

Linking Risk Models to Microeconomic Indicators

Jesse Anttila-Hughes

Mohan Sharma



WORLD BANK GROUP

Finance and Markets Global Practice Group

&

Global Facility for Disaster Reduction and Recovery

July 2015

Abstract

Catastrophe risk models are quantitative models used to estimate probabilistic loss distributions for a specified range of assets subject to a baseline level of disaster risk. While cat risk models are used extensively by the insurance and reinsurance industry to estimate expected losses to insured assets, their ability to estimate damages outside of a narrow range of physical assets such as buildings or infrastructure is still limited. This paper first provides a brief outline of cat risk models as they currently exist, and then outlines the major econometric issues involved in incorporating research from the growing literature on

the microeconomic impacts of disasters into a cat model framework. Attention is specifically drawn to issues arising from the generally low recurrence frequencies of disasters, the likely role of difficult-to-document indirect damages in influencing total disaster costs, and issues related to generalizing disaster response functions across different domains. The paper ends by noting the large discrepancy between the current state of the literature on disaster impacts on microeconomic indicators and the level needed for adequate cat risk model performance, and suggests means of closing that gap as well as potential areas for future research.

This paper is a product of the Disaster Risk Financing and Insurance Program (DRFIP), a partnership of the World Bank's Finance and Markets Global Practice Group and the Global Facility for Disaster Reduction and Recovery, with funding from the UK Department For International Development. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The authors may be contacted at jesse.anttilahughes@gmail.com and 2mohan.sharma@gmail.com.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

LINKING RISK MODELS TO MICROECONOMIC INDICATORS¹

Jesse Anttila-Hughes and Mohan Sharma

JEL Classification:C180 Methodological Issues: General; Q54 Climate; Natural Disasters; Global Warming; O12 Microeconomic Analyses of Economic Development

Keywords: Applied econometrics; Climate Change; Natural Disasters; Risk Modeling; Insurance

¹ This research has been completed as part of a series of papers developed under the Sovereign Disaster Risk Financing and Insurance Impact Appraisal project. The project is led by the World Bank Disaster Risk Financing and Insurance Program and funded through a cooperation between the UK Department for International Development (DFID) and the Global Facility for Disaster Reduction and Recovery (GFDRR).

Section I: Overview of Catastrophe Risk Models

Catastrophe risk, or "cat risk", models are quantitative models used to generate estimates of expected insurance losses due to catastrophic events. A cat model's primary goal is to generate probabilistic loss distributions for assets that may be exposed to possible future events. Plausible events are drawn from a stochastic event set, which is derived from the historical (instrumental and non-instrumental) record of severity and frequency of modeled perils, and is typically complemented

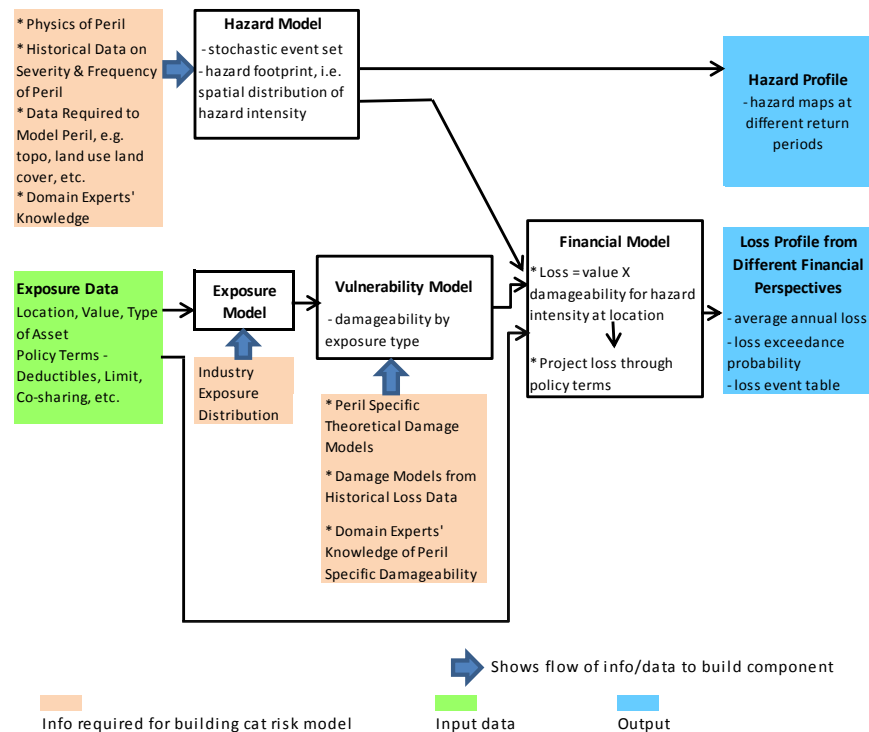


Figure 1. Graphical Representation of Cat Risk Models

by expert knowledge where historical data are insufficient to constrain the severity-frequency distribution. The stochastic event set addresses the issue of where events are likely to occur and if so what the likely magnitude and frequency of events will be, and may include change in distribution of future events as apt for some of the hazards with nonstationary distributions, e.g., for climatic hazards under climate change. Each event in the stochastic event set is associated with (i) an annual rate of occurrence and (ii) a spatial footprint, meaning the spatial distribution of the intensity of the disaster being modeled, e.g., wind speed in a tropical cyclone or peak ground acceleration for an earthquake. The spatial footprint of an event is modeled using the physics of the disaster, and generally requires the expertise of relevant natural and physical scientists. The footprint is then used to determine the expected exposure of assets in the exposure set, and a vulnerability module then yields the specific damages that can be expected for a certain severity of event. Vulnerability modules are heavily reliant on engineering and technical expertise, and are informed by a variety of sources ranging from building codes to topography to peer reviewed research; a specific vulnerability function in a vulnerability module might map cyclone wind speed intensity onto expected insurable damages for exposed commercial buildings built after a certain key building code item was passed, for example. As the vulnerability module is the part of the cat risk model that directly maps hazard outcomes onto socially relevant damages, it is the most natural place to expand cat risk models to include other potentially relevant outcomes, e.g., damages to microeconomic indicators.

Regardless of the type of damages yielded by the vulnerability module, total event loss for a given hazard intensity is calculated by summing the losses across all affected assets in the exposure. The probability

distribution of the loss can be calculated by using the estimated event loss and annual occurrence rate data for each event intensity in the stochastic event set. Once the loss probability distribution is estimated, various metrics of risk can be evaluated as for instance the expected loss, the losses at various return periods, etc. The fat tail of cat loss distributions and the covariate nature of the cat losses generally preclude the use of standard actuarial techniques of projecting historical losses. The usefulness of cat risk models thus lies in the fact that cat risk models provide objective, scientific methodologies for estimating loss probability distributions based on today's exposures. The losses calculated by cat models are losses arising from damages to physical infrastructure (houses, bridges, roads, etc.), human capital (as in injuries and fatalities), crops, livestock, etc. and are denoted as direct losses. These models are calibrated and validated using historical loss data when available. When properly calibrated and validated, cat risk models represent the best that can be done in modeling disaster loss distributions.

Cat risk models are relatively new (with the first national cat model released in the late 1980s) and data are still sufficiently sparse that major changes to model functionality and methodology are not uncommon after major disasters even today. It is important to note that as such cat risk models are still evolving, and can be expected to continue to change as they find wider use in (re)insurance markets for pricing, risk accumulation, reinsurance purchase, portfolio optimization, and the like. (See Grossi & Kunreuther, 2005 for an overview.)

Section II: Econometric Issues Related to Estimating the Impact of Disasters on Microeconomic Indicators

In recent years the economics and related social science literatures have begun generating a large number of estimates of the impacts of natural hazards on social outcomes (see Scott and Shepherd, 2014). In order to effectively model disaster damages to microeconomic indicators, cat models must be able to incorporate these microeconometrically-derived estimates of hazard impacts on specific indicators into existing damage estimation frameworks. We may conceptualize this task by noting that in a standard cat risk model (e.g., as diagrammed in Figure 1), exposure maps onto context-specific vulnerability functions in a vulnerability module, which generate output interpretable by the financial model. In order for cat risk models to generate estimates of damages to microeconomic indicators, they must thus generate a new suite of vulnerability functions that output damages to indicators, which may then be interpreted by analysts directly rather than fed into a costs model.

Updating existing risk models to generate estimates for microeconomic losses thus involves two distinct tasks: (1) generating context-specific vulnerability functions that generate probabilistic damage estimates for indicators of interest based on hazard incidence based on microeconomic estimates of hazard damages to those indicators and (2) incorporating data on the demographic and economic characteristics of exposed household populations by context to generate exposure estimates. The main econometric issues of interest revolve around the former: constructing appropriate vulnerability functions involves taking econometric estimates of micro-level damages caused by specific disaster outcomes in specific locations and aggregating them up into meaningful general form relationships between disaster incidence and damages in a specific context of interest. Identifying at-risk populations and their spatial distribution (and hence exposure to disaster risk) is largely a data concern, though we briefly outline potential issues with that below as well.

Estimates of disaster damages to social outcomes which could be used to generate meaningful vulnerability functions are generally hampered by two major econometric issues. The first is internal validity, or the extent to which damages attributed to a disaster outcomes are actually the result of the disaster per se, or merely a proxy for other variables (e.g., sorting by poorer segments of society into less desirable disaster-prone areas). The second is external validity, or the extent to which assessed impacts can be generalized out of sample, either to other contexts or locations, to other time periods (and resulting different institutional and development profile), or to disaster magnitudes or outcomes which are not directly estimated in the data. Internal validity of estimates needs to be pursued in order to infer that loss functions subject to a given hazard treatment are being estimated without bias; external validity of estimates needs to be pursued to ensure that behaviors inferred for non-studied contexts are meaningful.

Internal Validity: Concerns about internal validity in estimating vulnerability functions for specific microeconomic indicators echo general concerns about estimating causally interpretable² quantitative relationships between variables. Briefly, we may conceptualize the problem of estimating a damage function for a microeconomic indicator of interest as recovering the function

$$(1) \quad y = D(h, s)$$

where y is our outcome variable of interest (perhaps indexed to a specific subregion, demographic, time period, etc.), h is a measure of hazard outcome intensity, and s is a vector of relevant demographic or economic characteristics for an exposed microeconomic unit, e.g., a household. In an ideal experiment, we would recover the functional form of $D(\cdot)$ by randomly assigning households (or groups of households, if concerned about aggregate or equilibrium impacts) different outcomes of h , and observing how y varies both unconditionally as well as conditionally on s . Random assignment would allow us to obviate concerns about potentially omitted variables which may be biasing our estimates of $D(\cdot)$, and with sufficiently large sample size a general form relationship for a given household of type s could be recovered. The real world is, clearly, a far cry from the ideal experiment (which would of course not be defensible on ethical grounds even if it were physically plausible), and delineating these differences allows us to understand what is and is not econometrically feasible in current research.

It is first important to note that the task of econometrically estimating disaster losses is greatly aided by the generally unexpected, and hence plausibly random, arrival of disasters. While most disasters have well known spatial distributions which are correlated with a variety of local conditions, and are hence non-random, the *timing* of disaster events is generally stochastic, meaning that econometric estimates that rely on panel data capturing variation in disaster exposure and microeconomic indicators *within* an area of interest can generally be considered internally valid subject to a host of concerns, most notably including spatial and temporal fixed effects within panels to capture covariates constant within a relevant catchment area or period³. Multiple threats to internal validity nonetheless exist, and we outline them briefly below.

² A meaningful discussion of econometrics' current understanding of causal inference is beyond the scope of this paper. For an in-depth treatment, please see Holland, 1986, or Angrist and Pischke, 2008.

³ For a more in-depth discussion of the identification advantages of combining panel data with stochastically varying natural phenomena, particularly the climate, see Dell, Jones, and Olken, 2013.

The primary threat to the internal validity of empirical derived relationships between hazard incidence and microeconomic outcomes is the low interoccurrence time of most hazards. This presents two fundamental issues. The first is that many hazards which are of economic interest occur too infrequently to show up in modern datasets; this may be partly due to specific hazards being relatively rare (e.g., tsunamis), or because particularly extreme outcomes of certain hazard types are themselves rare (e.g., 1 in 100 year storms). Data availability issues may be further exacerbated by local institutional capacity or lack thereof; many developing contexts, for example, have very sparse historical data on either disaster intensities or microeconomic indicators, and may be severely limited even today. Regardless, absence of observations clearly obviates the ability to generate internally valid results describing outcomes of interest, meaning that vulnerability functions for such events need to be extrapolated, generating potential issues of external validity. The second issue is that even if data exist, disasters are often sufficiently infrequent that sample sizes are small, making estimates of damage functions fairly imprecise and reducing econometricians' ability to control for other factors. This becomes particularly concerning with estimating equilibrium social responses to disaster incidence (e.g., indirect responses) or similar complex processes which generally require more controls and/or greater variation in potentially confounding correlates than simple reduced form estimates. A tertiary concern with disasters' low interoccurrence times arises only for disasters with slow onset profiles or noisy ("fuzzy") spatial footprint. If disasters are relatively sharply delineated in space and time, assigning disaster incidence treatment to households is easy; if household location and duration of residence in that location is known, one may simply check to see for temporal and spatial overlap with a disaster measure and assignment is unambiguous. For slow onset or fuzzy-border disasters, however, assignment is not so easy, and often relies either on functional form assumptions regarding the impact of damages, which may be difficult to defend, or increasing the sample size to empirically determine treatment behaviors, which may not be possible due to low sample size.

The secondary major threat to the internal validity of estimates of disaster damages arises from disasters' disruptive impact on society. Most saliently, disasters, or even the perceived risk of disaster in and of itself, may cause individuals to sort, e.g., by migrating away from disaster-incident areas. Since sorting may be the result of a variety of difficult-to-understand social processes, estimates which do not account for sorting risk generating biased estimates of a disaster's impact. To illustrate, if one of the major effects of flooding is to drive relatively more economically prosperous households to relocate to higher ground, a comparison of average pre- and post-flood household characteristics will show an upwards-biased (by which we mean larger) estimate of the poverty impacts of the flood on households, as richer households are exiting the sample during the post period. Failing to account for sorting or similar attrition of observations, or at least bounding its influence on estimates, may result in seriously biased estimates of disasters' impacts. This concern is particularly heightened for disasters with low local occurrence times (i.e., where disaster incidence is expected by the populace), and this may encourage ex-ante sorting based on tastes for or ability to pay for risk exposure.

The last major threat to internal validity arises in estimating disasters' indirect damages to microeconomic indicators, e.g., by changing adaptive behaviors. Indirect damages are generally diffuse and slower to manifest; the capital loss of a building is easy to document once a hazard has destroyed it, but the indirect damages resulting from disruption in economic activity, changes in local investment behavior, and similar, occur in less immediately observable fashion, and over time, making estimation difficult. Estimates of indirect (and possibly lagged / dynamic) damages thus generally require richer datasets with larger sample sizes, typically panel in nature, in order to yield tractable estimates, which again pushes up against

concerns about low sample size in the empirical hazard literature. Moreover, estimates of equilibrium social response are often conditional on institutional and economic conditions at the time of disaster incidence, which may be difficult to capture in data, further threatening the validity of results. Policy makers, for example, may elect to respond to certain types of events differently at different times, with potentially drastically different end effects on social welfare (e.g., after switching from an ineffective to effective response regime).

External Validity: While threats to internal validity risk misestimating the vulnerability function describing microeconomic damages due to hazard exposure, threats to external validity instead risk misestimating reasonable vulnerability functions for domains in which internally valid estimates are unavailable. In the parlance of equation (1), above, external validity risks miss-specifying the damage function $D(\cdot)$ for values of s which were unavailable for analysis. Threats to external validity can be thought of as general threats to extrapolating results from one context to another.

The most serious threat to external validity again arises from data paucity; if we wish to estimate, for example, the impact a major earthquake in central Mexico would likely have on microeconomic indicators of interest, one would have to infer it by referring to estimates of less serious earthquake damages in Mexico as well as estimates of the damages caused by major earthquakes in other contexts, simply because central Mexico has not recently experienced a major earthquake. This concern becomes particularly salient when comparing contexts across strongly varying institutional, political, or development contexts; economists and other social science researchers are still only beginning to uncover mechanisms driving microeconomic losses, especially for indirect damages, and since these mechanisms can often plausibly be linked to social context, changes in that context will result in differing vulnerability functions. While some initial strides are being made in the literature to understand, or at least approximate, conditional damage functions, such work is still in its early stages. Moreover, it is worth noting that disasters are often sufficiently notable events for a country or region that they can directly result in changes in government policy, often as a result of difficult-to-model political negotiation. Endogenizing this sort of response and including it in hazard models is a Herculean task beyond the scope of any presently feasible quantitative models, necessitating increased reliance on the subjective judgment of analysts and reference to economic or political theory.

There are several additional concerns which arise when considering generating externally valid estimates of disaster damage functions. Recurrence times between events may influence event outcomes (e.g., two bad storms in a row may be worse than two bad storms with a ten year lag between them), and while data may allow reasonably valid estimates of damage functions as a function of hazard intensity for single events, accounting for possible compound effects adds an additional degree of freedom that makes demands on sample size even more extreme. Events may vary in fundamental ways according to geography or other state variables; volcanic eruptions, for example, behave in generally idiosyncratic ways even for the same volcano, much less across volcanoes, making estimates of "average" volcano damages very noisy. Lastly, it is important to note that regardless of all other external validity concerns, internal validity of retrospectively estimated results does not guarantee similar behavior in the future. Contexts change over time, and prospective estimates of damages from events will always be subject to the caveat that prior results do not guarantee future performance.

Section III: Reconciling the Current State of the Microeconomic Literature on Disasters with the Needs of Cat Modelers

Issues of internal and external validity established, we can now identify the disparity between what cat models would ideally like from the currently existing economics literature and what the literature can currently provide, as well as an outline potential steps for remediating that gap.

In an ideal world, the economics and broader social science literature would provide a sufficiently large number of estimates of microeconomic damages due to hazard exposure such that, for a given set of microeconomic indicators in a given country context, cat risk modelers would be able to aggregate damage estimates into vulnerability functions that map a given hazard realization (e.g., 1 in 20 year storm, or an earthquake with peak ground acceleration of 0.3 m/s) onto microeconomic indicator outcomes (e.g., a 10% increased risk of stunting among affected families). This damage function would capture both direct and indirect losses (e.g., both damages to physical capital as well as lost income due to disruption), dynamic impacts of disasters (i.e., how losses change or accumulate with time after event onset), and response-contingent results (e.g., policy response of the government). While this ideal is far from universally attainable, recent progress has been made in a variety of domains which has made it more feasible⁴. We outline currently confounding issues below.

Data availability: Several of the most notable concerns with generating ideal cat risk vulnerability functions have to do not with informing the function through reference to microeconomic results, but rather with simply ensuring the vulnerability function is being correctly used in the context of the cat model. Most notably, any vulnerability function must necessarily take as input some measure of hazard intensity (h in equation 1) which can then be mapped to losses. Ensuring that relevant hazards have been modeled correctly, specifically ensuring that hazard models reflect the actual temporal and spatial intensity of representative hazards, is thus paramount. Similarly, cat risk models must be certain to include detailed and up-to-date data on the spatial distribution of households, as well as their demographic characteristics and any relevant contextual data (s in equation 1), in order to ensure that exposure has been properly modeled. Where existing data are sparse, cat risk modelers may take a cue from recent trends in the development economics literature which have leveraged nontraditional data sources; see, for example, the recent use of cellular phone to infer both migration response and informal insurance network response following natural disasters in Rwanda and Haiti (Blumenstock, Eagle, and Fafchamps, 2012, and Bengtson et al., 2011).

Relative youth of the disaster impacts field: In addition to data availability, a major reason for the lack empirical estimates which can be used to construct vulnerability functions is the relative youth of the disaster impacts field. The computational and data storage needs of many micro level disaster impact projects are fairly large, meaning that it has only recently become feasible to easily run models with the necessary number of degrees of freedom and sample size capacity to estimate disaster incidence at the micro level in large data sets. At the same time, the econometric framework remains fairly recently developed and is continually being refined, e.g., through establishing the validity of novel sources of identification of exogenous variation. As a result, even where data on both historical hazard exposure and microeconomic outcomes exist, established research will often document relationships only for a subset of indicators, or only in specific countries, regions, or contexts.

⁴ See recent advances in the climate and conflict literature (Hsiang, Burke, and Miguel, 2013) and climate and health literature (Graff Zivin and Neidell, 2013) for examples of attempts to infer generalized vulnerability function analogues for climate variability.

Lack of standardization of results: Youth of field, differences in needs of academic economists versus policy makers, and multiplicity of disciplines involved in field combine to produce inconsistent estimate reporting. For example, papers in economics alone may report disaster damages due to tropical cyclones in terms of elasticities with respect to wind speed, categorical treatment (e.g., the Saffir-Simpson scale), power dissipation, etc. This heterogeneity is exacerbated by the fact that the econometric advantages and disadvantages of different treatments types still being debated and actively uncovered by practitioners in current research. While in many cases direct translations of treatment and outcome variables may be trivial (e.g., categorical Saffir-Simpson scale cyclone intensity maps fairly directly onto maximum wind speed measures), others may require reference to physical models, historical data, or discussion with original researchers in order to map estimates from the literature onto empirically usable vulnerability functions.

Heterogeneity of impact estimability by disaster type and context: Variation in hazard intensity is a necessary precondition for econometric estimation of disasters' impacts, and as such the fundamental estimability of vulnerability functions mapping hazard incidence onto microeconomic indicators can be expected to vary heavily with disaster type and context. Estimates of disaster impacts on indicators can be expected to be more reliable for a disasters which occur more frequently, vary highly intensity, are better observed, and are more sharply defined in time. Much ground has been covered in recent years for disasters which satisfy those conditions, e.g., in the literature on the long-run human capital impacts of drought, and in such areas of research the gap between cat modelers' needs and extent literature can be expected to close fairly rapidly as research progresses. Disasters which fail to satisfy those conditions will present additional challenges, which may vary greatly depending on the specific outcomes being examined, directness vs. indirectness of mechanisms, and the like. In some cases internal validity can be enhanced by exploiting plausibly exogenous spatial variation in the impact of relatively rare events⁵; in others it may be appropriate to compare disaster impacts that are physically similar across disaster types, e.g., by using estimates of storm surge damages for coastal areas hit by cyclones (which are reasonably frequent and thus more estimable) to infer probable damages from tsunamis (which are very rare).

Direct vs. indirect damages: As outlined in Section II above, indirect damages which result from human response to disasters, such as income loss due to suppressed market activity or disinvestment in children's education due to disaster-induced income shocks, can generally be expected to vary more strongly with local social and economic conditions than direct losses such as physical depreciation of assets. Indirect damages thus generally ask more of available data, and will be harder to estimate than direct damages, though they can be plausibly estimated subject to sufficient data availability and heterogeneity in disaster incidence. Estimating indirect damages, while difficult, is likely important, especially given that recent research (e.g., Anttila-Hughes and Hsiang, 2013) suggest that indirect damages may dwarf direct ones in the long run.

Given the role that mediating social and economic factors (s in Equation 1) play in affecting disaster impacts, we may generally put more stock in estimates of indirect damages that result from equilibrium adaptive behavior that are not changed dramatically by the disaster in question. Typhoons are sufficiently common in the Philippines, for example, that their arrival, while devastating, is fairly expected, resulting in relatively little large scale social upheaval that would contaminate estimates of average losses due to

⁵ See, for example, Hallstrom and Smith, 2005, or for a manmade disaster, Almond, Edlund, and Palme, 2009.

storm incidence. Rare but devastating events, such as volcanic eruptions or tsunamis, can be expected to cause larger shifts in equilibrium social organization, e.g., by fomenting conflict, changing political behavior, and similar, making estimates of indirect damages more difficult to infer. Where social contexts are uncertain, understanding mechanisms of indirect loss and establishing how they may be expected to vary with covariates across contexts is paramount. Burgess et al., 2013, for example, note that credit market access may partially be driving mortality effects of heat waves among farmers in India; agriculturally damaging events in similar contexts may thus be expected to have worse indirect effects where credit markets are more weakly developed.

Lastly, it should be noted that new general equilibrium models that could estimate equilibrium market response to disaster damages would greatly improve cat modelers' ability to infer indirect damages, though a large amount of work is still needed in this area. While general equilibrium modeling has been used to varying degrees of success in order to predict impacts of certain policies, e.g., in public finance, their ability to model economies' response to physical damages remains poor⁶.

Out of sample extrapolation, interpolation, and subjective judgment: Regardless of whether due to data availability, youth of the field, or sundry other issues, empirically well-identified vulnerability functions linking hazard exposure to microeconomic outcomes are sparse and will likely remain so for the near term future. Cat modelers seeking to link microeconomic indicators to existing hazard models must thus expect to use a good deal of subjective judgment in arriving at best estimates of damage functions based on comparable estimates while the empirical literature closes the gap.

In general, extrapolating and interpolating estimates for vulnerability functions reliably is a function of identifying which social mechanisms or pathways are causing observed variation in the indicators. If the income loss effect of droughts in agricultural regions is empirically sensitive to local market access, for example, assigning a vulnerability function derived from one low-market access context to another may be the best possible practice until a local estimate is determined. It is difficult to arrive at a general set of heuristic techniques that may be followed at this point, and in general the cat modelers' familiarity with context will be the best guide to whether an out-of-sample extrapolation or interpolation is apt. There is a particularly helpful role of less-well identified field research, case studies, "quasi case studies" (i.e., techniques such as propensity score matching, synthetic controls⁷, or similar which attempt to introduce plausible internal validity in otherwise not well identified data), historical comparables⁸, and similar at this stage.

It is important to note that in many cases context-specific estimates which are beyond the scope of data density limits may be resolved by pooling estimates. For particularly rare types / intensities of disasters, it may be best to first estimate a "global average vulnerability function" across contexts to determine general form relationships; from this platform local similarities to global outcomes can be inferred or tested.

⁶ See Pindyck, 2013, for a related discussion of the difficulty of incorporating general equilibrium frameworks into integrated assessment models for modeling the expected social costs of climate change.

⁷ See Cavallo et al. 2013

⁸ E.g., Hornbeck and Naidu, 2012, or Hornbeck, 2012.

In certain cases reference to the physical or natural sciences community may ease comparison of disaster impacts across contexts and event types. Recent advances in the agronomic literature, for example, have demonstrated nonlinearity of impacts of heat waves on crop yields in the United States, an area of fairly rich data availability (Schlenker and Roberts 2009); these estimates can then be used to infer the likely impact of climate change on African agriculture, where data are sparser (Schlenker and Lobell, 2010).

Section IV: Recommendations for Next Steps

In light of the relatively large gap between the immediate needs of the cat risk modeling community, especially the specific needs of the World Bank DRFI, we propose the following recommendations for immediate action steps, ranked in order of difficulty / attainability from easy to hard:

1. Synthesize existing areas of research where a large number of impact estimates exist, e.g., the literature on drought impacts on human capital, and standardize hazard exposure measures to get to generalizable vulnerability functions.
2. Identify areas with both rich survey data and rich variation in disasters to establish basic outcomes where they do not yet exist, i.e., make up for lack in coverage that is simply due to the field being young.
3. Use professional judgment and growing understandings of mechanisms to establish external validity of results to similar contexts, preferably generating secondary testable hypotheses to infer generalizability / appropriateness.
4. Refer to case studies, quasi-case studies, and similar where well identified estimates do not exist, and rely on subjective judgment as needed.
5. In the intermediate term, identify or commission new data sources to better identify effects, and encourage development of theory models that may provide testable inferences where well-identified estimates are implausible.

References

- Almond, Douglas, Lena Edlund, and Mårten Palme. "Chernobyl's subclinical legacy: prenatal exposure to radioactive fallout and school outcomes in Sweden." *The Quarterly Journal of Economics* 124.4 (2009): 1729-1772.
- Angrist, Joshua D., and Jörn-Steffen Pischke. *Mostly harmless econometrics: An empiricist's companion*. Princeton university press, 2008.
- Anttila-Hughes, Jesse K., and Solomon M. Hsiang. "Destruction, disinvestment, and death: Economic and human losses following environmental disaster." (2013) Available at SSRN: <http://ssrn.com/abstract=2220501>
- Bengtsson, Linus, et al. "Improved response to disasters and outbreaks by tracking population movements with mobile phone network data: a post-earthquake geospatial study in Haiti." *PLoS medicine* 8.8 (2011): e1001083.
- Blumenstock, Joshua, Nathan Eagle, and Marcel Fafchamps. "Risk and reciprocity over the mobile phone network: evidence from Rwanda." (2012).
- Burgess, Robin, et al. "The Unequal Effects of Weather and Climate Change: Evidence from Mortality in India." (2013) <http://www.iae.csic.es/investigatorsMaterial/a1330908215765036.pdf>
- Cavallo, Eduardo, et al. "Catastrophic natural disasters and economic growth." *Review of Economics and Statistics* 95.5 (2013): 1549-1561.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. "What Do We Learn from the Weather? The New Climate-Economy Literature." No. w19578. National Bureau of Economic Research, 2013.
- Grossi, Patricia and Kunreuther, Howard. "Catastrophe Modelling: A New Approach to Managing Risk", Huebner International Series on Risk, Insurance and Economic Security, 2005.
- Graff Zivin, Joshua, and Mathew Neidell. "Environment, health, and human capital." *Journal of Economic Literature* 51 (2013): 689-730.
- Hallstrom, Daniel G., and V. Kerry Smith. "Market responses to hurricanes." *Journal of Environmental Economics and Management* 50.3 (2005): 541-561.
- Holland, Paul W. "Statistics and causal inference." *Journal of the American statistical Association* 81.396 (1986): 945-960.
- Hornbeck, Richard. "The Enduring Impact of the American Dust Bowl: Short-and Long-Run Adjustments to Environmental Catastrophe." *The American Economic Review* 102.4 (2012): 1477-1507.
- Hornbeck, Richard, and Suresh Naidu. "When the levee breaks: Labor mobility and economic development in the American south." NBER Working Paper 18296 (2012).
- Hsiang, Solomon M., Marshall Burke, and Edward Miguel. "Quantifying the influence of climate on human conflict." *Science* 341.6151 (2013): 1235367.

Pindyck, Robert S. Climate Change Policy: What do the models tell us?. No. w19244. National Bureau of Economic Research, 2013.

Schlenker, Wolfram, and David B. Lobell. "Robust negative impacts of climate change on African agriculture." *Environmental Research Letters* 5.1 (2010): 014010.

Schlenker, Wolfram, and Michael J. Roberts. "Nonlinear temperature effects indicate severe damages to US crop yields under climate change." *Proceedings of the National Academy of Sciences* 106.37 (2009): 15594-15598.