Annual Review of Resource Economics

Understanding the Improbable: A Survey of Fat Tails in Environmental Economics

Marc N. Conte[1] and David L. Kelly[2]

[1]Department of Economics, Fordham University, Bronx, New York 10458, USA; email: mconte7@fordham.edu

[2]Department of Economics, University of Miami, Coral Gables, Florida 33146, USA; email: dkelly@miami.edu

Annu. Rev. Resour. Econ. 2021. 13:17.1–17.22

The *Annual Review of Resource Economics* is online at resource.annualreviews.org

https://doi.org/10.1146/annurev-resource-102020-094143

JEL codes: I18, Q51, Q54, Q58

Keywords

fat tails, climate sensitivity, natural disasters, environmental economics, infectious disease, dismal theorem, existence value

Abstract

We survey the growing literature on fat-tailed distributions in environmental economics. We then examine the theoretical and statistical properties of such distributions, focusing especially on when these properties are likely to arise in environmental problems. We find that a number of variables are fat tailed in environmental economics, including the climate sensitivity, natural disaster impacts, spread of infectious diseases, and stated willingness to pay. We argue that different fat-tailed distributions arise from common pathways. Finally, we review the literature on the policy implications of fat-tailed distributions and controversies over their interpretation. We conclude that the literature has made great strides in demonstrating when fat tails matter for optimal environmental policy. Yet, much is less well understood, including how alternative policies affect fat-tailed distributions, the optimal policy in a computational economy with many fat-tailed problems, and how to account for imprecision in empirical tests for fat tails.

1. INTRODUCTION

The long-run average daily rainfall in Beaumont, Texas in late August is about 0.2 inches. In 48 hours in August 2017, 39.23 inches of rain fell in Beaumont, with 22.33 inches falling on August 30 alone. This staggering amount of rainfall was due to Hurricane Harvey, which made landfall in Texas and Louisiana in late August 2017. Due to its combination of extreme rainfall and landfall in a heavily populated area, Hurricane Harvey caused $125 billion (2017 USD) in damages, on par with Hurricane Katrina in 2005.

Decision making under uncertainty is an essential area of research in many fields across the economics discipline. Such work has long recognized that risk affects optimal decisions. However, risk comes in many forms, including both variability and extreme events. Extreme value theory recognizes that low-probability, extreme events can be meaningful determinants of behavior in many contexts.

In many fields, a natural focus exists on extreme events. It is essential that a building remain standing after a substantial seismic event, that a levee holds after a torrential downpour, and that a portfolio of investments remains solvent after an extreme market shock. Extreme events are becoming more salient to economists and policy makers in the areas of climate change, natural disaster management, and other areas related to the environment.

Here, we review active areas of research within environmental and resource economics that might benefit from continued exploration and incorporation of fat-tailed distributions, in which extreme events are more common than predicted by traditional thin-tailed distributions, such as the normal distribution. We first introduce a working definition of fat-tailed distributions and their properties. We then discuss the empirical challenge of determining whether a data set is drawn from a fat-tailed distribution. Next, we review cases in environmental economics that feature fat-tailed distributions. Climate change is the most explored example, largely owing to Martin Weitzman's work, and we devote much of our review to an exploration of the implications of fat tails for the optimal climate change policy. We also consider natural disasters, infectious diseases, and the preferences for nonmarket environmental amenities, each of which is an active area of research associated with fat tails. We then turn to a description of the mechanisms responsible for fat tails and the implications of fat tails for policy makers. Finally, we highlight a few important, though currently unexplored, questions relating to fat tails in environmental and resource economics.

2. THEORY

2.1. Definition and Properties

A variety of distributions satisfy some notion of a thick upper tail, and terminology varies in the literature. Here, we follow the literature (e.g., Nordhaus 2011, Millner 2013, Conte & Kelly 2018) that defines a distribution that is asymptotically equivalent to a Pareto distribution as fat upper tailed. That is, the probability density eventually declines according to a polynomial.[1]

Fat-upper-tailed distributions have a number of interesting properties. Of course, the defining distinction between fat and thin tails is the frequency of tail events (e.g., Nordhaus 2011). Tail events are extremely rare with a normal distribution. For example, an event larger than 3σ from the mean occurs only 0.13% of the time. In contrast, an event larger than 3σ from the mean for a

[1]Heavy-tailed distributions decline slower than exponential distributions but faster than distributions that decline according to a polynomial. Heavy-tailed distributions such as the log-normal distribution are also common in environmental economics problems (Ott 1990, Kampas & White 2004). For conciseness, heavy-tailed distributions are not part of this review.

Pareto distribution with a shape parameter equal to three occurs with a probability of 1.5%: more than ten times as often. Conte & Kelly (2018) estimate the distribution of US hurricane damages using both a normal distribution and the potentially fat-tailed generalized Pareto distribution (GPD).[2] The best fit normal distribution has a 1% upper quantile of $31.1 billion, but three of the last 71 storms through 2017 had damages exceeding $31.1 billion, indicating that the normal distribution does not fit the upper tail of the damage distribution well. In fact, Conte & Kelly estimate the hurricane damage distribution and reject thin tails.[3]

A second interesting property of fat-tailed distributions is that the largest observation in the data is a poor predictor of subsequent tail events. The mean excess, or the mean conditional on being greater than a threshold, is increasing in the threshold. Therefore, the expected new record largest observation is unlikely to be close to the previous record observation. This provides an immediate intuition as to which data sets are likely generated from a fat-tailed distribution. For example, new records in athletic events such as the 100-m dash are typically close to previous records, indicating a thin-tailed distribution.[4] Similarly, the most intense hurricane on record, Patricia in 2015, had central pressure of 872 hPA (hectopascal pressure units). This beat the previous record of Wilma in 2004 by 1.1% (882 hPA), which in turn was 0.7% lower than the previous record. Hurricane intensity appears to be thin tailed. Yet, interestingly, the most damaging hurricane since 1980 is Katrina, with normalized damages that are 10% greater than the next most damaging hurricane, which is in turn 44% larger than the next most damaging hurricane (Weinkle et al. 2018). We argue below that this property potentially explains some findings in the risk perceptions literature.

A third property of fat-tailed distributions is that some moments are infinite. Suppose a fat-tailed distribution is asymptotically equivalent to a Pareto distribution with tail index k or shape parameter $\frac{1}{k}$. That is, the distribution declines according to a polynomial of degree k. Then, moments greater than or equal to k are infinite. For example, Roe & Baker (2007) and Kelly & Tan (2015) show that the climate sensitivity distribution has tail index $k = 2$, indicating infinite variance. Similarly, Conte & Kelly (2018) show that the hurricane damage distribution also has infinite variance. Practically, upper bounds on the distribution that bound the moments must exist, so an estimated infinite moment just means the data have few or no observations near the bound. Researchers therefore sometimes impose an upper bound (see Section 3.2).

2.2. Statistical Theory

Given the policy implications, it is important for applied researchers to diagnose whether data are drawn from thin-tailed or fat-tailed distributions. The GPD, which nests both thin and fat tails, is particularly attractive here. The Pickands–Balkema–de Haan theorem (Embrechts et al. 2005) shows that if the tail of a distribution is defined as the density greater than a threshold, then the GPD approximates the tail of the distribution to arbitrary accuracy as the threshold increases. If interest is in the tail of the distribution, it is sufficient to estimate a GPD. Still, difficulties arise because, by definition, observations in the tail of the distribution are rare, and most estimators focus on matching the moments of a distribution rather than the tail.

Empirical work can begin with visual inspection of the data. Quantile–quantile plots provide a useful visual comparison between the empirical data and a specified distribution

[2]The GPD is discussed in detail in Section 2.2.

[3]Nordhaus (2011) gives other economic phenomena for which the frequency of tail events makes a normal distribution unlikely, such as stock and oil prices.

[4]We thank Spencer Banzhaf for providing us with this example.

(Das & Resnick 2008). Likewise, a plot of the mean excess of the data that is increasing in the threshold provides evidence of a fat-tailed distribution.

de Zea Bermudez & Kotz (2010) introduce a variety of methods for estimating the parameters of a GPD. Maximum likelihood estimation (MLE) is a common method, despite its many challenges.[5] The method described by Grimshaw (1993) carefully handles several of these known difficulties. A particular challenge exists in that, after normalizing the data, only two parameters describe the entire distribution. Because tail observations are rare, MLE using the entire data set often results in estimates that fit the middle of the distribution better than the tail. Hence, when the interest is in accurate estimation of the tail, researchers use only a subset of the data lying above a chosen threshold. For this reason, the finite-sample properties of MLE, including small-sample bias, have garnered attention in the literature (Cox & Snell 1968, Hosking & Wallis 1987, Firth 1993, Cribari-Neto & Vasconcellos 2002). Giles et al. (2016) provide a method of correcting MLE parameter estimates for small-sample bias.

A second strategy is to estimate the GPD parameters using the mean-excess function, which records the expected value of the difference between a random variable and a selected threshold, conditional on the random variable being larger than the threshold. The mean-excess function is linear in the threshold if and only if the data are drawn from a GPD with a finite mean (Ghosh & Resnick 2010; Embrechts et al. 2005, 2013). If the mean-excess plot is close to linear for large thresholds, then there is no evidence against use of a GPD model (Davison & Smith 1990, Hogg & Klugman 2009, Embrechts et al. 2013).[6] For larger thresholds, few data points exist to estimate the mean excess, and so the estimation must correct for heteroskedasticity. Finally, the slope function of the mean-excess function is not necessarily continuous in the shape parameter, so consistency of this estimator does not follow from consistency of the mean-excess function. Consistency of the slope function is an active area of research (Ghosh & Resnick 2010).

While several empirical methods exist to test whether or not the data generating process is fat tailed, a first-order concern is the number of extreme events available from which to estimate the tail. Nonetheless, in a surprising number of contexts, environmental economists are confronted by fat-tailed data. We now turn to a description of some of these contexts.

3. CLIMATE SENSITIVITY

The best-known example of a fat-tailed distribution in environmental economics is the prior belief distribution of the climate sensitivity. The climate sensitivity is an elasticity that measures the steady-state temperature change from a doubling of CO_2 concentrations above preindustrial levels. The climate sensitivity is a key parameter in integrated assessment models (IAMs) of the climate and economy.[7]

It is well known that the climate sensitivity is uncertain (e.g., Kelly & Kolstad 1999, Weitzman 2009b, Kelly & Tan 2015, Lemoine & Traeger 2016), and one can estimate a prior belief distribution using physical global circulation models (GCMs; e.g., Roe & Baker 2007), the modern temperature and CO_2 records (e.g., Forest et al. 2002, Gregory et al. 2002, Schwartz 2012,

[5] See Castillo & Hadi (1997), Chaouche & Bacro (2006), Del Castillo & Daoudi (2009), and Zhang & Stephens (2009) for rich descriptions of the challenges of using MLE to estimate parameters of the GPD.

[6] However, Ghosh & Resnick (2010) show that the mean-excess plot is inconsistent if the tail index is less than or equal to one.

[7] For example, Kayaratna et al. (2017) simulate the DICE model using draws from a fitted distribution of the climate sensitivity and find the standard deviation of the social cost of carbon in 2020 is \$23.48, or 81% of the mean. Pycroft et al. (2011) obtain a similar result using the PAGE model.

Skeie et al. 2014, Lewis & Curry 2015), more ancient proxy data such as ice cores and tree rings (e.g., Hegerl et al. 2006), or some combination.

In a series of influential papers, Martin Weitzman (2009a,b,c; 2011; 2014), argues that the prior distribution for beliefs about the climate sensitivity is fat tailed, with profound implications for climate policy. Weitzman (2009b) shows that a Pareto prior combined with observational data yields a fat-tailed Student-*t* posterior belief distribution.

Weitzman (2009b) summarizes the policy implications with the so-called dismal theorem. The dismal theorem states that if (*a*) the coefficient of relative risk aversion is positive, (*b*) the consumption growth rate distribution can be normalized into a mean plus a scale parameter times a standardized random variable, (*c*) the scale parameter is uncertain with a Pareto prior with a positive tail index, and (*d*) a finite number of observations exist to estimate the scale parameter, then the stochastic discount factor (SDF) approaches infinity as the rate of substitution between consumption and mortality risk (value of statistical life or VSL) increases.

The SDF is the price of a security that delivers one unit of consumption with certainty in the future. Weitzman draws two policy conclusions. First, Weitzman (2009b,c) argues that conventional cost-benefit analysis is unlikely to yield precise results, because the optimal policy will be highly sensitive to the VSL or any alternative. Second, Weitzman (2009b) argues that potentially catastrophic outcomes drawn from fat-tailed distributions will dominate policy results, being more important than the discount factor and other parameters typically viewed as critical.

Like many of Weitzman's ideas, the sharp result for the SDF was created from a model free from distractions, which makes clear what drives the result. A number of subsequent papers examine whether or not the dismal theorem continues to hold under more complex models and alternative conditions. These, in turn, have led to a deeper understanding of the policy implications of fat-tailed uncertainty. We turn next to several of the most important cases and extensions.

3.1. Is the Climate Sensitivity Really Fat Tailed?

Weitzman (2009b) argues that fat-tailed beliefs arise naturally from a fat-tailed Pareto prior and a series of observations. If so, empirical estimates of the climate sensitivity distribution should have a fat-tailed posterior.

Estimation of the climate sensitivity is difficult. First, many feedback effects and the heat absorption of the deep ocean play out over centuries (Held et al. 2010), a time span that far exceeds the modern temperature record. Second, many variables of interest, including upper and deep ocean temperatures and forcing from aerosols, are either unobserved, imperfectly observed, or have a shorter record than the atmospheric temperature. This seems to support Weitzman's idea that a fat tail is a necessary component of a prior given little or no data and a wide range of theoretical models. Nonetheless, a number of studies try to estimate the climate sensitivity from the modern temperature record (e.g., Forest et al. 2002, Gregory et al. 2002, Schwartz 2012, Skeie et al. 2014, Lewis & Curry 2015). This literature has made significant progress in narrowing the uncertainty, especially in the upper tail (Skeie et al. 2014).

However, given that an ideal data set is lacking (Schwartz 2012), such research makes a number of compromises. First, the literature estimates a transient response over the modern record, typically fewer than 100 years. The transient response is then used to calculate a steady-state response. Because the thickness in the upper tail is primarily related to uncertain long-run feedbacks, the posterior is sensitive to the assumed relationship between transient and long-run feedbacks (Armour et al. 2013).[8]

[8] Skeie et al. (2014) show that explicitly accounting for decadal climate effects also adds to the posterior tail thickness.

Second, the estimation is in some ways more similar to what economics would call calibration, where many parameters and GCM modeling choices (e.g., the spatial resolution and whether features like cloud formation are included) are assumed certain, and other parameters are given assumed prior distributions (e.g., Forest et al. 2002, Huber & Knutti 2012). The posterior distribution is then formed by simulating the GCM over candidate parameter vectors and assessing the posterior probability of each vector by combining the prior probability and the fit of the model output to the data. Such techniques are necessary due to data limitations, but they obscure how the upper tail of the posterior depends on modeling choices in the GCM. Further, most studies use a uniform prior, not the Pareto prior advocated by Weitzman.

Third, the empirical literature typically reports statistics computed from the (nonclosed form) posterior distribution; typically the interest is in the mean and 90% and/or 95% quantiles (e.g., Forest et al. 2002, Lewis & Curry 2018). However, large quantiles may result from either fat tails or thin tails with high variance.[9] To our knowledge, no study has tested the empirical posterior distributions to see if the asymptotic decline is polynomial or exponential, which Weitzman argues is important for policy.

Roe & Baker (2007) and others provide an alternative, theoretical justification using GCMs. Consider a simple climate model in which radiative forcing from CO_2 causes warming, which induces ice melt (reducing the surface albedo of the planet) and other feedbacks that induce further warming. The climate sensitivity is a steady-state statistic and thus includes the cumulative direct and indirect effects of CO_2. Then, thin-tailed feedback uncertainties have a multiplicative effect on the climate sensitivity distribution, generating fat tails in the climate sensitivity distribution.

Although considerable work exists regarding the climate sensitivity distribution, the dismal theorem is ultimately about consumption, which is affected by the climate sensitivity through the damage function. Unfortunately, estimating an implied damage or consumption distribution is difficult, because little to no data exist on the performance of most economies following temperature changes of 7°C or more.

In summary, data limitations prevent estimation of a precise empirical distribution for the climate sensitivity, a problem magnified for the resulting consumption distribution. Whether or not existing studies based on short-/medium-run data are sufficient to declare the consumption distribution thin or fat tailed has unfortunately not been established (Pindyck 2011). Nonetheless, certainly nothing in the empirical literature contradicts Weitzman's characterization of a fat tail derived from diverse GCM model outputs and limited data.

3.2. Bounds on Distributions

We live in a finite world. The dismal theorem is superficially about infinite marginal utility resulting from zero consumption due to an infinitely high temperature that occurs in the infinite future. Thus, many researchers (including Weitzman himself) consider the robustness of the dismal theorem to finite bounds.

A bounded distribution is no longer asymptotically equivalent to a Pareto distribution, and thus no longer fat tailed. Nonetheless, a distribution with a sufficiently high bound is arbitrarily close

[9]Of course, a thin-tailed distribution with a high variance may not fit other quantiles well. Weitzman (2011) chooses the mean and variance of both the normal and Pareto distributions to match the mean and upper 15% quantiles of an average of GCMs considered by the Intergovernmental Panel on Climate Change (IPCC 2007). Given these constraints, the Pareto distribution fits the upper 5% quantile better than the normal distribution. Also, Nordhaus (2011) notes that the average of the quantiles is a biased estimator of the true quantile unless the studies represent draws from an identical population.

to a fat-tailed distribution. Weitzman (2009c) argues that imposing an arbitrary bound makes the optimal policy sensitive to the bound, because the dismal theorem holds approximately for a sufficiently large bound. Ideally then, such bounds should have empirical support, but this is difficult, as any bound is almost certainly outside the modern historical experience.

Pindyck (2011) considers an upper bound on marginal utility. The marginal willingness to pay (WTP) for consumption risk reduction then depends on the likelihood of the disaster, the cost of the risk-reducing activity, risk aversion, and other parameters. Indeed, he constructs examples for which expected marginal WTP is higher for thin-tailed distributions. Still, the issue of sensitivity to the assumed upper bound remains.

Solar energy is finite, and therefore the steady-state temperature must be bounded. Costello et al. (2010) show that the WTP to eliminate all damage is relatively insensitive (3–5% of consumption) to temperature bounds between 20 and 50°C. The key assumption is that damages are calibrated such that an 11°C temperature increase causes a 20% loss of gross domestic product (GDP). If a 20% loss of GDP occurs at 8.4°C, then WTP remains insensitive for risk aversion coefficients less than two.[10] However, Weitzman (2009c, p. 10) argues that even 10–20°C may cause 100% consumption loss ("would effectively destroy planet Earth as we know it"). Millner (2013) also eloquently makes this point, noting that a temperature increase of 85°C would boil blood.

Newbold & Daigneault (2009) consider both a maximum temperature increase and a minimum consumption level. Once the minimum consumption bound becomes active, the temperature bound is irrelevant. Newbold & Daigneault (2009) show that if the lower bound on consumption is about 0.5% or 10^{-7} of current levels, then it is optimal to pay 10% or 99%, respectively, of consumption per year to eliminate all temperature changes. Thus, WTP is sensitive only to very low minimum consumption bounds.

Millner (2013) replaces the constant relative risk aversion (CRRA) utility function with a hyperbolic absolute risk aversion (HARA) function. Utility is then a function of consumption plus a parameter that governs how household risk aversion decreases with consumption. If the HARA parameter is positive, utility and marginal utility are bounded at zero consumption. Imposing HARA utility has a number of advantages over imposing a bound on marginal utility (e.g., Weitzman 2009b, Pindyck 2011). First, HARA utility is fully differentiable and is therefore straightforward to adapt to computational models. Second, the parameter has economic meaning and can therefore be calibrated.

Millner (2013) shows that the SDF and WTP are relatively insensitive to the HARA parameter for values greater than one. Some estimates of the HARA parameter are far greater than one (e.g., Levaggi & Menoncin 2013). However, other estimates find negative values (e.g., Menoncin & Nembrini 2018), and more research is needed to determine whether or not the HARA parameter in this context is in the insensitive region.

Weitzman (2009b) himself considers a lower bound on both utility and marginal utility. The lower bound of consumption can then be calibrated using VSL estimates. Weitzman then shows that the dismal theorem holds to an arbitrarily small approximation error for a sufficiently large VSL. The SDF becomes driven by an uncertain VSL parameter.

[10]In a finite world, climate policy depends on a fat-tailed climate sensitivity indirectly. Thus, parameters and modeling assumptions such as the shape of the damage function, risk aversion, the discount rate, and the heat capacity of the ocean affect the sensitivity of climate policy to the fat tail. For example, Costello et al. (2010) consider alternative damage parameters and risk aversion, Calel et al. (2015) consider alternative damage functions and heat capacity parameters, and Weitzman (2011) considers alternative damage functions.

Finally, Hwang et al. (2013, 2016) consider a bound on the climate sensitivity and show that the optimal near-term carbon tax is a concave function of the variance of the fat-tailed climate sensitivity distribution (and thus the fat tail has a diminishing effect on policy), even for very high bounds.

In conclusion, one can assign lower bounds on consumption loss and utility, through limits on temperature change, utility, damages (Dietz 2011), a finite VSL, or other means. Indeed, one might speculate that human ingenuity and effort will find ways of limiting the impact of even extreme climate change to some large but finite value. Nevertheless, in the DICE model a carbon emissions–free economy currently costs only 7.4% of world GDP (Nordhaus 2017) and is declining. Thus, if a more realistic model with finite bounds indicates it is optimal to pay at least 7.4% of GDP to reduce climate sensitivity risk, then the conclusions are in essence the same as the dismal theorem: Both recommend zero emissions.[11,12] A WTP of about 10% of GDP obtains with a minimum consumption bound of 0.5% of current levels in the study by Newbold & Daigneault (2009) and with a temperature bound of 45°C, with higher damage and risk aversion coefficients in Costello et al.'s (2010) example. If these bounds are realistic,[13] fat tails do indeed have strong climate policy consequences.

3.3. Policy Translation

Under the dismal theorem, the SDF, and therefore marginal WTP (for the first unit of consumption), are infinite. However, abatement is not generally equivalent to a security that delivers one unit of consumption. For example, markets might be incomplete in the sense that a transfer mechanism might not be available for the exact future states of interest: If the world ends, any consumption transferred to the future is presumably lost.

Abatement technologies transfer consumption from the present to the future by reducing future damages from climate change. Because damages are typically modeled as reducing total factor productivity, abatement increases future productivity and therefore consumption. However, as the climate sensitivity is uncertain, so are future damages, and thus so is the ability of abatement to deliver future consumption.

Karp (2009) argues importantly that if it is possible to transfer a certain level of consumption to the future, then future consumption and marginal utility are bounded, and society would never transfer more than half of consumption to the future. He shows that even the weaker argument that potential disasters may be more important for policy than other parameters such as the discount rate is not universally true, at least within the context of a disaster distribution that is thin tailed.

However, abatement, which lowers the steady-state carbon concentrations, reduces damages more when the climate sensitivity is high, and so it flattens the tail rather than transfers certain consumption. Horowitz & Lange (2014) develop conditions on transfer technologies for which the transfer of 100% of consumption is optimal in a fat-tailed environment. If a transfer technology exists that is sufficiently safe, namely, it is able to deliver consumption in sufficiently bad states with sufficiently high probability, then the planner does not optimally transfer all consumption. By transferring a finite amount of consumption, the planner ensures that future consumption is sufficiently likely to be positive, while at the same time avoiding zero current consumption. Conversely, as the probability of the technology transferring consumption

[11]Hwang et al. (2016) also emphasize the importance of emissions control.

[12]Strictly speaking, the dismal theorem might recommend additional measures, such as deployment of (or R&D on) carbon capture to achieve negative emissions and/or geoengineering.

[13]Realistic means that in reality the probability of exceeding the bounds is zero, not that there is a high probability that the bound is the actual climate sensitivity.

in bad states decreases, the planner optimally allocates more and more current consumption to the future, until eventually all consumption is transferred.

Millner (2013) constructs a model in which abatement flattens the tail of the probability density of future temperature states. The optimal abatement is at least the level that flattens the tail enough (depending on damage convexity and risk aversion) so that total WTP is finite. If abatement is sufficiently effective at flattening the tail, then little abatement is required, even if marginal WTP is infinite. But if abatement is sufficiently ineffective, then total WTP becomes infinite.

In conclusion, the relationship between abatement policy and the SDF depends critically on how abatement technologies affect future consumption through the damage function.[14] If abatement technologies satisfy the above conditions, then total WTP is infinite. Yet, whether or not practical abatement technologies satisfy these conditions is less well understood. For example, technologies such as carbon capture and solar geoengineering may shift the consumption distribution differently, especially given that the effectiveness of such technologies is uncertain.

3.4. Learning and Mid-Course Corrections

Given prior beliefs that the true climate sensitivity is drawn from a Pareto distribution, after observing data (or other information), the posterior is the fat-tailed Student-*t* distribution, regardless of the number of observations (Weitzman 2009b). However, after introducing bounds, modeling assumptions, including learning, matter. Furthermore, learning might affect optimal abatement policy by flattening the tail sufficiently so that total WTP falls, even if marginal WTP remains infinite.

Timing also matters. Indeed, Nordhaus (2011) argues that if the climate sensitivity is extremely high, then we should make a mid-course correction after observing rapidly escalating temperatures in the next 50 years. This assumes short-run feedback effects.[15] Allen & Frame (2007) similarly argue that short-run feedbacks, which determine the abatement level that stabilizes the short-run climate, are most important for near-term policy.[16] The ability to make a correction in time to avoid catastrophic temperature changes is a computational question.

Kelly & Tan (2015) consider a DICE-type IAM where the planner is uncertain about a climate feedback parameter but learns over time by observing temperature changes and radiative forcing. As in Roe & Baker's (2007) study, uncertain feedbacks create a fat-tailed prior climate sensitivity distribution. Kelly & Tan show that more abatement is optimal initially with uncertainty and learning, relative to certainty. However, learning flattens the tail of the climate sensitivity distribution considerably in the first few decades, allowing for mid-course corrections until the optimal policy is not much different than the policy under certainty.[17]

Hwang et al. (2017) consider a related model in which the interaction between feedbacks and parameters such as the ocean heat uptake are known and the remaining uncertain feedbacks are incorporated into temperature changes more slowly. In this case, learning about the remaining

[14]Nordhaus (2011) argues similarly that no one relationship exists between the SDF and abatement policy.

[15]Mahadevan & Deutch (2010) show that both uncertain short- and long-run feedbacks generate fat tails, and Roe & Bauman (2013) show that smaller, though significant, short-run uncertainties are more important for abatement policy.

[16]However, Fitzpatrick & Kelly (2017) demonstrate that stabilization of the climate is difficult even in the short run due to the lag between policy actions and climate response.

[17]Bistline (2015) considers a related model in which learning exogenously resolves uncertainty after a fixed number of periods. He also finds that more abatement is optimal initially with uncertainty and learning and that the amount of extra abatement is sensitive to when learning resolves, indicating that mid-course corrections are possible if learning resolves quickly enough.

uncertainty proceeds more slowly, and abatement is less sensitive to beliefs. Nonetheless, the authors also find that the planner can make mid-course corrections, albeit after a longer period of time.

Importantly, in the studies by Hwang et al. (2017) and Fitzpatrick & Kelly (2017), the optimal temperature change is increasing in the climate sensitivity. At a higher climate sensitivity, the abatement cost required to achieve a particular temperature increases, but the benefit of a particular temperature is constant. Mid-course corrections are thus easier, as the planner need not correct the temperature all the way back to the original optimum after learning that the climate sensitivity is higher than previously thought. Indeed, Fitzpatrick & Kelly (2017) show that the optimal target increases with the climate sensitivity, and maintaining a constant target (e.g., 2°C) results in a welfare loss.

In conclusion, learning does not make a fat tail thin. Yet, mid-course corrections are possible. For short-run feedbacks that raise the temperature quickly, learning is relatively fast, and mid-course corrections are possible. Long-run feedbacks raise the temperature more slowly, resulting in slower learning, but the planner has more time to adjust, and so mid-course corrections are still possible.

4. NATURAL DISASTERS

Natural disasters pose significant threats to society through property damage, loss of life, and reduced quality of life. Research shows that a variety of natural disaster–related variables have fat-tailed distributions, including the economic damages from tropical cyclones (Conte & Kelly 2018), the discharge from flood events (Malamud & Turcotte 2006), the magnitude and intensity of volcanic eruptions (Pyle 2015), the volume of landslides (Guzzetti et al. 2009), the area burned by wildfires (Malamud et al. 2005), and the size frequency of near-Earth asteroids (Chapman 2004).

Economists interested in risk amelioration policies and actions have been drawn to the issue of natural disasters, and there is an active literature that continues to explore these issues (e.g., Botzen & van den Bergh 2009, Michel-Kerjan 2010, Bakkensen & Mendelsohn 2016, Bakkensen & Barrage 2017, Sheldon & Zhan 2019). For example, several studies (e.g., Browne & Hoyt 2000, Gallagher 2014) find that the purchase of insurance through the National Flood Insurance Program in the United States is correlated with flood losses during the prior year. Atreya et al. (2013) and Bin & Landry (2013) find decreases in property values in affected communities immediately following significant disasters, though the price differentials fade quickly (4–6 years). Gallagher (2014) finds an uptick in policy purchases in the wake of flood events, although Kousky (2017) notes that most of the observed increase in coverage following hurricanes is due to requirements associated with the acceptance of disaster relief.

These results are somewhat surprising, particularly if damages from such events are drawn from a normal distribution, as there are centuries of experience from which to draw on in forming expectations about exposure to damages from such events. Given a large, normally distributed data set, decision makers should assume a second disaster is highly unlikely and not update beliefs or change decisions significantly. In this context, it is difficult to explain the lack of adequate coverage by private homeowners and changes in property values following flood events. Gallagher (2014) proposes an explanation for this sensitivity of beliefs: Property owners overly discount past events.

Fat tails in the distribution of damages may offer an additional mechanism to explain changes in beliefs and actions following natural disasters. Because tail events are rare and because a fat-tailed distribution has the property that tail events are poor predictors of future tail events, a flood or other disaster can cause a large change in beliefs and decision making, even with hundreds of years of data. Similarly, a period of time without a disaster can also cause beliefs to optimally decline, which can explain the fading issue in the literature (Conte & Kelly 2020).

Born & Viscusi (2006) highlight the challenges that natural disasters pose to private insurers, including bankruptcy risk and the rate increases required to support capital reserves held in case of a tail event. Kousky & Cooke (2009) emphasize that damages from natural disasters are characterized by three challenging features: fat-tailed damages, tail dependence (the tendency of dependence between two random variables to concentrate in the extreme values), and microcorrelations (positive correlations between variables that can be overlooked). Natural disasters can generate damages through multiple pathways (e.g., tropical cyclones can cause damages through high winds, flooding, and storm surge), and ignoring these tail dependencies can lead to underestimation of the damage from catastrophic events. Additionally, correlations between damages caused by natural disasters in different locations are intensified by aggregation, undermining the management of risk exposure via diversification. Kousky & Cooke (2009) use simulations to show that that these features of natural disasters can drive a wedge between the price that private insurers must charge to maintain adequate reserves and the price that homeowners are willing to pay for full insurance coverage, meaning that full insurance is not optimal for many at-risk households. Empirical analysis of the dynamics of private insurance markets for disaster coverage given fat-tailed damages would be a meaningful contribution to this literature.[18]

The fat-tailed nature of damages from natural disasters causes challenges in estimating risk and exposure for households and insurers alike. This challenge is compounded by the possibility that these distributions are nonstationary, due to climate change. The search for time trends in damages caused by natural disasters has been frequently undertaken in the literature, with mixed results. One potential explanation is that such efforts have looked for changes in the mean damage per event across time, typically using ordinary least squares, which might lead trends in extreme events to go unnoticed. Using quantile regression, Coronese et al. (2019) find strong evidence of tail fattening in the damage distribution over time. The presence of a shifting distribution, along with fat tails, creates an additional challenge for insurers, let alone households, to accurately estimate risk exposure.

5. INFECTIOUS DISEASES

Fat tails are an important consideration for researchers and policy makers seeking to understand the spread of pathogens and to prepare for disease arrival. The health impacts of disease depend on the characteristics of both the pathogen (e.g., mortality rate) and the community (e.g., population size). Compartmental epidemiological models are a standard approach to determining the threat posed by infectious diseases (e.g., Brauer 2008, Shaman & Karspeck 2012).

In these epidemiological models, a system of differential equations tracks the dynamics of subsets of the population, with a common class of models focusing on susceptible (not yet infected and susceptible to disease), infected (carrying the disease and infectious), and recovered (neither able to spread the disease nor be reinfected) population segments (SIR models). Traditional SIR models assume a deterministic contact rate, which is one component of the equation determining the spread of the disease from infected to susceptible individuals. These models assume that individuals from various population segments mix uniformly and randomly with each other, with the disease spreading when a susceptible individual interacts with an infected individual.

Assuming uniform contact rates with infected individuals is a modeling convenience that does not align with observed heterogeneity in contact rates, which has substantial impacts on disease dynamics at the population level (Bansal et al. 2007). Outbreaks of gonorrhoea, measles, and SARS

[18] See Conte & Kelly (2020) for an exploration of this issue in Florida between 2003 and 2016.

have been linked to the behavior of small groups of individuals (Hethcote & Yorke 1984, Cent. Dis. Control Prev. 2000, Poutanen et al. 2003). If the distribution of contact rates across time and space is fat tailed, then not only is modeling heterogeneity important for predicting disease spread, but the policy focus immediately shifts to preventing superspreader (tail) events.

For a given probability, the magnitude of extreme events is much greater for fat-tailed distributions relative to thin-tailed distributions. Public-health officials and other policy makers must therefore consider that extreme events play an important role in the efficient response to the disease. Heterogeneous contact rates have been explored using refined SIR models (e.g., Ball et al. 1997, Bjørnstad et al. 2002, Grenfell et al. 2002) and contact network models (e.g., Newman 2002, Pastor-Satorras & Vespignani 2002, Meyers et al. 2005, Shirley & Rushton 2005), which allow for individual-level resolution.

Some network models use scale-free networks, in which contacts follow a power-law distribution, with a small fraction of very highly connected hubs (Barabási & Albert 1999). Hosts in these networks with anomalously high numbers of potential disease-causing contacts are called superspreaders (e.g., Shen et al. 2004, Lloyd-Smith et al. 2005). Such networks, which may have large, or infinite, variance in the degree of contacts, are characterized by low or nonexistent epidemic thresholds, meaning that even very sparse networks are highly vulnerable to epidemics (Bansal et al. 2007).

Although scale-free networks have received substantial attention in the epidemiological literature (e.g., May & Lloyd 2001, Dezso & Barbasi 2002), relatively limited support exists for their relevance in epidemiological contact patterns (Liljeros et al. 2001). Bansal et al. (2007) explore empirical support for the adoption of scale-free networks in different epidemiological contexts and find that, in their systems of study, networks have exponentially distributed contact patterns, which generate epidemiological behavior that is much closer to the predictions of homogeneous-mixing SIR models than the scale-free network models.

Scale-free network models allow for investigation of the impacts of fat tails on epidemiological outcomes, although developing an exact contact network is not practical in large populations. For airborne diseases, it is perhaps more valuable to identify communities in which conditions exist that correspond to fat-tailed distributions of contact rates. Conte et al. (2020) use cellphone data to define a stochastic contact rate in US counties based on visitation rates to places of interest and observe fat tails in this distribution that are positively correlated with cases and deaths due to COVID-19. Wong & Collins (2020) demonstrate fat tails in COVID-19 superspreading events, and Chang et al. (2021) indicate that superspreader points of interest are responsible for a large majority of COVID-19 infections in the United States.

6. PREFERENCES FOR ENVIRONMENTAL AMENITIES

Stated preference methods of valuing nonmarket environmental amenities play a critical role in environmental policy development. For example, species on the brink of extinction are not resource-providing goods with an observable market price, requiring nonmarket valuation methods to estimate their value in pursuit of policies dictating appropriate investment in species preservation. Krutilla (1967) suggests that the existence value of species is a key component of their value, which, if true, implies that stated preference methods are required to recover their full nonmarket valuation.

Fat-tailed WTP values, as manifest in unexpectedly high yes-response rates at the highest available bid, are not an unusual result in the literature. Parsons & Myers (2017) survey 86 contingent valuation studies published in eight of the top environmental economics journals between 1990 and 2015 and find that 60% of the studies include at least one scenario in which the yes-response rate at the greatest bid is at least 20%, and nearly 50% of the studies have at least one scenario in which the yes-response rate at the largest bid is at least 30%. Because the frequency of tail responses exceeds the frequency expected assuming normally distributed responses, researchers sometimes truncate bids, drastically reducing estimates of mean WTP. In response to the challenge posed by fat-tailed WTP distributions for environmental amenities, it is essential that the offered bids allow for well-defined tails of the WTP distribution, which are needed to ensure reliable measures of central tendency of WTP (Haab & McConnell 2002).

McFadden & Leonard (1993) develop a model that maps the compensating variation associated with an increase in the level of public good provision into the responses given by participants in a contingent valuation study with a dichotomous choice format. The authors consider four potential families of distributions to describe consumer tastes for the environmental resource, with the families including the normal (thin tailed), log-normal (heavy tailed), gamma (heavy tailed), and Weibull (fat tailed). In the empirical analysis, the authors find that the assumption of normally distributed preferences for the environmental resource generates the worst fit of the four candidate distribution families in terms of the likelihood criterion. Importantly, the authors find this to be the case for WTP estimates from both dichotomous-choice and open-ended questions, which are not subject to anchoring bias or yea-saying, two issues facing contingent valuation due to the hypothetical nature of the questions that might be explanations for unexpectedly high yes-response rates at the highest bids (Boyle 2003, Parsons & Myers 2017). These results suggest that researchers working with contingent valuation methods might expect relatively high rates of yes-response at extreme bid levels.

In contrast to McFadden & Leonard (1993), Desvousges et al. (1993) find meaningful differences between participants responding to dichotomous-choice versus open-ended questions regarding the frequency of WTP values drawn from the upper tail of the distribution (while these frequencies were in fact quite similar for values in the rest of the distribution). Participants responding to discrete-choice questions had WTP values in the tail of the distribution more often than participants responding to open-ended question. In fact, 34% of dichotomous-choice respondents offered a WTP of $1,000, versus 3% of open-ended respondents. Given these results, the authors conclude that the high frequency of yes responses among dichotomous-choice participants is evidence of anchoring bias: Respondents indicated a WTP that exceeds their true preferences due to the anchor provided by the value mentioned in the question.

An alternative explanation consistent with fat tails is that nonresponse rates are higher for open-ended questions for respondents whose WTP lies in the tail of the distribution. Also, Seller et al. (1985) argue that tail respondents in open-ended questions might shade their responses in order to avoid revealing their divergence from the norm.

Note that these questions about the shape of the preference function for nonmarket environmental amenities also raise questions for value estimation based on choice experiments, which have emerged as an alternative stated preference method to address many of the concerns about contingent valuation (Adamowicz et al. 1994, Hanley et al. 1998). In this context, the concern is that WTP estimates are sensitive to the maximum bid level used in the payment attribute of the choice experiment.

Fat-tailed WTP for environmental amenities has several interesting policy implications. The mean exceeds the median for fat-tailed distributions, so cost values exist such that a majority are not willing to fund the amenity if the cost is shared equally, yet the amenity cost can be supported by placing a greater burden on those with strong preferences for the amenity. Thus, for amenities with meaningful use value, user fees can result in a more efficient allocation than proportional taxes when WTP is fat tailed. For amenities whose value derives primarily from existence, an income tax can result in a more efficient allocation given strong correlation between income and WTP.

A second policy implication concerns the option value of environmental amenities. Option values are increasing in the level of uncertainty. If the WTP for preservation of an environmental amenity is fat tailed, then the difficulty of estimating fat-tailed distributions noted above implies that the tail of the WTP distribution is typically highly uncertain. Thus, with a fat-tailed WTP distribution, the option value of preservation is likely to be substantially higher than the option value given thin-tailed distributions, where uncertain tail events have a negligible effect.

7. COMMON FEATURES

Sections 3–6 show that fat tails arise in a number of seemingly disparate contexts: prior beliefs about the climate sensitivity, damages from natural disasters, the size of networks in the spread of infectious diseases, and WTP for environmental amenities. Yet, a close reading leads to the conclusion that fat-tailed distributions that arise in environmental economics share commonalities, in particular, pathways by which fat tails arise and policy implications. By understanding these common features, we might identify other environmental problems where fat tails are likely relevant.

7.1. Population Centers

The size of population centers follows a fat-tailed distribution known as Zipf's law (see Gabaix 1999 for an explanation). The size of the largest city in the United States, New York, is twice as large as the second largest city, Los Angeles. The second largest city size is a poor predictor of the largest city size, and indeed, Gabaix (1999) finds city size to follow a Pareto distribution with a tail index of approximately one.

Such a population probability density is likely to have profound implications for environmental impacts and environmental policy. Natural disasters that are thin tailed in terms of physical characteristics like hurricane pressure, might, nonetheless, have fat-tailed damages because the damage likely depends on the population size of the area impacted by the disaster (Conte & Kelly 2018). Similarly, network connections are likely much larger in larger population centers reliant on communal modes of transportation. Thus, we might expect a fat-tailed distribution of contact rates, with implications for the outbreak intensity in highly populated areas. One might even expect increases in WTP for environmental amenities in highly populated areas, if survey respondents understand that their contribution to a public good benefits more people in a particular area. Despite the clear relationship between environmental impacts and population size, little work has been done to date understanding this relationship, beyond simply looking at impacts per capita.[19]

Finally, fat tails in a population distribution have important policy consequences. For example, Muller & Mendelsohn (2009) show that a cap-and-trade system that accounts for the distribution of population has optimal prices that vary by orders of magnitude depending on the emissions location.

7.2. Stock/Decay Problem

If x is thin tailed, $1/x$ can be fat tailed. Even if the probability that $x = 0$ approaches zero as x approaches zero, the values of $1/x$ get arbitrarily large, skewing the distribution. This problem is likely to arise in a number of environmental contexts. As noted in Section 3.1, Roe & Baker (2007) show that the climate sensitivity distribution becomes fat tailed because the normally distributed strength of feedback effects have a multiplicative effect on the temperature. Small differences in feedback effects repeated year after year eventually cause large changes in temperature.

[19]For example, how policies such as closures and social distancing affect a fat-tailed contact rate is unclear.

This pathway is possible in other environmental problems. Consider stock pollutants with small but uncertain decay rates, such as heavy metals. Small, thin-tailed uncertainties in the decay rate can lead to fat-tailed uncertainty in the stock and therefore damages. Consider also relatively small uncertainties in the parameters of an infectious disease, such as the time during which an individual is contagious or the number of contacts. Repeated interactions result in much larger uncertainties in the stock of infections after a period of time.

Owing to articles by Weitzman (2009b,c) and the subsequent literature, the policy implications of the denominator problem are better understood. Yet, apparent small uncertainties in feedback parameters in other contexts often lead to overconfidence in both prediction of the stock and the optimality of policy, precisely what Weitzman warned against.

7.3. Correlated Micro Events

Similar to the stock/decay problem, microcorrelations can lead to an aggregate effect that is much larger than expected. Kousky & Cooke (2009) demonstrate how microcorrelations can cause fat tails in aggregate natural disaster damages. Microcorrelations are possible in other contexts. Small correlations in susceptibility to adverse health consequences from pollution emissions or infectious diseases can lead to much larger impacts than might be empirically estimated from a small-scale study.

8. UNDEREXPLORED AREAS

8.1. Multiple Fat Tails

The policy implications of multiple fat tails remains largely unexplored. Nordhaus (2011) argues that several potential catastrophes might be fat tailed, including (prophetically!) emerging diseases, asteroids, and earthquakes, but he provides evidence only for earthquakes and does not discuss how optimal policy might be affected.[20]

Pindyck (2011) and Martin & Pindyck (2015) show that multiple possible catastrophes with fat-tailed distributions reduce expected future consumption. Thus, the marginal value of transferring consumption into the future rises, and so too does WTP. However, a particular catastrophe causes a smaller drop in future consumption because expected future consumption is lower due to all the other catastrophes. For example, if a series of hurricanes is expected, it is also expected that fewer buildings will be left standing to be protected against sea level rise. For a risk aversion coefficient greater than one, the first effect dominates and the existence of a second catastrophe increases WTP for the other. Martin & Pindyck (2015) conduct a simple calibration,[21] which shows that 5 of 7 catastrophes should be averted.

However, Weitzman (2011) argues that other potential disasters might be very damaging but are not necessarily fat tailed. A computational analysis that carefully calibrates other catastrophes with potentially fat-tailed distributions in a capital accumulation model would be especially interesting, because some disasters like pandemics reduce human lives, whereas others like sea level rise primarily destroy capital.

8.2. Alternative Policies

As noted in Section 3.3, considerable effort has been devoted in the literature to understanding the policy implications of a fat-tailed climate sensitivity (Karp 2009, Millner 2013,

[20]Fat tails are only policy relevant if the damage causes fat tails in the consumption distribution.
[21]Interestingly, they predict a 20% chance of a pandemic in 10 years.

Horowitz & Lange 2014). This work mainly looks for conditions under which a theoretical abatement technology might have an infinite total WTP. Still, it remains unclear if practical abatement technologies satisfy these conditions. For example, current abatement technologies can achieve zero emissions at far less than 100% of current GDP. But even with zero emissions, carbon concentrations and temperatures will remain elevated for many years. The tail of the consumption distribution will potentially vary year by year, initially widening as feedback effects continue to raise the temperature and then falling as concentrations return to preindustrial. It is unclear if the year-by-year tail flattening behavior of abatement satisfies the conditions in the literature. Consider also adaptation and emerging technologies such as geoengineering and carbon capture. The tail-flattening behavior of such technologies remains poorly understood, especially how fat tails affect optimal R&D spending on such technologies.

Natural disaster policies in a fat-tailed environment are also a fertile area for further study. It is unclear if policies such as adaptation investments, more strict building codes, and reducing subsidies for coastal development have the ability to flatten tails, and whether fat tails strengthen the case for such policies at the margin or in a more significant way. It does seem clear that natural disaster policies should be focused on vulnerable areas with high population density, as tail disasters are more a function of the vulnerability of the area than the strength of the storm or other disaster (Conte & Kelly 2018).

8.3. Empirical Measurement of Fat Tails

Section 3.1 shows that considerable work exists that estimates the upper quantiles of the climate sensitivity distribution (especially the upper 5% quantile). However, the polynomial decay rate is what defines a fat-tailed distribution, not specifically the value of a particular quantile. For example, a distribution might have a large upper 5% quantile, and yet be thin tailed if the distribution is bounded or has a large variance. Testing specifically for fat tails would help differentiate between these possibilities and provide data-based evidence for the upper bound of the uncertainty distribution.

Another understudied empirical problem is the estimation of the tail of a distribution when data are sparse. In particular, little guidance exists in defining the threshold above which a realization qualifies as extreme. This subset of the full data set is the relevant subset for estimation of the tail index and the shape parameter of the GPD. Using a threshold is important, because the shape parameter and tail index, in addition to determining tail thickness, determine the mean of the distribution. So, if the threshold is set too low, the estimate of the shape parameter may fit the center of the distribution well, but poorly fit the upper tail of the data. Setting the threshold too high means that only a handful of data points exist from which to estimate the parameters, resulting in noisy estimates that may not distinguish between fat tails and thin tails. While researchers in certain fields (e.g., finance, epidemiology) often have millions of observations and may not be concerned about setting too high of a threshold, natural disaster researchers may have only a few hundred observations in the full data set. A data-based threshold would allow the literature to move beyond simply examining alternative thresholds as a robustness check.

Though reliable data exist to measure the distribution of the climate sensitivity, many of the other variables described above must contend with an additional complexity in identifying fat tails: measurement error. Natural disaster damages, infections and resultant mortality of infectious diseases, and WTP for nonmarket environmental amenities are all measured with error in a way that challenges confidence in the presence of fat tails. Correlations between the intensity of the disaster or infectious disease and the resultant damages might lead to incorrect classification of a distribution as thin or fat tailed. Furthermore, the truncation point in the stated preference survey might have a similar effect. Additional research on the implications of measurement error on the presence of fat tails and on the methods currently available to identify fat tails would be a valuable contribution to the literature.

8.4. Disparate Impacts

Finally, fat-tailed impacts from natural disasters are unlikely to be distributed evenly across the population. First, because reserve costs drive insurance costs above actuarially fair levels, low-income property owners are less likely to be insured (Picard 2008). Second, low-income residents can less likely afford risk-reducing investments.[22] An interesting question is to what extent fat-tailed dollar damages underestimate the welfare cost when borne by residents with a high marginal utility of consumption.

9. CONCLUSIONS

The work of Martin Weitzman has sparked a growing literature exploring the implication of fat tails for climate policy, and we have presented several other areas of active research by environmental economists, including natural disasters, infectious diseases, and WTP for environmental amenities. We argue that fat-tailed distributions of these variables result from similar pathways such as fat tails in the population distribution and from uncertainties in the long-run decay rates of stock pollutants.

Subsequent to the work of Weitzman, a developing literature has focused on understanding the implications of fat tails. Considerable progress has been made understanding the conditions under which fat tails have strong policy consequences. However, many policy implications remain unexplored, including how alternative policies affect fat-tailed distributions, a computational analysis of an economy with multiple fat tails, and imprecision in empirical tests for fat tails resulting from sparse data and other issues. Nonetheless, other than the paucity of data on extreme events, few theoretical or empirical barriers exist to answering these questions. We expect that research in these areas will continue to expand to answer these important questions.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We would like to thank Christopher Costello and Charles D. Kolstad for useful comments and suggestions. D.L.K. would like to thank Martin Weitzman for many deep conversations about fat tails, every one of which led me to new insights and a better understanding of the issue. M.N.C. is grateful for financial support from the National Science Foundation (award number 1924378).

LITERATURE CITED

Adamowicz W, Louviere J, Williams M. 1994. Combining revealed and stated preference methods for valuing environmental amenities. *J. Environ. Econ. Manag.* 26(3):271–92

Allen MR, Frame DJ. 2007. Call off the quest. *Science* 318(5850):582–83

[22] See Kousky (2019) for a review of the empirical literature on the role of natural disaster insurance in improving resilience to these events.

Armour KC, Bitz CM, Roe GH. 2013. Time-varying climate sensitivity from regional feedbacks. *J. Clim.* 26(13):4518–34

Atreya A, Ferreira S, Kriesel W. 2013. Forgetting the flood? an analysis of the flood risk discount over time. *Land Econ*. 89(4):577–96

Bakkensen LA, Barrage L. 2017. *Flood risk belief heterogeneity and coastal home price dynamics: Going under water?* NBER Work. Rep. 23854. **https://www.nber.org/papers/w23854**

Bakkensen LA, Mendelsohn RO. 2016. Risk and adaptation: evidence from global hurricane damages and fatalities. *J. Assoc. Environ. Resour. Econ.* 3(3):555–87

Ball F, Mollison D, Scalia-Tomba G. 1997. Epidemics with two levels of mixing. *Ann. Appl. Probab.* 7(1):46–89

Bansal S, Grenfell BT, Meyers LA. 2007. When individual behaviour matters: homogeneous and network models in epidemiology. *J. R. Soc. Interface* 4(16):879–91

Barabási AL, Albert R. 1999. Emergence of scaling in random networks. *Science* 286(5439):509–12

Bin O, Landry CE. 2013. Changes in implicit flood risk premiums: empirical evidence from the housing market. *J. Environ. Econ. Manag.* 65(3):361–76

Bistline JE. 2015. Fat-tailed uncertainty, learning, and climate policy. *Clim. Change Econ.* 6(2). **https://doi.org/10.1142/S2010007815500098**

Bjørnstad ON, Finkenstädt BF, Grenfell BT. 2002. Dynamics of measles epidemics: estimating scaling of transmission rates using a time series SIR model. *Ecol. Monogr.* 72(2):169–84

Born P, Viscusi WK. 2006. The catastrophic effects of natural disasters on insurance markets. *J. Risk Uncertain.* 33(1–2):55–72

Botzen WW, van den Bergh JC. 2009. Bounded rationality, climate risks, and insurance: Is there a market for natural disasters? *Land Econ*. 85(2):265–78

Boyle K. 2003. Contingent valuation in practice. In *A Primer on Non-Market Valuation*, ed. PA Champ, KJ Boyle, TC Brown, pp. 111–70. Dordrecht, Neth.: Kluwer Acad.

Brauer F. 2008. Compartmental models in epidemiology. In *Mathematical Epidemiology*, ed. F Brauer, P van den Driessche, J Wu, pp. 19–79. Berlin: Springer

Browne MJ, Hoyt RE. 2000. The demand for flood insurance: empirical evidence. *J. Risk Uncertain.* 20(3):291–306

Calel R, Stainforth DA, Dietz S. 2015. Tall tales and fat tails: the science and economics of extreme warming. *Clim. Change* 132(1):127–41

Castillo E, Hadi AS. 1997. Fitting the generalized pareto distribution to data. *J. Am. Stat. Assoc.* 92(440):1609–20

Cent. Dis. Control Prev. 2000. Measles outbreak—Netherlands, April 1999-January 2000. *Morb. Mortal. Wkly. Rep.* 49(14):299–303

Chang S, Pierson E, Koh PW, Gerardin J, Redbird B, et al. 2021. Mobility network models of COVID-19 explain inequalities and inform reopening. *Nature* 589(7840):82–87

Chaouche A, Bacro JN. 2006. Statistical inference for the generalized Pareto distribution: maximum likelihood revisited. *Commun. Stat. Theory Methods* 35(5):785–802

Chapman CR. 2004. The hazard of near-earth asteroid impacts on earth. *Earth Planet. Sci. Lett.* 222(1):1–15

Conte MN, Gordon M, Sims CB. 2020. *Fat-tailed contact rate distribution linked to more intense COVID-19 outbreak*. Work. Pap., Fordham Univ., New York

Conte MN, Kelly DL. 2018. An imperfect storm: Fat-tailed tropical cyclone damages, insurance, and climate policy. *J. Environ. Econ. Manag.* 92:677–706

Conte MN, Kelly DL. 2020. *How beliefs change following a natural disaster: evidence from wind insurance data*. Work. Pap., Univ. Miami

Coronese M, Lamperti F, Keller K, Chiaromonte F, Roventini A. 2019. Evidence for sharp increase in the economic damages of extreme natural disasters. *PNAS* 116(43):21450–55

Costello CJ, Neubert MG, Polasky SA, Solow AR. 2010. Bounded uncertainty and climate change economics. *PNAS* 107(18):8108–10

Cox DR, Snell EJ. 1968. A general definition of residuals. *J. R. Stat. Soc. B* 30(2):248–65

Cribari-Neto F, Vasconcellos KL. 2002. Nearly unbiased maximum likelihood estimation for the beta distribution. *J. Stat. Comput. Simul.* 72(2):107–18

Das B, Resnick SI. 2008. Qq plots, random sets and data from a heavy tailed distribution. *Stoch. Models* 24(1):103–32
Davison AC, Smith RL. 1990. Models for exceedances over high thresholds. *J. R. Stat. Soc. B* 52(3):393–425
de Zea Bermudez P, Kotz S. 2010. Parameter estimation of the generalized Pareto distribution—part I. *J. Stat. Plan. Inference* 140(6):1353–73
Del Castillo J, Daoudi J. 2009. Estimation of the generalized Pareto distribution. *Stat. Probab. Lett.* 79(5):684–88
Desvousges WH, Johnson FR, Dunford RW, Hudson SP, Wilson KN, Boyle KJ. 1993. Measuring natural resource damages with contingent valuation: tests of validity and reliability. In *Contingent Valuation: A Critical Assessment*, ed. JA Hausman, pp. 91–164. Bingley, UK: Emerald Group
Dezso Z, Barbasi A. 2002. Halting viruses in scale-free networks. *Phys. Rev. E* 65(5):055103
Dietz S. 2011. High impact, low probability? An empirical analysis of risk in the economics of climate change. *Clim. Change* 108(3):519–41
Embrechts P, Klüppelberg C, Mikosch T. 2013. *Modelling Extremal Events for Insurance and Finance*. New York: Springer Sci. & Bus.
Embrechts P, McNeal A, Frey R. 2005. *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton, NJ: Princeton Univ. Press
Firth D. 1993. Bias reduction of maximum likelihood estimates. *Biometrika* 27–38
Fitzpatrick LG, Kelly DL. 2017. Probabilistic stabilization targets. *J. Assoc. Environ. Resour. Econ.* 4(2):611–57
Forest CE, Stone PH, Sokolov AP, MR Allen, Webster MD. 2002. Quantifying uncertainties in climate system properties with the use of recent observations. *Science* 295:113–17
Gabaix X. 1999. Zipf's law for cities: an explanation. *Q. J. Econ.* 114(3):739–67
Gallagher J. 2014. Learning about an infrequent event: evidence from flood insurance take-up in the United States. *Am. Econ. J. Appl. Econ.* 206–33
Ghosh S, Resnick S. 2010. A discussion on mean excess plots. *Stoch. Proc. Appl.* 120(8):1492–517
Giles DE, Feng H, Godwin RT. 2016. Bias-corrected maximum likelihood estimation of the parameters of the generalized Pareto distribution. *Commun. Stat. Theory Methods* 45(8):2465–83
Gregory JM, Stouffer RJ, Raper SC, Stott PA, Rayner NA. 2002. An observationally based estimate of the climate sensitivity. *J. Clim.* 15:3117–21
Grenfell BT, Bjørnstad ON, Finkenstädt BF. 2002. Dynamics of measles epidemics: scaling noise, determinism, and predictability with the TSIR model. *Ecol. Monogr.* 72(2):185–202
Grimshaw SD. 1993. Computing maximum likelihood estimates for the generalized Pareto distribution. *Technometrics* 35(2):185–91
Guzzetti F, Ardizzone F, Cardinali M, Rossi M, Valigi D. 2009. Landslide volumes and landslide mobilization rates in Umbria, central Italy. *Earth Planet. Sci. Lett.* 279(3–4):222–29
Haab TC, McConnell KE. 2002. *Valuing Environmental and Natural Resources: The Econometrics of Non-Market Valuation*. Cheltenham, UK: Edward Elgar
Hanley N, Wright RE, Adamowicz V. 1998. Using choice experiments to value the environment. *Environ. Resour. Econ.* 11(3–4):413–28
Hegerl GC, Crowley TJ, Hyde WT, Frame DJ. 2006. Climate sensitivity constrained by temperature reconstructions over the past seven centuries. *Nature* 440:1029–32
Held IM, Winton M, Takahashi K, Delworth T, Zeng F, Vallis GK. 2010. Probing the fast and slow components of global warming by returning abruptly to preindustrial forcing. *J. Clim.* 23(9):2418–27
Hethcote HW, Yorke JA. 1984. *Gonorrhea Transmission Dynamics and Control*. Lect. Notes Biomath. Vol. 56. Berlin: Springer Verlag
Hogg RV, Klugman SA. 2009. *Loss Distributions*. Hoboken, NJ: John Wiley & Sons
Horowitz J, Lange A. 2014. Cost-benefit analysis under uncertainty-a note on Weitzman's dismal theorem. *Energy Econ*. 42:201–3
Hosking JR, Wallis JR. 1987. Parameter and quantile estimation for the generalized Pareto distribution. *Technometrics* 29(3):339–49
Huber M, Knutti R. 2012. Anthropogenic and natural warming inferred from changes in Earth's energy balance. *Nat. Geosci.* 5(1):31–36

Hwang IC, Reynès F, Tol RS. 2013. Climate policy under fat-tailed risk: an application of dice. *Environ. Resour. Econ.* 56(3):415–36

Hwang IC, Reynès F, Tol RS. 2017. The effect of learning on climate policy under fat-tailed risk. *Resour. Energy Econ.* 48:1–18

Hwang IC, Tol RSJ, Hofkes MW. 2016. Fat-tailed risk about climate change and climate policy. *Energy Policy* 89:25–35

IPCC (Intergov. Panel Clim. Change). 2007. *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the IPCC*, ed. S Solomon, D Qin, M Manning, M Marquis, K Averyt, et al. Cambridge, UK: Cambridge Univ. Press

Kampas A, White B. 2004. Administrative costs and instrument choice for stochastic non-point source pollutants. *Environ. Resour. Econ.* 27:109–33

Karp LS. 2009. *Sacrifice, discounting and climate policy: five questions*. CESifo Work. Pap. 2761. **https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1458887**

Kayaratna K, McKitrick R, Kreutzer D. 2017. Empirically constrained climate sensitivity and the social cost of carbon. *Clim. Change Econ.* 8(2):1750006

Kelly DL, Kolstad CD. 1999. Bayesian learning, pollution, and growth. *J. Econ. Dyn. Control* 23(4):491–518

Kelly DL, Tan Z. 2015. Learning and climate feedbacks: optimal climate insurance and fat tails. *J. Environ. Econ. Manag.* 72:98–122

Kousky C. 2017. Disasters as learning experiences or disasters as policy opportunities? Examining flood insurance purchases after hurricanes. *Risk Anal.* 37(3):517–30

Kousky C. 2019. The role of natural disaster insurance in recovery and risk reduction. *Annu. Rev. Resour. Econ.* 11:399–418

Kousky C, Cooke RM. 2009. *The unholy trinity: fat tails, tail dependence, and micro-correlations*. Tech. Rep. 09-36, Resour. Future, Washington, DC

Krutilla JV. 1967. Conservation reconsidered. *Am. Econ. Rev.* 57(4):777–86

Lemoine D, Traeger CP. 2016. Economics of tipping the climate dominoes. *Nat. Clim. Change* 6(5):514–19

Levaggi R, Menoncin F. 2013. Optimal dynamic tax evasion. *J. Econ. Dyn. Control* 37:2157–67

Lewis N, Curry J. 2018. The impact of recent forcing and ocean heat uptake data on estimates of climate sensitivity. *J. Clim.* 31(15):6051–71

Lewis N, Curry JA. 2015. The implications for climate sensitivity of ar5 forcing and heat uptake estimates. *Clim. Dyn.* 45(3–4):1009–23

Liljeros F, Edling C, Amaral L, Stanley H, Aberg Y. 2001. The web of human sexual networks. *Nature* 411:907–8

Lloyd-Smith JO, Schreiber SJ, Kopp PE, Getz WM. 2005. Superspreading and the effect of individual variation on disease emergence. *Nature* 438(7066):355–59

Mahadevan L, Deutch JM. 2010. Influence of feedback on the stochastic evolution of simple climate systems. *Proc. R. Soc. A* 466(2116):993–1003

Malamud BD, Millington JD, Perry GL. 2005. Characterizing wildfire regimes in the United States. *PNAS* 102(13):4694–99

Malamud BD, Turcotte DL. 2006. The applicability of power-law frequency statistics to floods. *J. Hydrol.* 322(1–4):168–80

Martin IW, Pindyck RS. 2015. Averting catastrophes: the strange economics of Scylla and Charybdis. *Am. Econ. Rev.* 105(10):2947–85

May RM, Lloyd AL. 2001. Infection dynamics on scale-free networks. *Phys. Rev. E* 64(6):066112

McFadden D, Leonard G. 1993. Issues in the contingent valuation of environmental goods: Methodologies for data collection and analysis. In *Contingent Valuation: A Critical Assessment*, ed. JA Hausman, pp. 165–215. New York: North-Holland

Menoncin F, Nembrini S. 2018. Stochastic continuous time growth models that allow for closed form solutions. *J. Econ.* 124(3):213–41

Meyers LA, Pourbohloul B, Newman ME, Skowronski DM, Brunham RC. 2005. Network theory and SARS: predicting outbreak diversity. *J. Theor. Biol.* 232(1):71–81

Michel-Kerjan EO. 2010. Catastrophe economics: the national flood insurance program. *J. Econ. Perspect.* 24(4):165–86

Millner A. 2013. On welfare frameworks and catastrophic climate risks. *J. Environ. Econ. Manag.* 65(2):310–25

Muller NZ, Mendelsohn R. 2009. Efficient pollution regulation: getting the prices right. *Am. Econ. Rev.* 99(5):1714–39

Newbold SC, Daigneault A. 2009. Climate response uncertainty and the benefits of greenhouse gas emissions reductions. *Environ. Resour. Econ.* 44(3):351–77

Newman ME. 2002. Spread of epidemic disease on networks. *Phys. Rev. E* 66(1):016128

Nordhaus WD. 2011. The economics of tail events with an application to climate change. *Rev. Environ. Econ. Policy* 5(2):240–57

Nordhaus WD. 2017. Revisiting the social cost of carbon. *PNAS* 114(7):1518–23

Ott W. 1990. A physical explanation of the lognormality of pollution concentrations. *J. Air Waste Manag. Assoc.* 40:1378–83

Parsons G, Myers K. 2017. Fat tails and truncated bids in contingent valuation: an application to an endangered shorebird species. In *Contingent Valuation of Environmental Goods*, ed. D McFadden, K Train, pp. 17–42. Cheltenham, UK: Edward Elgar

Pastor-Satorras R, Vespignani A. 2002. Epidemic dynamics in finite size scale-free networks. *Phys. Rev. E* 65(3):035108

Picard P. 2008. Natural disaster insurance and the equity-efficiency trade-off. *J. Risk Insur.* 75(1):17–38

Pindyck RS. 2011. Fat tails, thin tails, and climate change policy. *Rev. Environ. Econ. Policy* 5:258–74

Poutanen SM, Low DE, Henry B, Finkelstein S, Rose D, et al. 2003. Identification of severe acute respiratory syndrome in Canada. *N. Engl. J. Med.* 348(20):1995–2005

Pycroft J, Vergano L, Hope C, Paci D, Ciscar JC. 2011. A tale of tails: uncertainty and the social cost of carbon dioxide. *Economics* 5:2011–22

Pyle DM. 2015. Sizes of volcanic eruptions. In *The Encyclopedia of Volcanoes*, ed. H Sigurdsson, pp. 257–64. Amsterdam: Academic

Roe GH, Baker MB. 2007. Why is climate sensitivity so unpredictable? *Science* 318(5850):629–32

Roe GH, Bauman Y. 2013. Climate sensitivity: Should the climate tail wag the policy dog? *Clim. Change* 117(4):647–62

Schwartz SE. 2012. Determination of earths transient and equilibrium climate sensitivities from observations over the twentieth century: strong dependence on assumed forcing. *Surv. Geophys.* 33(3–4):745–77

Seller C, Stoll JR, Chavas JP. 1985. Validation of empirical measures of welfare change: a comparison of non-market techniques. *Land Econ.* 61(2):156–75

Shaman J, Karspeck A. 2012. Forecasting seasonal outbreaks of influenza. *PNAS* 109(50):20425–30

Sheldon TL, Zhan C. 2019. The impact of natural disasters on US home ownership. *J. Assoc. Environ. Resour. Econ.* 6(6):1169–203

Shen Z, Ning F, Zhou W, He X, Lin C, et al. 2004. Superspreading SARS events, Beijing, 2003. *Emerg. Infect. Dis.* 10(2):256–60

Shirley MD, Rushton SP. 2005. The impacts of network topology on disease spread. *Ecol. Complex.* 2(3):287–99

Skeie RB, Berntsen T, Aldrin M, Holden M, Myhre G. 2014. A lower and more constrained estimate of climate sensitivity using updated observations and detailed radiative forcing time series. *Earth Syst. Dyn.* 5(1):139–75

Weinkle J, Landsea C, Collins D, Musulin R, Crompton RP, et al. 2018. Normalized hurricane damage in the continental United States 1900–1917. *Nat. Sustain.* 1:808–13

Weitzman ML. 2009a. Additive damages, fat-tailed climate dynamics, and uncertain discounting. *Economics* 3:2009–39

Weitzman ML. 2009b. On modeling and interpreting the economics of catastrophic climate change. *Rev. Econ. Stat.* 91(1):1–19

Weitzman ML. 2009c. *Some basic economics of extreme climate change*. Work. Pap., Harvard Univ., Cambridge, MA. **http://scholar.harvard.edu/files/weitzman/files/basiceconomicsclimatechange.pdf?m=1360041775**

Weitzman ML. 2011. Fat-tailed uncertainty in the economics of catastrophic climate change. *Rev. Environ. Econ. Policy* 5(2):275–92

Weitzman ML. 2014. Fat tails and the social cost of carbon. *Am. Econ. Rev.* 104(5):544–46

Wong F, Collins JJ. 2020. Evidence that coronavirus superspreading is fat-tailed. *PNAS* 117(47):29416–18

Zhang J, Stephens MA. 2009. A new and efficient estimation method for the generalized Pareto distribution. *Technometrics* 51(3):316–25