_Ideology_i$ for individual i ($\alpha = 0.76$) and is normalized to have mean of zero and variance of one. $University_i$ is a binary indicator variable that =1 if individual i's highest schooling completed is greater than or equal to a Bachelor's degree. L_{R_1} is the likelihood value for the completely restricted (i.e., just a constant) model. L_{R_2} is the likelihood value for the restricted model excluding just the treatment status variables. *p<0.1; **p<0.05; ***p<0.01

Table 7: Estimation results of two-limit Tobit model of WTP observing motivation.

	Coefficient	Standard Error	Marginal Effect
(Treatment 1)$_i$	-13.345^{**}	5.948	-12.109
(Treatment 2a)$_i$	-1.479	6.003	-1.342
(Treatment 2b)$_i$	9.039	6.227	8.202
$Situational\_Worry_i$	5.469^{**}	2.658	4.962
$Situational\_Prepare_i$	10.455^{***}	2.765	9.487
$Global\_Financial_i$	-5.111^{**}	2.504	-4.638
$Global\_Safe_i$	6.690^{***}	2.542	6.071
$FloodExp_i$	17.676^{***}	4.985	16.040
$ClimatePrior_i$	-2.159	2.486	-1.959
$Political_i$	-0.795	0.703	-0.723
Age_i	-0.267	0.207	-0.242
$Children_i$	5.881^{**}	2.564	5.336
$Income_i$	0.00000	0.0001	0.000002
$University_i$	-2.309	4.787	-2.095
$DetachedHome_i$	-8.258	5.249	-7.493
$CoastalState_i$	0.842	4.417	0.764
Constant	67.954^{***}	10.486	–
σ	36.885^{***}	1.046	–
Observations			320
Log Likelihood			$-1{,}411.343$
$-2\log(L_{R_1}/L_U)$			87.012^{***}
$-2\log(L_{R_2}/L_U)$			13.318^{***}

Note: *(Treatment 1)$_i$*, *(Treatment 2a)$_i$*, and *(Treatment 2b)$_i$* are binary variables indicating individual i's treatment assignment. The motivation variables are normalized to have mean of zero and variance of one. $ClimatePrior_i$ is generated as the sum of $ClimateChange_i$ and SLR_i for individual i ($\alpha = 0.84$) and is normalized to have mean of zero and variance of one. $Political_i$ is generated as the sum of $Political\_Party_i$ and $Political\_Ideology_i$ for individual i ($\alpha = 0.76$) and is normalized to have mean of zero and variance of one. $University_i$ is a binary indicator variable that =1 if individual i's highest schooling completed is greater than or equal to a Bachelor's degree. L_{R_1} is the likelihood value for the completely restricted (i.e., just constant) model. L_{R_2} is the likelihood value for the restricted model excluding just the treatment status variables. *p<0.1; **p<0.05; ***p<0.01

Figures

Figure 1: Mean and 95% confidence intervals for WTP and CWTP across treatment and control groups.

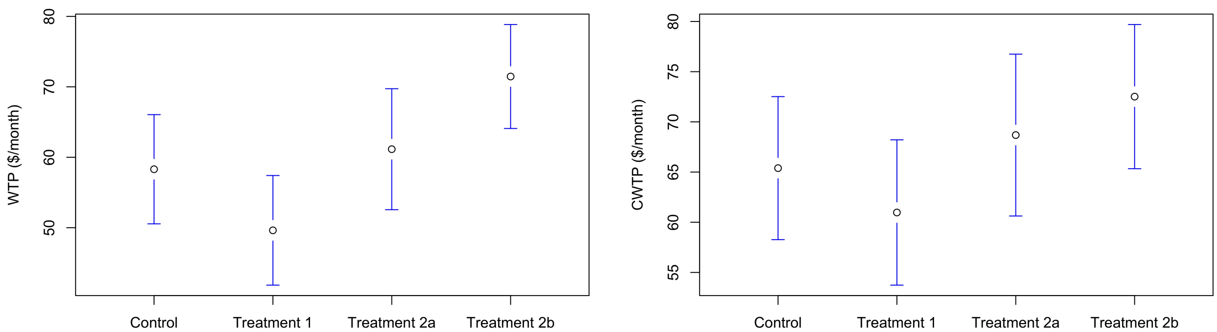

Figure 2: Distribution of WTP values for full sample ($n = 320$). A normal distribution with moments equal to the sample moments is overlaid to show the censoring at the tails of the sample distribution.

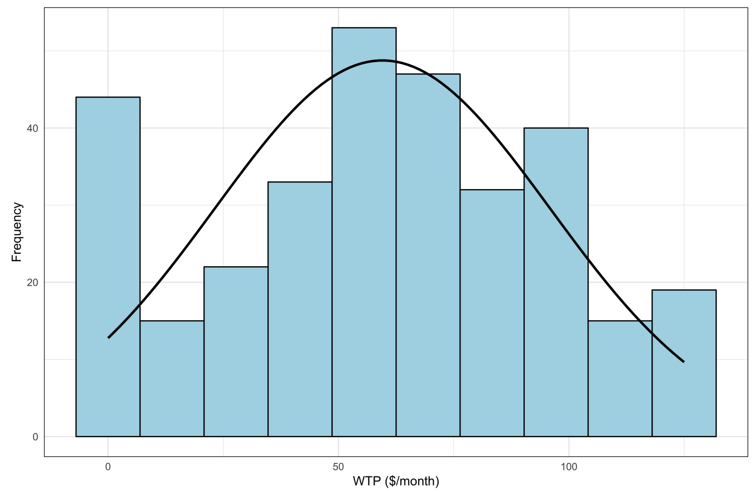

Figure 3: Mean and 95% confidence intervals for the four motivation variables across treatment and control groups: worry (situational), preparation (situational), safety (global), and financial security (global). The motivational variables are transformed into z-scores.

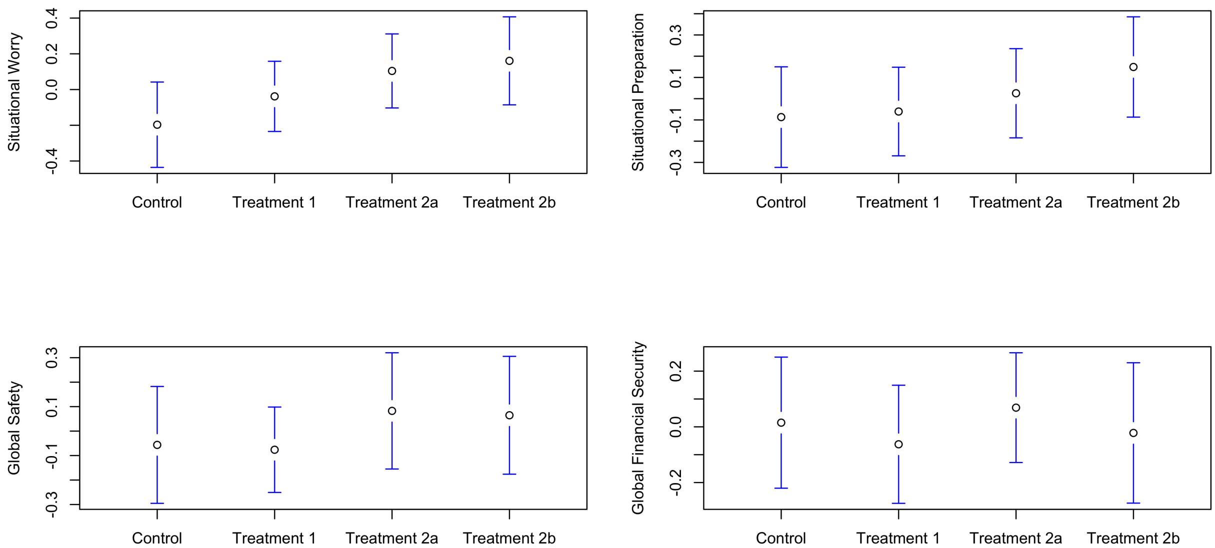

Figure 4: The effect of situational motivation on WTI and mean WTP. The situational motivation categorical variable is generated by: summing the values of the situational worry and prevention variables, generating the associated z-score for the resulting sum, and dividing individuals into 7 equally-sized quantiles. Motivation values range from 1 = lowest z-score to 7 = highest z-score.

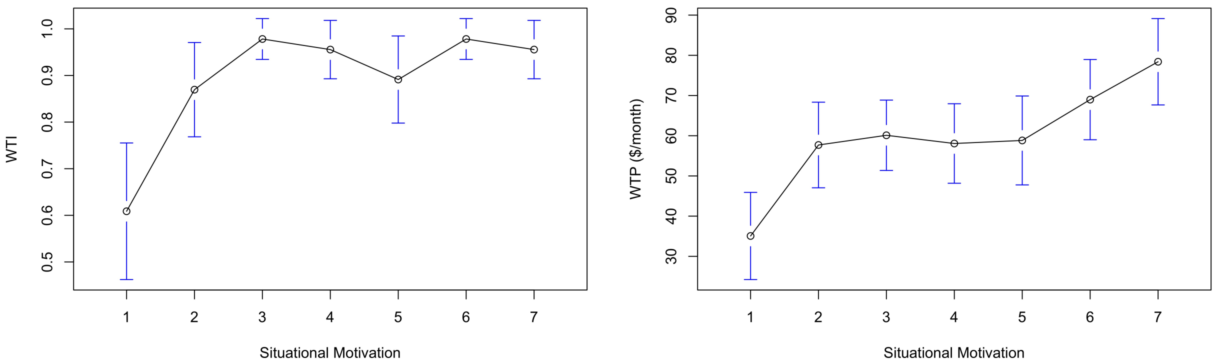

Appendix A. Hypothetical Insurance Scenarios

The hypothetical scenarios presented to respondents in each of the four treatment and control groups are presented below. After being presented with one of the scenarios below, respondents are each asked the following three questions:

1. "Consider an annual insurance policy with a coverage amount of $75,000. What do you believe is the monthly cost for such a policy?"
2. "Would you be interested in purchasing an annual flood insurance policy for your Coastalville home?"
3. "As a Coastalville homeowner, what is the highest amount you would be willing to pay per month to purchase an annual insurance policy with a coverage amount of $75,000?"

Control Contition
Respondents in the Control Condition are presented with the following hypothetical scenario:

> We will now ask you to consider a hypothetical scenario and then indicate your interest in different insurance policies that cover damages due to coastal flooding. Imagine the following scenario:
>
> **Coastal Area:** You currently own a standalone home in Coastalville, which is located in a floodplain in the coastal U.S. The total value of your home is estimated to be $300,000. There is a 1% chance, in any given year, that you will experience flooding resulting in approximately $75,000 worth of damages.
>
> An annual insurance policy that will cover the cost of damages associated with severe flooding is available in exchange for a monthly payment. An insurance policy's "coverage amount" indicates the maximum monetary amount you can be reimbursed under the policy.

Treatment 1
Before being presented with the hypothetical scenario, respondents in Treatment 1 are provided with the following information:

> We will now ask you to consider a hypothetical scenario and then indicate your interest in different insurance policies that cover damages due to coastal flooding. Before we do so, we would like to provide you with some information on the principles of probability that may prove useful as you think through the questions in this section.
>
> > It is often helpful to think about the probability of an event occurring as the frequency with which that event will occur in the long run. For example, the probability of landing on "heads" when tossing a fair coin is 50%, or 1/2. We can interpret this probability as implying that in a large number of tosses—say, 100 tosses—the frequency with which "heads" actually occurs will be approximately 1/2: for each toss resulting in "heads", we expect a toss to result in "tails." Thus, flipping our fair coin 100 times, we expect the coin to land on "heads" approximately 50 times and "tails" approximately 50 times.
> >
> > This interpretation of probability applies not only to coin flips, but also to a wide array of problems that individuals face on a regular basis. An easy rule for converting probabilities to expected long run frequencies is to convert the probability from a percentage or fraction to a frequency by multiplying the given probability by 100. This allows you to estimate the expected number of times an event will occur over 100 iterations. For example, if the probability of your friend not returning the umbrella that you lend him/her is 1%, or 1/100, each time that you do so, then you would expect that if you lend him/her the umbrella 100 times, he/she will not return the umbrella on 1 of those occasions

Respondents in Treatment 1 are then presented with the following hypothetical scenario:

> Imagine the following scenario:
>
> > **Coastal Area:** You currently own a standalone home in Coastalville, which is located in a floodplain in the coastal United States. The total value of your home is estimated to be $300,000. There is a 1% (or 1 in 100) chance, in any given year, that you will experience flooding resulting in approximately $75,000 worth of damages. This implies that across 100 different floodplains in the coastal United States with the same 1% probability of experiencing this type of severe flood, we expect that 1 of those floodplains would experience this type of severe flood in a given year. Similarly, this implies that over the course of 100 years, we expect Coastalville to experience 1 severe flood of this type.
> >
> > An annual insurance policy that will cover the cost of damages associated with severe flooding is available in exchange for a monthly payment. An insurance policy's "coverage amount" indicates the maximum monetary amount you can be reimbursed under the policy.

Treatment 2a
Respondents in Treatment 2a are presented with the following hypothetical scenario:

> We will now ask you to consider a hypothetical scenario and then indicate your interest in different insurance policies that cover damages due to coastal flooding. Imagine the following scenario:
>
> **Coastal Area:** You currently own a standalone home in Coastalville, which is located in a floodplain in the coastal United States. The total value of your home is estimated to be $300,000. There is a 1% chance, in any given year, that you will experience flooding resulting in approximately $75,000 worth of damages. While the probability of experiencing a severe flood of this type over the course of a year is relatively low, over the course of 30 years of home ownership, the probability of experiencing a severe flood of this type is approximately 26%.
>
> An annual insurance policy that will cover the cost of damages associated with severe flooding is available in exchange for a monthly payment. An insurance policy's "coverage amount" indicates the maximum monetary amount you can be reimbursed under the policy.

Treatment 2b

Before being presented with the hypothetical scenario, respondents in Treatment 2b are provided with the following information:

> We will now ask you to consider a hypothetical scenario and then indicate your interest in different insurance policies that cover damages due to coastal flooding. Before we do so, we would like to provide you with some information about coastal flooding that may prove useful as you think through the questions in this section.
>
> > In the United States, more than 8.6 million Americans live in areas susceptible to coastal flooding, with more than $1 trillion of property and structures located within just a few feet of current sea level. Moreover, the problem of coastal flooding is expected to worsen as the sea level rises and major storm events intensify due to climate change. By 2050, experts predict that a majority of US coastal areas are likely to be threatened by 30 or more days of flooding each year. Higher sea levels and more intense storms will contribute to more damaging flooding events like Hurricanes Harvey and Irma, which made landfall in Texas and Florida, respectively, during the 2017 hurricane season. In total, the 2017 hurricane season was the costliest on record, with total damages of at least $282 billion.

The above information is also accompanied by the following photos:

Respondents in Treatment 2b are then presented with the same hypothetetical scenario as that shown to respondents in Treatment 2a.

Appendix B. Motivation Elicitation

Following the WTP elicitation, respondents were then asked to consider the hypothetical scenario (varying by treatment status; see Appendix A) and indicate their level of agreement or disagreement with the following statements:

- *[Situational]* If I lived in Coastalville, then I would worry about the possibility of experiencing flooding.
- *[Situational]* If I lived in Coastalville, then I would make sure that I am prepared for possible flooding.
- *[Global]* A safe environment is important to me; I prefer to avoid risky situations and I prefer to protect myself from natural hazards.
- *[Global]* Financial security is important to me.

Appendix C. Estimation Method

Respondents are asked to report their maximum WTP for an annual flood insurance policy in a hypothetical scenario on a continuous sliding scale from \$0/month-\$125/month. The resulting positive, continuous measure of respondents' WTP is therefore censored from below and above by 0 and 125, respectively.[8] Thus, WTP_i given by Eq. (7) is treated as an unobserved latent variable. Instead, draws of censored data $(WTP_i^*, X_i, D_i, C_I, P_i, T_i)$ are observered. WTP_i^* is characterized by:

$$WTP_i^* = \begin{cases} 0 & \text{if } WTP_i \leq 0 \\ WTP_i & \text{if } 0 < WTP_i < 125 \\ 125 & \text{if } WTP_i \geq 125 \end{cases} \tag{9}$$

Thus, for $WTP_i \in (0, 125)$, the observed WTP_i^* is given by Eq. (7). Note that, in the survey instrument employed herein, there is no heterogeneity in D_i and P_i, so these are not included in the estimation of Eq. (7) since they are captured by the inclusion of a constant term. Given the structure of Eq. (9), the two-limit Tobit model is employed.[9] Define:

$$Z_i = \begin{bmatrix} 1 & X_i & C_i & T_i \end{bmatrix} \qquad \beta = \begin{bmatrix} \beta_0 & \beta_1 & \tilde{\beta}_3 & \gamma \end{bmatrix}^T \tag{10}$$

It can be shown that:

$$\Pr\{WTP_i^* = 0 | Z_i\} = \Phi\left(\frac{-Z_i\beta}{\sigma_\epsilon}\right) \tag{11}$$

$$\Pr\{WTP_i^* \in (0, 125) | Z_i\} = \Phi\left(\frac{Z_i\beta}{\sigma_\epsilon}\right) \tag{12}$$

$$\Pr\{WTP_i^* = 125 | Z_i\} = \Phi\left(\frac{-(125 - Z_i\beta)}{\sigma_\epsilon}\right) \tag{13}$$

[8] This assumption allows for WTP values that are less than \$0/month and greater than \$125/month, which, given the data observed, is reasonable. Of course, it is possible that there are valid WTP responses elicited that equal \$0 or \$125 precisely. However, in the sample used herein, no respondents who indicated interest in purchasing the insurance product indicated a WTP equal to \$0.

[9] The Tobit model was first developed by Tobin (1958). Amemiya (1984) provides a useful classification of different forms of the Tobit model. The two-limit Tobit is a natural extension of the Type I Tobit to the case of left- and right-censoring. For a helpful reference on the two-limit Tobit, see Wooldridge (2010). A sample selection model, such as Heckman's two-stage estimator, is deemed inappropriate in this setting as (1) the assumption of negative WTP values and WTP values greater than 125 is reasonable; and (2) those same factors affecting the decision to enter the insurance market explain WTP elicited via the survey and these factors operate in the same direction (Amemiya, 1984). Assuming that agents employ a decision rule for market entrance that states that they will enter the insurance market (i.e., indicate interest in the hypothetical insurance product) if and only if their WTP for insurance is non-negative, then using the Tobit model instead of a sample selection model is reasonable. This decision rule is viewed as reasonable in the current setting.

where $\Phi(\cdot)$ is the standard normal cumulative density function. It follows that the log-likelihood for a random draw i is:

$$\begin{aligned}\log f(WTP_i^*|Z_i;\beta) &= \mathbb{1}\{WTP_i^* = 0\}\log\left[\Phi\left(\frac{-Z_i\beta}{\sigma_\epsilon}\right)\right] \\ &+ \mathbb{1}\{WTP_i^* = 125\}\log\left[\Phi\left(\frac{-(125 - Z_i\beta)}{\sigma_\epsilon}\right)\right] \\ &+ \mathbb{1}\{WTP_i^* \in (0,125)\}\log\left[\frac{1}{\sigma_\epsilon}\phi\left(\frac{WTP_i^* - Z_i\beta}{\sigma_\epsilon}\right)\right] \qquad (14)\end{aligned}$$

where $\phi(\cdot)$ is the standard normal probability density function. The two-limit Tobit model maximizes the log-likelihood function for the sample:

$$\hat{\beta} = \arg\max_{\beta} \frac{1}{N}\sum_{i=1}^{N}\log f(WTP_i^*|Z_i;\beta) \qquad (15)$$

It follows that

$$\mathbb{E}[WTP_i^*|Z_i, WTP_i^* \in (0,125)] = Z_i\beta + \sigma\left[\frac{\phi\left(\frac{-Z_i\beta}{\sigma_\epsilon} - \phi\left(\frac{-Z_i\beta}{\sigma_\epsilon}\right)\right)}{\Phi\left(\frac{125-Z_i\beta}{\sigma_\epsilon} - \Phi\left(\frac{-Z_i\beta}{\sigma_\epsilon}\right)\right)}\right] \qquad (16)$$

and the unconditional expectation of WTP_i^* is given by:

$$\begin{aligned}\mathbb{E}[WTP_i^*|Z_i] = &\Pr\{WTP_i^* \in (0,125)\}\mathrm{E}[WTP_i^*|Z_i, WTP_i^* \in (0,125)] \\ &+ 125 \times \Phi\left(\frac{-(125 - Z_i\beta)}{\sigma_\epsilon}\right) \qquad (17)\end{aligned}$$

The partial effect of covariate Z^j simplifies to:

$$\frac{\partial\mathbb{E}[WTP^*|Z]}{\partial Z^j} = \left[\Phi\left(\frac{125 - Z\beta}{\sigma_\epsilon} - \Phi\left(\frac{-Z\beta}{\sigma_\epsilon}\right)\right)\right]\beta_j \qquad (18)$$

The average partial effects are therefore given as:

$$APE(Z^j) = \left(\frac{1}{N}\sum_{i=1}^{N}\left[\Phi\left(\frac{125 - Z\hat{\beta}}{\hat{\sigma}_\epsilon} - \Phi\left(\frac{-Z\hat{\beta}}{\hat{\sigma}_\epsilon}\right)\right)\right]\right)\hat{\beta}_j \qquad (19)$$

for continuous regressor Z^j. The results of the two-limit Tobit model characterized by Eq. (9) are reported in Table 6.

Appendix D. Distribution of WTP by Treatment Group

The figure below shows the distribution of WTP by treatment/control status (group mean shown as vertical dashed line). Though there is heterogeneity across groups, there is clear left- and right-censoring at $0/month and $125/month, respectively.