Essays in Environmental Economics

by

Tatyana Deryugina

B.S. Environmental Economics and Policy
B.A. Applied Mathematics
UC Berkeley, 2006

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2012

Author ....
Department of Economics
February 15, 2012

Certified by ....
Amy Finkelstein
Professor of Economics
Thesis Supervisor

Certified by ....
Michael Greenstone
3M Professor of Environmental Economics
Thesis Supervisor

Accepted by ....
Michael Greenstone
3M Professor of Environmental Economics
Chairman, Departmental Committee on Graduate Studies

Essays in Environmental Economics

by

Tatyana Deryugina

Submitted to the Department of Economics
on February 15, 2012, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis examines various aspects of environmental economics. The first chapter estimates how individuals' beliefs about climate change are affected by local weather fluctuations. Climate change is a one-time uncertain event with no opportunities for learning; the belief updating process may not be fully Bayesian. Using unique survey data on beliefs about the occurrence of the effects of global warming, I estimate how individuals use local temperature fluctuations in forming these beliefs. I test for the presence of several well-known psychological heuristics and find strong evidence for representativeness, some evidence for availability and no evidence for associativeness. I find that very short-run temperature fluctuations (1 day – 2 weeks) have no effect on beliefs about the occurrence of global warming, but that longer-run fluctuations (1 month – 1 year) are significant predictors of beliefs. Only respondents with a conservative political ideology are affected by temperature abnormalities.

In the second chapter, I examine the economic impacts of natural disasters by estimating the effect of hurricanes on US counties' economies 0-10 years after landfall. Overall, I find no substantial changes in a county's population, earnings, or the employment rate. The largest empirical effect of a hurricane is observed in large increases in government transfer payments to individuals, such as unemployment insurance. The estimated magnitude of the extra transfer payments is large. While per capita disaster aid averages $356 per hurricane in current dollars, I estimate that in the eleven years following a hurricane an affected county receives additional non-disaster government transfers of $67 per capita per year. Private insurance-related transfers over the same time period average only $2.4 per capita per year. The fiscal costs of natural disasters are thus much larger than the cost of disaster aid alone. Because of the deadweight loss of taxation and moral hazard concerns, the benefits of policies that reduce disaster vulnerability, such as climate change mitigation and removal of insurance subsidies, are larger than previously thought. Finally, the substantial increase in non-disaster transfers suggests that the lack of changes in other economic indicators may be in part due to various social safety nets.

In the third chapter, I estimate the extent of adverse selection in area yield insurance. Despite a long-run decrease in developed countries' vulnerability to weather shocks, agriculture worldwide remains susceptible to weather fluctuations. If climate change increases the frequency and intensity of extreme weather events, as it is predicted to do, food prices will likely become more volatile. A well-functioning insurance market is key to keeping the agricultural sector stable. I discuss the institutional and empirical features of the US crop insurance market. I outline the ways in which market designers have attempted to minimize adverse selection and moral hazard, as well as the remaining ways in which the market remains vulnerable to these. I then test for a particular form of adverse selection: whether public information (last year's average yield in the county) that is not explicitly priced by crop insurance companies predicts takeup of area yield insurance plans. I find no evidence that the recent yield influences takeup. I then perform another reduced-form test, using end-of-growing season yields as predictors of insurance takeup at the beginning of the growing season, and find that area yield insurance takeup is higher when average yields are higher. This suggests that the net selection into area yield plans favors providers, not buyers of insurance. In some specifications, the total demand for crop insurance is affected by current and past yields as well, potentially due to changes in the desirability of other plans.

Thesis Supervisor: Amy Finkelstein
Title: Professor of Economics

Thesis Supervisor: Michael Greenstone
Title: 3M Professor of Environmental Economics

Acknowledgments

This work could not have been completed without the help and guidance of many many people. First and foremost, I would like to thank my advisers, Amy Finkelstein and Michael Greenstone, who taught me much of what I know about the research process (both by example and direct instruction) and who spent many hours patiently giving me feedback. I am grateful for their guidance through these years.

When I entered graduate school, I was told that the most important presence in my life would be my classmates. As I leave, I can fully attest to that. They helped me navigate the first two years of courses, inspired and challenged me with their feedback and comments on everything from nascent ideas to drafts, and provided much-needed breaks from work. I would especially like to thank Jason Abaluck, whose frank and poignant criticism has made me a better researcher, and Mar Reguant-Rido for being my "sister-in-arms" and an amazing officemate.

I gratefully acknowledge the financial support of the National Science Foundation's Graduate Research Fellowship and the MIT Energy Initiative. Additional support for Chapter 1 was provided by the MIT Shultz Fund.

I would like to thank Stefano DellaVigna, who introduced me to research in economics when I was an undergraduate. I can confidently say that I would not have gone to graduate school had it not been for him.

My sister Mariya, who was also my roommate for three years of graduate school, has always been a great source of support. She listened patiently to hours of economics talk and didn't for a moment seem disinterested or annoyed. I am lucky to have her as a sister and a friend.

I am extremely grateful for having Julian Reif in my life. He pushed me when I needed to be pushed and comforted me when I needed comforting. I thank him for his wisdom, kindness, loving encouragement and unceasing support.

Finally, I would like to thank my parents, Vladimir and Irina, who worked incredibly hard to give me every opportunity in the world. They taught me to value education and to persevere no matter what. I dedicate this thesis to them.

Contents

Chapter 1

How do people update? The effects of local weather fluctuations on beliefs about global warming

1.1 Introduction

The hypothesis that increased greenhouse gas concentrations may lead to a rise in global temperatures first emerged in the 1960's (Peterson et al. 2008). The overwhelming majority of climate scientists now agree that the evidence for anthropogenic global warming is strong (Rosenberg et al. 2010).

Climate change may be one of the most disruptive events of the 21st century. Predictions about average temperature changes and economic costs of climate change are uncertain, but generally bleak: for increases of 5-6 °C, which is a "Business as Usual" scenario, the predicted economic loss is 5-10% of global GDP (Stern 2007).

Despite efforts of scientists, the general public first became concerned about global warming only in 1988, after the US experienced what back then was the hottest year on record (1987) and an extreme drought. However, the public's attention soon waned (Ungar 1992). To date, the international collaboration necessary to reduce warming has not been achieved. The Kyoto protocol, an international agreement established in 1997 to curb greenhouse emissions, did not affect several of the largest emitters today (such as the US, which has not ratified it, and China, which is exempt from compliance). The US has recently tried and failed to pass legislation that would have established a CO2 emissions trading scheme. Although there is consensus that large cuts in global emissions are necessary to avoid substantial harm (United Nations Framework Convention on Climate Change 2010, 2011), there is currently no international agreement that is expected to result in such cuts.

The implementation of effective public policy depends not only on climate science but on public perception of the occurrence and seriousness of climate change. Although the potential reasons for the lack of a strong international treaty are abundant, the lack of overwhelming public support may be a particularly important one. In a 2010 Gallup Environmental Poll, only 50% of respondents thought that the effects of global warming have already begun to happen, a further 20% thought they would never happen, and only 29% thought that global warming would be a significant threat to them or their lifestyle in their lifetime.

Global warming is a very visible issue, and thus some public consensus is necessary to implement policies that address it. For this reason, it is important to consider how beliefs about climate change are formed and updated. Many models with uncertainty often assume that agents update their beliefs using Bayes' rule. In some settings, this type of rationality is difficult to justify without empirical evidence. Climate change is a highly complex one-time event without opportunities for learning. How individuals use information in this or similar contexts is an empirical question.[1]

I use a large representative sample of US adults who were surveyed about the occurrence of global warming (the underlying cause of climate change) to test how local temperature fluctuations affect beliefs. The dataset is rich and spans multiple years, allowing me to include numerous controls. The question about the occurrence of global warming is straightforward and has categorical answers that fit easily into a regression framework. I consider the effects of both short (1 day – 2 weeks) and prolonged (1 month – 12 months) periods of abnormal temperatures.

Overall, beliefs are not affected by very short-term (on the order of days) temperature fluctuations, although longer periods of extremes do have an effect. Some features of updating are consistent with Bayesian belief formation. However, when I test for the presence of specific biases, I find evidence that representativeness, a well-known psychological heuristic, plays an important role in the observed updating patterns. Availability is present to the extent that individuals give significantly more weight to local temperatures than to national or global temperatures. However, there is no evidence for another type of availability, where individuals give more weight to recent temperature fluctuations than to less recent ones. I find no evidence for associativeness, where recent temperature fluctuations cause individuals to recall similar weather instances from the past and update based on the recalled rather than the true weather history. Finally, when I estimate the effects of weather by political ideology, I find that only those who are conservative or very conservative are affected by temperature fluctuations.

I contribute to the existing literature on the effects of environmental cues on beliefs about global warming by focusing on the effects of longer-run temperature abnormalities; previous studies have only looked at temperature fluctuations over one day to one week. In addition, I test for the presence of specific biases, based on previous psychology studies unrelated to global warming. Finally, due to the large sample size, I am able to test for differences in updating between conservatives, liberals,

[1] The empirical evidence on updating is mixed. Evidence for various forms of irrational updating includes DeBondt and Thaler (1984) in finance, Terrell (1994) and Clotfelter and Cook (1993) in lottery play, and Egan and Mullin (2009), Risen and Critcher (2009), and Cameron (2005) in climate change beliefs.

and moderates.

Section 2 describes the conceptual framework and links different biases to expected relationships between temperature and beliefs. Section 3 describes the data and the construction of the regression variables. The empirical framework and results are presented and discussed in Section 4. Section 5 concludes.

1.2 Conceptual Framework

Let G and $\neg G$ represent the states of the world with and without global warming, respectively. Let E be some evidence observed by someone who is a Bayesian updater. Beliefs about whether global warming is happening or not will then be determined by Bayes' formula:

$$Pr\left(G|E\right) = \frac{Pr\left(E|G\right)Pr\left(G\right)}{Pr\left(E|G\right)Pr\left(G\right) + Pr\left(E|\neg G\right)\left(1 - Pr\left(G\right)\right)}$$

In general, "evidence" can include global or local weather, a news story on melting glaciers, an Intergovernmental Panel on Climate Change (IPCC) report, or long-run climate data. For the effect of local temperatures to be detectable empirically, people must observe them more precisely than temperature elsewhere (otherwise any change in beliefs will be absorbed by year fixed effects).

The change in beliefs following the updating of evidence from E to E' is $Pr(G|E') - Pr(G|E)$. The extent to which beliefs change depends on the prior, $Pr(G|E)$, and the relative likelihood of the observed evidence when there is global warming, $\frac{Pr(E'|G)}{Pr(E'|\neg G)}$. The further this ratio is from one, the larger the change in posteriors.

Following a period of abnormal local temperatures, a rational Bayesian updater may significantly change his beliefs about global warming if the ratio $\frac{Pr(t'|G)}{Pr(t'|\neg G)}$ (where t' represents newly observed temperatures) is large. For most local weather events, conditional on national weather, this ratio is likely to be close to one. Furthermore, this effect will be detectable empirically only if respondents observe local weather more precisely than national weather. If the updating process is largely Bayesian, the following patterns should be observed:

1. Longer periods of abnormal temperatures will have a larger effect than shorter periods.

2. More extreme temperatures will produce larger changes in beliefs.

3. Within a relatively short period of time, such as a year, whether extreme temperatures occurred more or less recently should not matter.

In addition to Bayesian updating, there are several heuristics that have been found to play a role in belief formation: associativeness, availability, and representativeness.

Under the availability heuristic, people use salient instances of an event to judge its likelihood. For example, someone who has witnessed a serious auto accident will judge the probability of such an accident to be higher that someone who has not seen one, even if both have identical statistical information (Kahneman and Tversky 1973; Kahneman et al. 1982). This bias predicts that people may be more likely to believe that global warming is occurring if they have experienced local fluctuations in temperatures, even if it is not rational to do so. The bias may be stronger if the temperature fluctuations are recent because recent events are presumably more salient.

Under the associativeness heuristic, current events cause past instances of similar events to be recalled (Mullainathan 2002). Abnormal temperatures could bias the recalled history toward similarly extreme events, leading the individual to conclude such events are more frequent than they are. In this case, E will be more likely to include events that are similar to recent events and less likely to include those that are not.

Associativeness and availability are similar, but can nevertheless be distinguished by statistical testing. If the availability heuristic is present, recent local temperature fluctuations will influence beliefs more than less recent fluctuations. If the associativeness heuristic is present, the interaction between recent and similar past temperature patterns will be a significant determinant of beliefs.

Representativeness is judging the probability of a sample by how much it resembles a salient feature of the population it came from (Kahneman et al. 1973; Grether 1980). For example, people judge the sequence HTTHTH to be more probable than the sequences HHHHTH and HHHTTT (Kahneman et al. 1972), although all three sequences are equally likely. Importantly, the representativeness of a sample is not affected by the sample size; therefore, neither are the subsequent probability estimates made by individuals. If the representativeness heuristic is involved in updating, then the length of the time over which temperatures are abnormal should not affect the magnitude of the effect.

1.3 Data

1.3.1 Gallup survey

For beliefs about global warming, I use Gallup's Environmental Poll for the years 2003-2010. Every March, about 1,000 US adults are surveyed within a 3-4 day window.[2]

The dependent variable in the subsequent regression analysis is the answer to the question of when the respondent believes the effects of global warming will start to happen. The exact wording is shown in Table 1, along with the breakdown of answers. The numerical value assigned to each answer for regression analysis is in parentheses following the answer.[3]

Overall, about 56.3% (out of 7,847) of respondents believe that the effects of global warming have already begun to happen, 12.9% think they will never happen, and the rest think they will

[2]The sample is representative of the US. Respondents are surveyed by phone. Global warming is not the sole focus of the survey: the topics include energy, the economy, US environmental policies, Arctic drilling, and environmental behaviors.

[3]"Refused" and "Don't know" are treated as missing in the regression analysis. These options never include more than 5% of the sample; for most of the questions, less than 3% of respondents chose these options.

happen sometime in the future. In the Electronic Supplemental Materials, I also show that beliefs about global warming vary significantly by characteristics such as gender, income, and education.

One possible objection to using the above question for assessing the effect of temperatures on beliefs is that the answers are categorical rather than expressed as probabilities. Although this does add some noise to the estimation, as long as there is some underlying continuous probability that the individual uses to answer the question, the effect of weather can still be seen using qualitative data.

To further check that any potential lack of significance is not due to noisy survey answers, I estimate how people update their beliefs about the country's economic conditions, shown in the Electronic Supplemental Materials. I find that respondents use local unemployment rates to make an inference about the economic conditions in the US, which supports the notion that the survey answers are not so noisy as to make statistical testing impossible.

Figure 1 shows the breakdown of beliefs by stated political ideology. The top graph shows the fraction of people of each political ideology that believes the effects of global warming have already begun to happen. The bottom shows the fraction that believes the effects of global warming will never happen.

Both graphs reveal considerable differences between people of different ideologies. Very conservative respondents are more likely to believe that the effects of global warming will never happen (nearly 38% of respondents reported this belief) than that they have already begun to happen (30% thought this). The probability that a respondent believes the effects of global warming have already begun to happen is highest for those who are liberal or very liberal (73% and 76%, respectively); these groups are also the least likely to report believing that the reports of global warming will never happen (2.5% and 3.5%, respectively). Political ideology is in theory multifaceted and stems from beliefs about many environmental and non-environmental issues. There is nevertheless a stark ideological divide on an issue that should be purely a scientific question.

1.3.2 Temperature fluctuations

This section outlines how abnormal temperature fluctuations are measured. For temperature data, I use National Climatic Data Center's daily weather station observations for maximum temperatures for 1949-2010, matched to counties.[4] These data were provided by Michael Greenstone and are used in Deschenes and Greenstone (2007a and 2007b).

The basic abnormality measure is the number of standard deviations from the long run average:

$$NumSD_{cd} = \frac{temp_{cd} - \overline{temp_{cd}}}{sd_{cm}} \tag{1.1}$$

d indexes the day of year; c indexes the county; m indexes the month. $temp_{cd}$ is the observed maximum temperature in county c on day d. $\overline{temp_{cd}}$ is the corresponding long-run average, con-

[4]If there are multiple weather stations in a county, I average their daily measurements.

structed by computing a seven-day running average across all years that precede the year of the survey. In other words, for respondents in county c taking the survey in year Y:

$$\overline{temp_{cd}} = \frac{1}{7\,(Y-1948)} \sum_{s=d-3}^{d+3} \sum_{y=1949}^{Y-1} temp_{cdy} \tag{1.2}$$

sd_{cm} is the standard deviation of maximum temperatures, constructed by computing the standard deviation of observed temperatures in that month and county between 1949 and 2000. I match each respondent's location and date of survey to the temperature data to determine the respondent's temperature deviations x days ago, where x ranges from 0 (day of the survey) to 364 (one year ago).[5]

To allow for a cumulative effect of longer stretches of abnormal temperatures, I construct variables that measure the fraction of days over a given time period on which the number of standard deviations was above a certain (high) quantile and the fraction of days on which it was below a low quantile. The formulas for these variables are:

$$FracAbove_{cnq} = \frac{1}{n} \sum_{t=0}^{n} 1\left\{numSD_{ct} \geq sd_q\right\} \tag{1.3}$$

$$FracBelow_{cnq} = \frac{1}{n} \sum_{t=0}^{n} 1\left\{numSD_{ct} \leq sd_{100-q}\right\} \tag{1.4}$$

where 1 is an indicator function. t is now relative to the day the respondent took the survey. n ranges from 7 to 360 days. q is a quantile of the number of standard deviations. I use q = 75, 90, and 95. Thus, the variables above measure the fraction of days on which temperature standard deviations were at or exceeded the 75th, 90th, and 95th quantiles and the fraction of days on which temperature standard deviations were at or below the 25th, 20th, and 5th quantiles. The resulting $FracBelow_{cnq}$ variables have means and standard deviations similar to their corresponding $FracAbove_{cnq}$ variables.

1.4 Effect of temperatures

1.4.1 Empirical framework

In this section, I outline the procedure for testing whether local temperature fluctuations significantly affect people's beliefs about global warming and, if so, whether this is due to psychological heuristics. Each regression specification is an ordered probit, which assumes that there is an underlying continuous outcome variable that is observed as categorical. In addition to computing the

[5]The respondents are called between 5pm and 9pm local time, making the inclusion of that day's temperatures reasonable.

average effect of an explanatory variable across all categories, it is also possible to estimate its effect for each category.[6]

I first regress beliefs on the number of standard deviations of maximum temperatures on the day of the survey:

$$Happening_{ict} = \beta numSD_{ct} + X_{ict}\gamma + \epsilon_{ict} \tag{1.5}$$

i indexes the individual; c indexes the county; and t indexes the survey date. $Happening_{ict}$ is the belief about the occurrence of the effects of global warming, equal to 5 if the respondent said they have already begun to happen, equal to 1 if the respondent said they will never happen, and taking on intermediate values for the other answer options.[7] X_{ict} is a set of flexible controls: sex, race, age, age squared, indicators for education level, income category, political ideology indicators, interactions of education and sex, interactions of political ideology and sex, and state and year fixed effects. $numSD_{ct}$ is the number of temperature standard deviations in the respondent's county on the day of the survey. I also estimate the effect of the number of standard deviations the day before the survey.

The second regression specification allows for the influence of longer periods of abnormal temperatures. This specification tests the predictions of the Bayesian updating model, as well as whether representativeness plays a role.

$$Happening_{ict} = \beta FracAbove_{cnq} + X_{ict}\gamma + \epsilon_{ict} \tag{1.6}$$

$FracAbove_{cnq}$ is the fraction of days over the n days before the survey on which the number of temperature standard deviations in county c exceeded the qth quantile, as described in Section 3.2. I also perform this test using the variable $FracBelow_{cnq}$ as the independent variable. This test also reveals whether extremely low temperatures the opposite effect of extremely high temperatures, in other words, whether the updating process is symmetric.

To test for associativeness, I interact the fraction of days over the past week on which temperatures exceeded a particular quantile with the fraction of days over the past n days on which temperatures exceeded the same quantile:

$$\begin{aligned} Happening_{ict} \;=\; & \beta FracAbove_{c7q} FracAbove_{cnq} \\ & + \theta FracAbove_{c7q} + \delta FracAbove_{cnq} \\ & + X_{ict}\gamma + \epsilon_{ict} \end{aligned} \tag{1.7}$$

The coefficient of interest is β. The idea is that even though people may not be affected by

[6]See Wooldridge (2002) for details about this estimation procedure.
[7]The full numerical coding of the questions is shown in Table 1.

last week's temperature abnormalities directly, a subset of them may be, as it causes them to recall similarly extreme weather over a longer period of time.[8] I also perform a version of this test using the variable $FracBelow_{c7q}$ and $FracBelow_{cnq}$ in the place of $FracAbove_{c7q}$ and $FracAbove_{cnq}$.

1.4.2 Results

Table 2 shows the marginal effects of a one standard deviation change in maximum temperatures on (a) the day of the survey and (b) the day before the survey. The effect of temperatures is computed for each of the five answer categories and can be interpreted as the additional probability that the respondent will choose a particular answer category for a one-unit change in the independent variable.

The estimated effects are small and insignificant. An effect of greater than one percentage point per one standard deviation of maximum temperatures can be ruled out for the answer category where the respondent believes the effects of global warming have already begun to happen.[9] The estimated effects on other answer categories are similarly small. Thus, very short-run fluctuations do not affect beliefs about global warming. This differs from some previous studies, which find that beliefs are affected by very short-run fluctuations (Egan and Mullin 2010; Joireman et al. 2010; Li et al. 2011; Schuldt and Schwarz 2009).[10] This divergence may be due to differences in the survey questions used to assess beliefs.

Figures 2-4 show the estimated effects of longer periods of abnormal weather on whether the respondent believes the effects of global warming have already begun to happen (answer value 5).[11] This is a natural category to focus on, as it should be most influenced by recent weather fluctuations. Moreover, the estimated effect of weather for this answer category generally has the opposite sign from the other four answer categories. Finally, the marginal effects on different answer categories are derived from a single estimated coefficient and a set of estimated thresholds and thus have nearly identical significance levels. The coefficients can be interpreted as the change in the probability that the respondent believes the effects of global warming have already begun to happen following a one-unit change in the fraction of abnormal days over the given time period.[12]

Figure 2 shows the estimated coefficients from regressions with the least extreme thresholds – standard deviations that are at the 25th percentile or lower and those at the 75th percentile or higher. The effect appears to be slightly asymmetric. Abnormalities over 7 and 14 days are insignificant. Persistently colder-than-normal weather over 30-360 days before the survey significantly decreases

[8] This is not the only possible formulation for associativeness. I test two other formulations, which produce similar results and are described in the Electronic Supplemental Materials.

[9] Looking at the effect of the average number of standard deviations over the week before the survey (shown in the Electronic Supplemental Materials) produces similar results.

[10] In a controlled experiment, Risen and Critcher (2009) find that indoor temperatures also affect beliefs.

[11] Point estimates can be found in the Electronic Supplemental Materials.

[12] Because the fraction of abnormal days theoretically varies from 0 to 1 (over longer periods, the fraction never reaches 1 in practice), this coefficient can also be interpreted as the effect of going from zero days having temperature deviations outside the defined thresholds to all days having temperature deviations outside the thresholds.

the probability that respondents will believe that the effects of global warming have already begun to happen, while persistently warmer-than-normal weather has the opposite (but insignificant) effect. A one-unit increase in the fraction of days of abnormally cold weather over 60 days (as defined in this specification) decreases the probability that the respondent believes that the effects of global warming have already begun to happen by 11.7 percentage points, while a one-unit increase over 180 days decreases it by 22.5 percentage points.

Figure 3 shows the results using slightly more extreme thresholds – standard deviations that are at the 10th percentile or lower and at the 90th percentile or higher. Abnormally warm weather now has a significant effect of 15.4 percentage points and 18.6 percentage points over 30 and 60 days, respectively. Abnormally cold weather has a significantly negative effect on beliefs over 60-180 days. Figures 2 and 3 suggest that the effects of abnormally warm and cold days are asymmetric: colder days decrease beliefs in global warming much more than warmer days increase them. Moreover, the magnitude of the coefficients does not increase for periods of more than 120 days. The last fact is indicative of representativeness.

Figure 4 shows the effects of deviations that lie outside the most extreme thresholds – 5th percentile or lower and 95th percentile or higher. Here, periods of abnormally warm weather significantly ($p<0.1$) increase the probability that the respondent says the effects of global warming have begun to happen for all period lengths. Extreme negative deviations, on the other hand, now have an insignificant effect on beliefs over the entire time period. The colder-than-normal estimates are now smaller (in absolute value) than their warmer-than-normal counterparts.

Next, I check for differences in updating by political ideology. Specifically, I separate the sample into conservatives, moderates, and liberals, as reported by the respondent. I estimate the effect on beliefs of the fraction of days on which the temperature standard deviations exceeded or were below a given threshold over periods of varying lengths for each of the groups.

Moderates and liberals are largely unaffected by this measure of abnormal weather (results are shown in Tables A4 and A5 in the Electronic Supplemental Materials). However, conservatives' beliefs are affected by weather fluctuations over various periods. These results are shown in Table 3. The response is asymmetric in that cooler-than-normal temperatures have no effect on beliefs, but warmer-than-normal temperatures do.

Why the response to abnormal weather is limited to conservatives is not immediately clear. Because conservatives are the least likely group to believe that the effects of global warming have already begun to happen, it's possible that they have a distribution of priors that is more likely to be affected by local weather abnormalities. In other words, weather fluctuations may make liberals' and moderates' beliefs that the effects of global warming are already occurring stronger, but they do not cause them to change their answer category. Alternatively, conservatives may have more distrust toward scientific reports and media and give that information less weight, relying relatively more on personal experience. The current data do not allow me to distinguish between these two hypotheses. However, this is an important area for future research.

Finally, Figure 5 shows the results of the associativeness test from equation 8. The coefficients are, for the most part, indistinguishable from zero. There is some indication that associativeness is present for temperature deviations at or below the 25th percentile, but given the lack of robustness with respect to other specifications, it seems more likely that the significance is spurious. This and two other specifications of associativeness (shown in the Electronic Supplemental Materials) suggest that this bias does not play a large role in updating in this context.

Overall, my findings resonate with Cameron (2005), who finds that updating exhibits both Bayesian and non-Bayesian attributes. The updating process exhibits some qualitative traits consistent with Bayesian updating: very short-run fluctuations for the most part do not affect beliefs and more extreme temperature abnormalities (for the medium-run measures) produce larger changes in beliefs. However, whether local temperature fluctuations should affect a Bayesian's belief about global warming is not straightforward. The Intergovernmental Panel on Climate Change notes that it is scientifically difficult to attribute a single event to the occurrence of global warming (IPCC 2007).

Moreover, if people observe weather everywhere in the US with nearly equal precision, local weather should be an insignificant predictor of beliefs, due to the inclusion of year fixed effects. The fact that local temperature plays any role in the updating process suggests that availability is also present, either through respondents observing local weather more precisely or giving it more weight in the updating process than weather elsewhere. Because coefficients do not consistently increase with the time period and because the effects of abnormally cool and abnormally hot temperatures are asymmetric, representativeness also seems to play an important role in the updating process.

There is little evidence for another kind of availability, which would predict that more recent temperature fluctuations have a larger effect than less recent ones. There is also little evidence of associativeness, which would predict that those who recently experienced abnormal temperatures and who had more abnormal weather in the less recent past would change their beliefs more than people who only experienced one of those two events.

1.5 Conclusion

Scientific estimates suggest that global warming may have catastrophic effects on the world's climate. The dire projections and overwhelming agreement in the scientific community that the time for mitigation is running out make immediate policy intervention increasingly necessary. However, international talks have to date failed to produce a comprehensive binding agreement to combat climate change. Although the potential reasons for this are abundant, the lack of public pressure may be an important contributing factor. It is thus essential to understand how individual beliefs about climate change are formed and what causes them to evolve.

Global warming is a highly uncertain event whose occurrence is very difficult to determine objectively, even for climate scientists. In addition, most people do not have all the information that a climate scientist does. Various violations of Bayesian updating have been found empirically. Biases such as representativeness, associativeness, and availability can cause individuals' updating processes to deviate from Bayesian models.

In this paper, I study the updating of beliefs about global warming. Using a multi-year survey, I test whether individuals use local temperature abnormalities to form inferences about global warming's occurrence. I find that very short-run fluctuations in temperatures over 1-2 days prior to the survey do not significantly affect beliefs. However, longer periods of abnormally warm or cold temperatures do change the probability that the respondent believes the effects of global warming have already begun to happen. This effect is limited to conservatives for reasons that are beyond the scope of this paper. Although some features of the updating process are Bayesian (more extreme temperature deviations produce larger changes in beliefs), the pattern of updating is also consistent with representativeness (beyond a 120-day period, longer periods of abnormal weather do not have a statistically larger effect).

The exact pathway through which these effects work is difficult to determine. Because I do not observe individuals' information set, I cannot rule out that individuals observe weather everywhere but irrationally give larger weight to local weather. It's also possible that the effects of temperatures are indirect. For example, more extreme temperatures could lead to more discussion of global warming in the local media and more exposure to other evidence about global warming, such as IPCC reports.[13]

Finally, the stark ideological divide in beliefs dwarfs any changes that plausible weather fluctuations can cause. Conservatives are much less likely than liberals to believe that the effects of global warming have already begun to happen and much more likely to believe that they will never happen. The exact reasons for this divide are beyond the scope of the paper, but should be an important subject of future research.

1.6 References

Belasen, A. and S. Polachek. 2008. "How Hurricanes Affect Employment and Wages in Local Labor Markets", IZA Discussions Paper No. 3407.

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

Clotfelter, C. and P. Cook. 1993. "The 'Gambler's Fallacy' in Lottery Play", *Management Science*, 39(12), 1521-1525.

DeBondt, W. and R. Thaler. 1984. "Does the Stock Market Overreact?", *The Journal of Finance*, 40(3), 793-805.

Deschenes, O. and M. Greenstone. 2007a. "The Economic Impacts of Climate Change: Evidence

[13]For example, Shanahan and Good (2000) find that climate issues were more likely to be covered in the New York Times during periods of unusually high temperatures.

from Agricultural Output and Random Fluctuations in Weather", *American Economic Review*, 97(1), 354-385.

Deschenes, O. and M. Greenstone. 2007b. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US", NBER Working Paper No. 13178.

Egan, P. and M. Mullin. 2010. "How Citizens Integrate Information without Ideological Cues: Local Weather and Americans' Beliefs about Global Warming", NYU working paper.

Grether, D. 1980. "Bayes Rule as a Descriptive Model: The Representativeness Heuristic", *Quarterly Journal of Economics*, 95(3), 537-557.

IPCC. 2007. "Climate Change 2007: The Physical Science Basis", Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change.

Solomon, S., Qin, D., Manning, M., Chen, Z., Marquis, M., Averyt, K.B., Tignor, M. and H.L. Miller (eds.) Cambridge University Press, Cambridge, UK and New York, NY, USA.

Joireman, J., Truelove, H. and B. Duell. 2010. "Effect of outdoor temperature, heat primes and anchoring on belief in global warming", *Journal of Environmental Psychology*, forthcoming.

Kahneman, D., Slovic P. and A. Tversky (eds.). 1982. "Judgment Under Uncertainty: Heuristics and Biases", New York: Cambridge University Press, reprinted in 2001.

Kahneman, D. and A. Tversky. 1972. "Subjective Probability: A Judgment of Representativeness", *Cognitive Psychology*, 4(3), 430-454.

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

Li, Y., Johnson, E., and L. Zaval. 2011. "Local Warming: Daily Temperature Change Influences Belief in Global Warming", *Psychological Science*, 22(4), 454-459.

Mullainathan, S. 2002. "A memory-based model of bounded rationality", *Quarterly Journal of Economics*, 117(3), 735-774.

Peterson, T., Connolley, W., and J. Fleck. 2008. "The Myth of the 1970s Global Cooling Scientific Consensus", *Bulletin of the American Meteorological Society*, 89, 1325-37.

Risen, J.L., and C. Critcher. 2009. "Visceral Fit: While in a Visceral State, Associated States of the World Seem More Likely", Manuscript in Preparation.

Rosenberg, S., Vedlitz, A., Cowman, D., and S. Zahran. 2010. "Climate change: a profile of U.S. climate scientists' perspectives", *Climatic Change*, 101 (3–4), 663-668.

Schuldt, J., Schwarz, N., and B. West. 2009. "Do people believe in global warming? Depends on what's the weather's like when you ask", Unpublished manuscript, University of Michigan.

Shanahan, J., and J. Good. 2000. "Heat and hot air: influence of local temperature on journalists' coverage of global warming", *Public Understanding of Science*, 9(3), 285-295.

Stern, N. 2007. "The economics of climate change: the Stern review", Cambridge University Press, Cambridge, UK.

Terrell, D. 1994. "A Test of the Gambler's Fallacy - Evidence from Pari-Mutuel Games", *Journal of Risk and Uncertainty*, 8(3), 309-317.

Ungar, S. 1992. "The Rise and (Relative) Decline of Global Warming as a Social Problem", *The Sociological Quarterly*, 33(4), 483-501.

United Nations Framework Convention on Climate Change. 2010. Report of the Conference of the Parties on its fifteenth session, held in Copenhagen from 7 to 19 December 2009. Addendum. Part Two: Action taken by the Conference of the Parties at its fifteenth session. United Nations Office at Geneva, Geneva (Switzerland).

United Nations Framework Convention on Climate Change. 2011. Report of the Conference of the Parties on its sixteenth session, held in Cancun from 29 November to 10 December 2010. Addendum. Part two: Action taken by the Conference of the Parties at its sixteenth session. United Nations Office at Geneva, Geneva (Switzerland).

Wooldridge, J. 2002. "Econometric Analysis of Cross Section and Panel Data", MIT Press, Cambridge, MA.

1.7 Electronic Supplemental Materials for: How do people update? The effects of local weather fluctuations on beliefs about global warming

1.7.1 Beliefs and Respondent Characteristics

In this section, I show the results of regressing stated beliefs about climate change on respondent demographics. The correlations between "Which of the following best reflects your view on when the effects of global warming will begin to happen?" and respondent characteristics reveal which groups are less likely to believe in the occurrence of climate change and support the notion that this belief is well-formed.

The included demographics are: sex, age, age squared, white indicator, log income, log income squared, education level indicators, political ideology indicators, and male-by-education-level and male-by-political-ideology indicators. Because the answers are categorical, the regression specification is an ordered probit.

The results are shown in Table 1. There is a quadratic relationship between age and beliefs about the timing of global warming, increasing in age and decreasing in age squared. The same pattern is true for income, implying that those with the highest and the lowest incomes are least likely to believe that the effects of global warming have already begun to happen.

More educated people are more likely to believe that the effects of global warming have already begun to happen, although males with college or graduate education are less likely to believe this than females with the same education level.

By far, the largest determinant of belief differences is political ideology: conservatives are much less likely to believe that the effects of global warming have begun to happen. Conservative males are even less likely to believe this than females, although the opposite is true for very liberal males.

1.7.2 Beliefs About Economic Conditions and Unemployment

Another way to address the concern that the survey answers may not reflect well-formed beliefs is to consider other beliefs that may be affected by measurable local information. Respondents to the Gallup Environmental Poll were also asked whether the state of the overall economy is "excellent", "good", "only fair" or "poor". Over the whole sample period, 31% of respondents said that economic conditions are good or excellent, 45.4% said they are fair, and 22.7% rated them as poor.

I use an ordered probit regression to examine the relationship between the respondent's assessment of economic conditions in the US and the average local (county) or state unemployment rate in the past month and in the 12 months before the survey. The results are shown in Table 2. An increase in the local or state unemployment rate over the past year has a large negative effect on the respondent's assessment of economic conditions. All the estimates are significant, and the inclusion of state or county fixed effects does not change the results qualitatively.

There is a smaller negative effect of last month's state and local unemployment rates on the probability that the respondent states that economic conditions are excellent or good and they are insignificant or marginally significant once county or state fixed effects are included. The larger significance of the annual unemployment rate is consistent with a Bayesian updater giving more weight to a larger number of observations. This provides evidence that people do sometimes use local information for updating beliefs and that the categorical survey answers are not so noisy as to make studying them empirically impossible.

1.7.3 Effect of Other Temperature Variables

In this section, I present supplementary tests of the effect of weather on beliefs about global warming. Specifically, I look at the effect of changes in temperatures and of the average number of standard deviations over the past week (rather than the number of standard deviations in the day of or before the survey). The results are shown in Table 3. The change in temperatures between the day of and the day before the survey is an insignificant predictor of beliefs. However, a one degree increase in temperatures between the day before the survey and two days before is estimated to decrease the probability that the respondent believes the effects of global warming have already begun to happen by 0.14 percentage points. This is counterintuitive, unless the temperature changes highlight to the respondent the variability of weather and cause him to discard other weather-based evidence that was increasing his belief in the occurrence of global warming. Alternatively, the correlation may be spurious. The average temperature standard deviation in the past week has no significant effect on beliefs. This also holds if deviations are measured in degrees. Thus, there is little evidence that people are systematically influenced by very recent fluctuations in weather.

1.7.4 Additional Analysis

Table 4 shows the point estimates corresponding to Figures 1, 2, and 3.

Table 5 shows the effects of longer-run abnormal temperatures on moderates' beliefs about whether the effects of global warming have already begun to happen.

Table 6 shows the effects of longer-run abnormal temperatures on liberals' beliefs about whether the effects of global warming have already begun to happen.

1.7.5 Other Associativeness Tests

In this section I present two alternative tests for associativeness. In the first alternative test, I interact the indicator for whether the number of temperature standard deviations exceeded a particular quantile on the day before the survey with the fraction of days over the past n days on which temperatures exceeded the same quantile:

$$\begin{aligned} Happening_{ict} &= \beta 1\left\{numSD_{c,t-1} \geq sd_q\right\} FracAbove_{cnq} + \theta FracAbove_{cnq} + \delta 1\left\{numSD_{c,t-1} \geq sd_q\right\} \\ &\quad + X_{ict}\gamma + \epsilon_{ict} \end{aligned} \tag{1.8}$$

The estimated coefficients β for the change in the probability that the respondent believes the effects of global warming have already begun to happen are shown in Figure A1. None of the estimates is statistically significant, supporting the notion that associativeness does not play a significant role in updating beliefs about global warming.

In the second alternative test, I interact the number of standard deviations on the day before the survey with the fraction of days over the past n days on which temperatures exceeded a particular quantile:

$$\begin{aligned} Happening_{ict} &= \beta numSD_{c,t-1} FracAbove_{cnq} + \theta FracAbove_{cnq} + \delta numSD_{c,t-1} \\ &\quad + X_{ict}\gamma + \epsilon_{ict} \end{aligned} \tag{1.9}$$

The results are shown in Figure A2. As in the previous test, there is no evidence for this particular form of associativeness playing a role in belief formation.

Tables and Figures

Table 1: Summary of responses, 2003-2010

Which of the following best reflects your view on when the effects of global warming will begin to happen?	
They have already begun to happen (5)	56.3%
They will start happening within a few years (4)	4.0%
They will start happening within your lifetime (3)	9.9%
They will not happen within your lifetime, but they will affect future generations (2)	16.8%
They will never happen (1)	12.9%
Observations	7,847

The coding of responses into numerical values to create dummy variables is in parentheses following the answer choice. Percentages may not add up to 100 due to rounding, people responding "I don't know" or refusing to answer the question.

Figure 1. Beliefs by political ideology

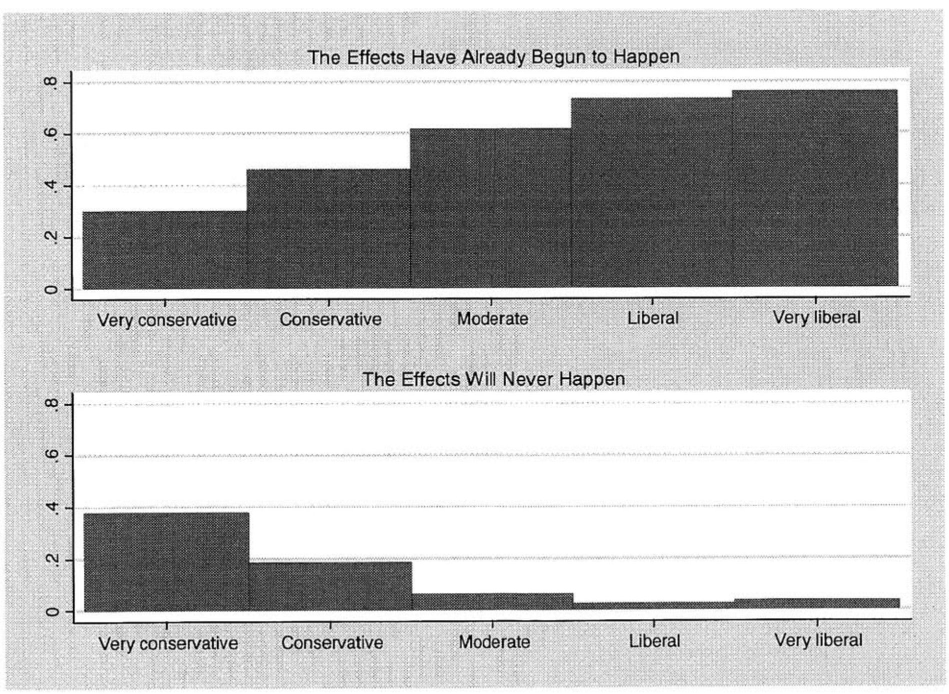

Table 4: Effect of recent temperature deviations on beliefs about global warming

	On day of survey	One day ago
Pr (Never Happen)	0.390	0.419
	(0.301)	(0.268)
Pr (Happen after lifetime)	0.359	0.382
	(0.284)	(0.247)
Pr (Happen within lifetime)	0.113	0.121
	(0.086)	(0.076)
Pr (Happen within few years)	0.026	0.027
	(0.019)	(0.017)
Pr (Already happening)	-0.887	-0.949
	(0.694)	(0.609)
Observations	5,448	5,443

The regression specification is an ordered probit. Marginal effects shown. Robust standard errors (clustered by state) in parentheses; * significant at 10%; ** significant at 5%; *** significant at 1%. Includes controls for respondent characteristics, state and year fixed effects. Probability is expressed in percentage points.

Figure 2. Effect of longer-run weather abnormalities 1. The solid line represents the point estimates. The dashed lines represent the 95 percent confidence interval

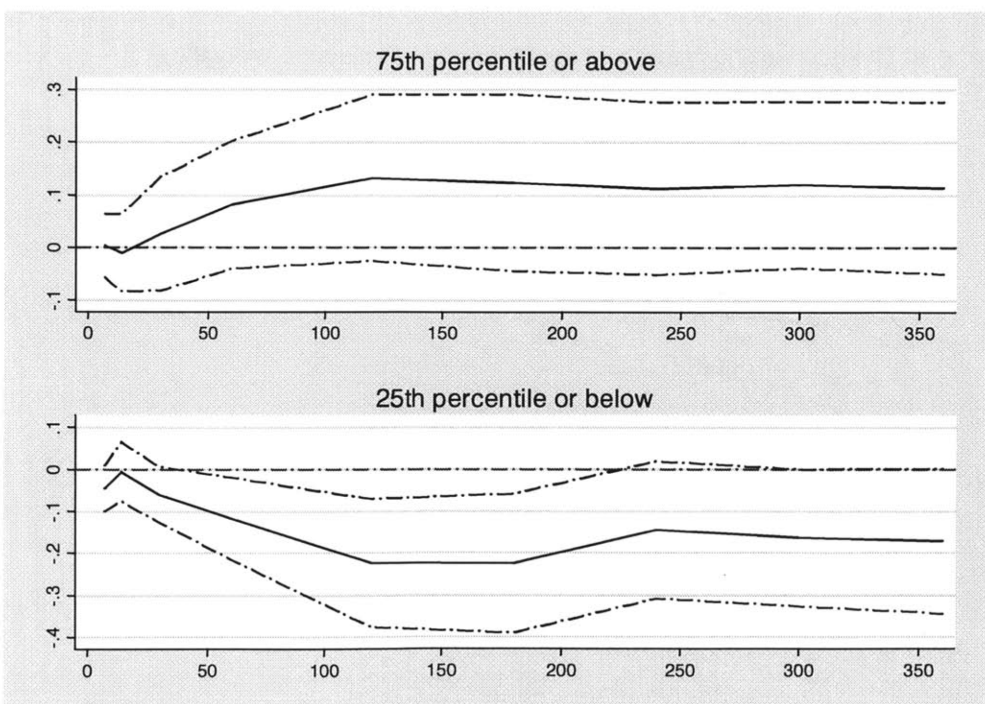

Figure 3. Effect of longer-run weather abnormalities 2. The solid line represents the point estimates. The dashed lines represent the 95 percent confidence interval

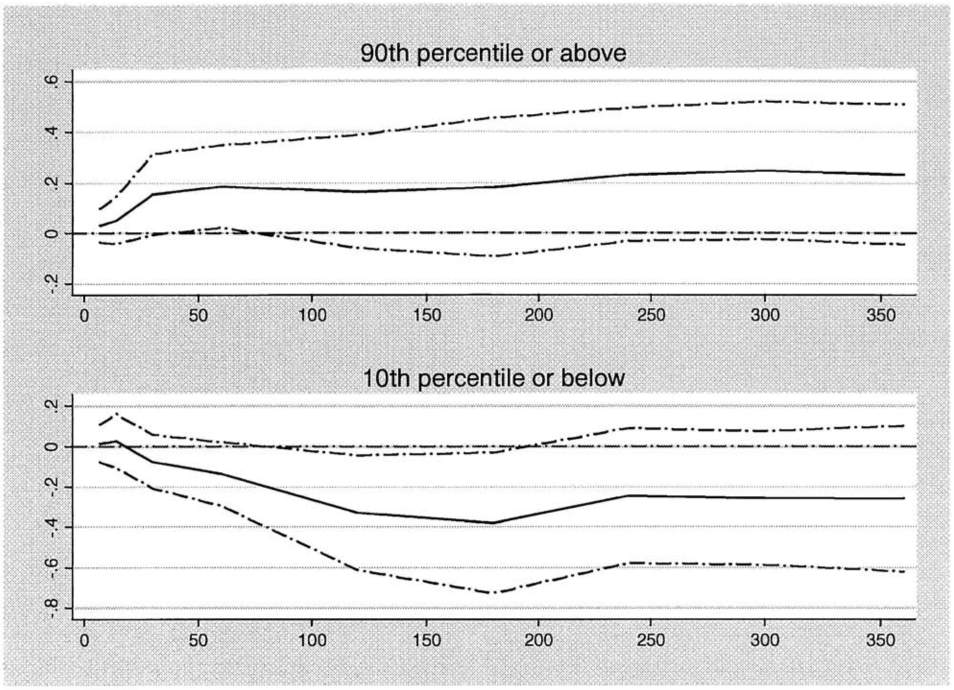

Figure 4. Effect of longer-run weather abnormalities 3. The solid line represents the point estimates. The dashed lines represent the 95 percent confidence interval

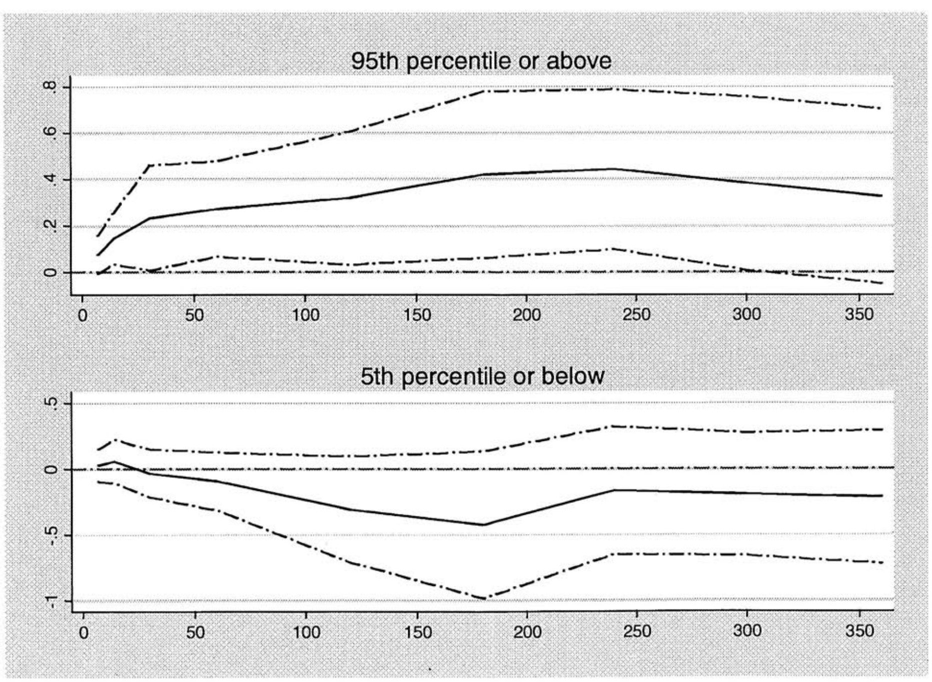

Table 3: Effect of longer-run abnormalities on conservative respondents

Quantile	25th or below	75th or above	10th or below	90th or above	5th or below	95th or above
0-7 days ago	-0.085	0.054	-0.029	0.092	0.013	0.182
	(0.032)**	(0.037)	(0.043)	(0.046)**	(0.052)	(0.052)***
0-14 days ago	-0.052	0.044	0.004	0.143	0.070	0.289
	(0.045)	(0.048)	(0.065)	(0.068)**	(0.079)	(0.078)***
0-30 days ago	-0.093	0.110	-0.077	0.297	0.023	0.481
	(0.048)*	(0.075)	(0.086)	(0.112)***	(0.119)	(0.149)***
0-60 days ago	-0.087	0.187	-0.082	0.350	0.047	0.450
	(0.059)	(0.089)**	(0.114)	(0.120)***	(0.131)	(0.130)***
0-120 days ago	-0.115	0.258	-0.182	0.425	-0.048	0.649
	(0.083)	(0.115)**	(0.173)	(0.180)**	(0.250)	(0.231)***
0-180 days ago	-0.136	0.261	-0.208	0.557	-0.146	0.873
	(0.098)	(0.119)**	(0.210)	(0.203)***	(0.349)	(0.282)***
0-240 days ago	-0.119	0.236	-0.145	0.530	-0.038	0.780
	(0.104)	(0.116)**	(0.212)	(0.206)***	(0.321)	(0.280)***
0-300 days ago	-0.112	0.222	-0.112	0.490	0.028	0.679
	(0.101)	(0.118)*	(0.209)	(0.193)**	(0.312)	(0.254)***
0-360 days ago	-0.115	0.217	-0.128	0.452	0.013	0.587
	(0.108)	(0.118)*	(0.216)	(0.194)**	(0.314)	(0.261)**
Observations	2,797	2,797	2,797	2,797	2,797	2,797

Note: The regression specification is an ordered probit. Robust standard errors (clustered by state) in parentheses; * significant at 10%; ** significant at 5%; *** significant at 1%. Includes controls for respondent characteristics, state and year fixed effects. Probability is expressed as a fraction. Marginal effects for one standard deviation change shown.

Figure 5. Associativeness test 1. The solid line represents the point estimates. The dashed lines represent the 95 percent confidence interval

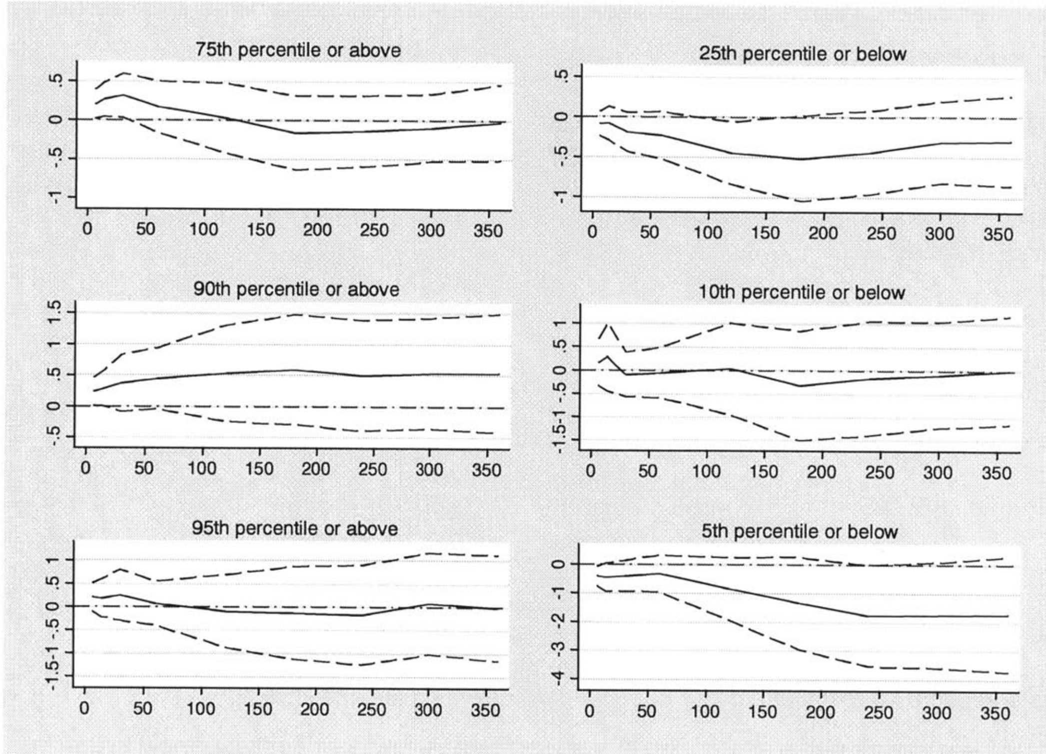

Supplementary Tables and Figures

Table 1: Beliefs about global warming and respondent characteristics

Male	-0.069	Male x college	-0.181
	(0.054)		(0.091)**
White indicator	-0.048	Male x grad. school	-0.189
	(0.055)		(0.066)***
Age	0.019	Very conservative	-0.787
	(0.005)***		(0.078)***
Age2	-2.4E-04	Conservative	-0.372
	(4.8e-05)***		(0.038)***
Log income	0.798	Liberal	0.356
	(0.155)***		(0.055)***
Log income2	-0.040	Very liberal	0.290
	(0.007)***		(0.103)***
Some college	0.107	Male x very conservative	-0.485
	(0.039)***		(0.103)***
College	0.258	Male x conservative	-0.250
	(0.062)***		(0.065)***
Graduate	0.408	Male x liberal	-0.081
	(0.058)***		(0.078)
Male x some college	0.005	Male x very liberal	0.277
	(0.052)		(0.141)**
Observations	7,021		

Regression specification is an ordered probit. Robust standard errors (clustered by state) in parentheses; * significant at 10%; ** significant at 5%; *** significant at 1%. Includes state and year fixed effects. Omitted categories are: moderate, high school education or less, male x moderate, male x high school or less.

Table 2: Beliefs about the state of the economy and local employment conditions

	Average unemployment rate over past year			Unemployment rate last month		
	Panel A: local unemployment rates					
Pr (excellent)	-0.003	-0.002	-0.001	-0.003	-0.001	-1.48E-04
	(0.001)***	(0.001)**	(0.000)*	(0.001)***	(0.001)	(0.000)
Pr (good)	-0.020	-0.013	-0.029	-0.017	-0.007	-0.006
	(0.004)***	(0.006)**	(0.014)**	(0.004)***	(0.004)	(0.014)
Pr (fair)	0.004	0.003	-0.001	0.003	0.001	-1.75E-04
	(0.001)***	(0.001)**	(0.001)	(0.001)***	(0.001)	(0.000)
Pr (poor)	0.019	0.012	0.031	0.016	0.006	0.007
	(0.004)***	(0.004)**	(0.014)**	(0.004)***	(0.004)	(0.014)
Fixed effects	none	state	county	none	state	county
Observations	3,066	3,059	2,102	3,066	3,059	2,102
	Panel B: state unemployment rates					
Pr (excellent)	-0.004	-0.005	-0.002	-0.004	-0.004	-0.002
	(0.001)***	(0.002)**	(0.001)*	(0.001)***	(0.002)**	(0.001)**
Pr (good)	-0.022	-0.031	-0.038	-0.022	-0.027	-0.036
	(0.004)***	(0.012)**	(0.017)**	(0.004)***	(0.012)**	(0.014)**
Pr (fair)	0.005	0.007	0.003	0.005	0.006	0.003
	(0.001)***	(0.003)**	(0.001)**	(0.001)***	(0.003)**	(0.001)**
Pr (poor)	0.021	0.029	0.037	0.021	0.025	0.035
	(0.004)***	(0.010)**	(0.017)**	(0.004)***	(0.010)**	(0.014)**
Fixed effects	none	state	county	none	state	county
Observations	8,132	8,132	3,066	8,132	8,132	3,066

Robust standard errors (clustered by state) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. All regressions control for age and age squared and include the following fixed effects: year, employment status, income category, race, male, political ideology and education. Marginal effects shown. Probability is expressed as a fraction.

Table 3: Effect of other temperature variables on beliefs about global warming

	Change in temperature: between today and yesterday	Change in temperature: between yesterday and two days ago	Average deviation over past week: in degrees	Average deviation over past week: in standard deviations
Pr (Never Happen)	0.001	0.081	0.013	0.245
	(0.032)	(0.028)***	(0.043)	(0.340)
Pr (Happen after lifetime)	0.001	0.074	0.012	0.223
	(0.028)	(0.028)***	(0.039)	(0.310)
Pr (Happen within lifetime)	0.000	0.023	0.004	0.070
	(0.008)	(0.008)***	(0.012)	(0.097)
Pr (Happen within few years)	0.000	0.005	0.001	0.016
	(0.002)	(0.002)***	(0.003)	(0.021)
Pr (Already happening)	-0.002	-0.184	-0.030	-0.554
	(0.071)	(0.067)***	(0.098)	(0.772)
Observations	5,431	5,429	5,468	5,468

The regression specification is an ordered probit. Marginal effects shown. Robust standard errors (clustered by state) in parentheses; * significant at 10%; ** significant at 5%; *** significant at 1%. Includes controls for respondent characteristics, state and year fixed effects. Probability is in percentages.

Table 4: Point estimates for Figures 1-3

	Figure 1		Figure 2		Figure 3	
Quantile	25th or below	75th or above	10th or below	90th or above	5th or below	95th or above
0-7 days ago	-0.045	0.003	0.014	0.031	0.028	0.074
	(0.007)	(0.008)	(0.008)	(0.007)	(0.007)	(0.007)*
0-14 days ago	-0.006	-0.010	0.028	0.051	0.060	0.144
	(0.007)	(0.008)	(0.007)	(0.007)	(0.006)	(0.006)**
0-30 days ago	-0.061	0.026	-0.076	0.154	-0.031	0.233
	(0.006)*	(0.008)	(0.007)	(0.008)*	(0.006)	(0.008)**
0-60 days ago	-0.117	0.081	-0.137	0.186	-0.094	0.271
	(0.007)**	(0.008)	(0.006)*	(0.008)**	(0.006)	(0.007)***
0-120 days ago	-0.224	0.132	-0.329	0.163	-0.306	0.318
	(0.007)***	(0.008)	(0.008)**	(0.008)	(0.008)	(0.007)**
0-180 days ago	-0.225	0.123	-0.379	0.180	-0.425	0.419
	(0.007)***	(0.007)	(0.008)**	(0.008)	(0.008)	(0.008)**
0-240 days ago	-0.145	0.112	-0.245	0.231	-0.164	0.442
	(0.007)*	(0.008)	(0.008)	(0.008)*	(0.008)	(0.007)**
0-300 days ago	-0.164	0.120	-0.256	0.247	-0.194	0.382
	(0.006)**	(0.007)	(0.008)	(0.008)*	(0.007)	(0.008)**
0-360 days ago	-0.172	0.113	-0.260	0.232	-0.215	0.327
	(0.006)*	(0.007)	(0.008)	(0.008)	(0.007)	(0.008)*
Observations	5,468	5,468	5,468	5,468	5,468	5,468

Note: The regression specification is an ordered probit. Robust standard errors (clustered by state) in parentheses; * significant at 10%; ** significant at 5%; *** significant at 1%. Includes controls for respondent characteristics, state and year fixed effects. Probability is expressed as a fraction. Marginal effects for one standard deviation change shown.

Table 5: Effect of longer-run abnormalities on moderate respondents

Quantile	25th or below	75th or above	10th or below	90th or above	5th or below	95th or above
0-7 days ago	0.073	-0.038	0.070	-0.030	0.082	-0.014
	(0.035)**	(0.043)	(0.041)*	(0.046)	(0.061)	(0.059)
0-14 days ago	0.120	-0.025	0.106	-0.014	0.101	0.034
	(0.037)***	(0.052)	(0.065)	(0.059)	(0.074)	(0.081)
0-30 days ago	0.067	0.066	0.003	0.119	0.010	0.129
	(0.048)	(0.079)	(0.086)	(0.093)	(0.108)	(0.141)
0-60 days ago	0.025	0.074	-0.069	0.081	-0.076	0.094
	(0.063)	(0.105)	(0.119)	(0.137)	(0.177)	(0.150)
0-120 days ago	0.026	0.093	-0.107	0.099	-0.174	0.156
	(0.096)	(0.112)	(0.202)	(0.175)	(0.310)	(0.230)
0-180 days ago	0.069	0.049	-0.093	-0.049	-0.142	0.064
	(0.101)	(0.108)	(0.216)	(0.171)	(0.347)	(0.239)
0-240 days ago	0.127	0.059	0.082	-0.014	0.175	0.106
	(0.097)	(0.097)	(0.208)	(0.157)	(0.319)	(0.229)
0-300 days ago	0.111	0.075	-0.005	0.017	-0.020	0.078
	(0.097)	(0.093)	(0.202)	(0.159)	(0.314)	(0.231)
0-360 days ago	0.134	0.058	0.045	-0.018	0.014	0.001
	(0.100)	(0.104)	(0.214)	(0.175)	(0.347)	(0.246)
Observations	2,732	2,732	2,732	2,732	2,732	2,732

Note: The regression specification is an ordered probit. Robust standard errors (clustered by state) in parentheses; * significant at 10%; ** significant at 5%; *** significant at 1%. Includes controls for respondent characteristics, state and year fixed effects. Probability is expressed as a fraction. Marginal effects for one standard deviation change shown.

Table 6: Effect of longer-run abnormalities on liberal respondents

Quantile	25th or below	75th or above	10th or below	90th or above	5th or below	95th or above
0-7 days ago	-0.003	-0.022	-0.003	0.006	0.046	-0.007
	(0.039)	(0.052)	(0.068)	(0.052)	(0.082)	(0.059)
0-14 days ago	-0.037	-0.055	-0.084	-0.023	0.059	-0.013
	(0.059)	(0.072)	(0.119)	(0.059)	(0.145)	(0.071)
0-30 days ago	-0.032	-0.104	-0.034	-0.063	0.078	0.022
	(0.068)	(0.108)	(0.133)	(0.130)	(0.150)	(0.153)
0-60 days ago	-0.101	0.026	-0.092	0.111	0.027	0.224
	(0.076)	(0.119)	(0.141)	(0.122)	(0.188)	(0.143)
0-120 days ago	-0.180	-0.024	-0.304	-0.014	-0.191	0.011
	(0.107)*	(0.130)	(0.228)	(0.174)	(0.303)	(0.209)
0-180 days ago	-0.173	0.015	-0.373	0.096	-0.402	0.276
	(0.115)	(0.123)	(0.259)	(0.167)	(0.393)	(0.254)
0-240 days ago	-0.062	-0.022	-0.198	0.101	0.089	0.308
	(0.119)	(0.125)	(0.250)	(0.201)	(0.365)	(0.282)
0-300 days ago	-0.060	0.001	-0.107	0.240	0.147	0.497
	(0.112)	(0.123)	(0.237)	(0.211)	(0.349)	(0.308)
0-360 days ago	-0.046	-0.026	-0.071	0.211	0.207	0.449
	(0.115)	(0.125)	(0.247)	(0.231)	(0.391)	(0.351)
Observations	1,493	1,493	1,493	1,493	1,493	1,493

Note: The regression specification is an ordered probit. Robust standard errors (clustered by state) in parentheses; * significant at 10%; ** significant at 5%; *** significant at 1%. Includes controls for respondent characteristics, state and year fixed effects. Probability is expressed as a fraction. Marginal effects for one standard deviation change shown.

Fig. A1 Associativeness test 2. The solid line represents the point estimates. The dashed lines represent the 95 percent confidence interval

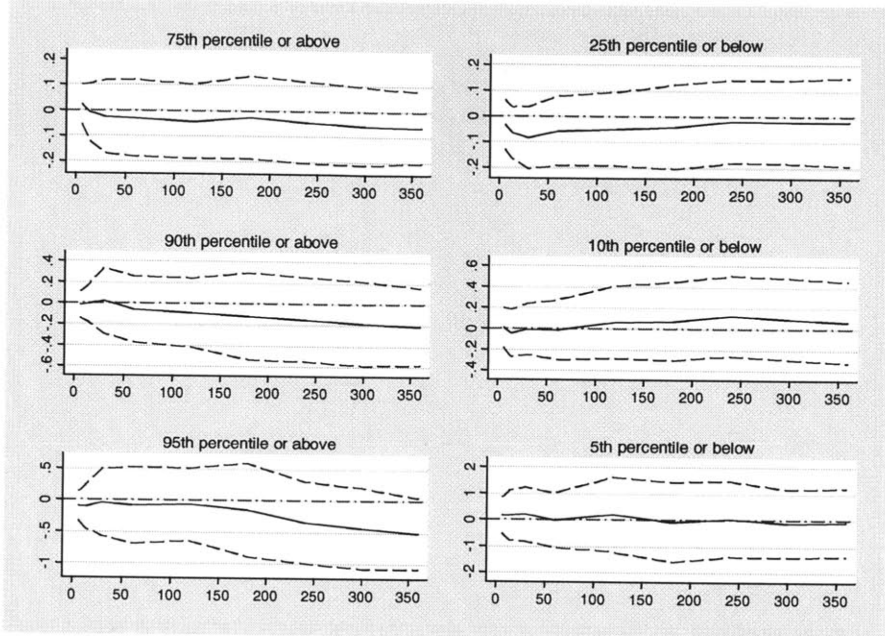

Fig. A2 Associativeness test 3. The solid line represents the point estimates. The dashed lines represent the 95 percent confidence interval

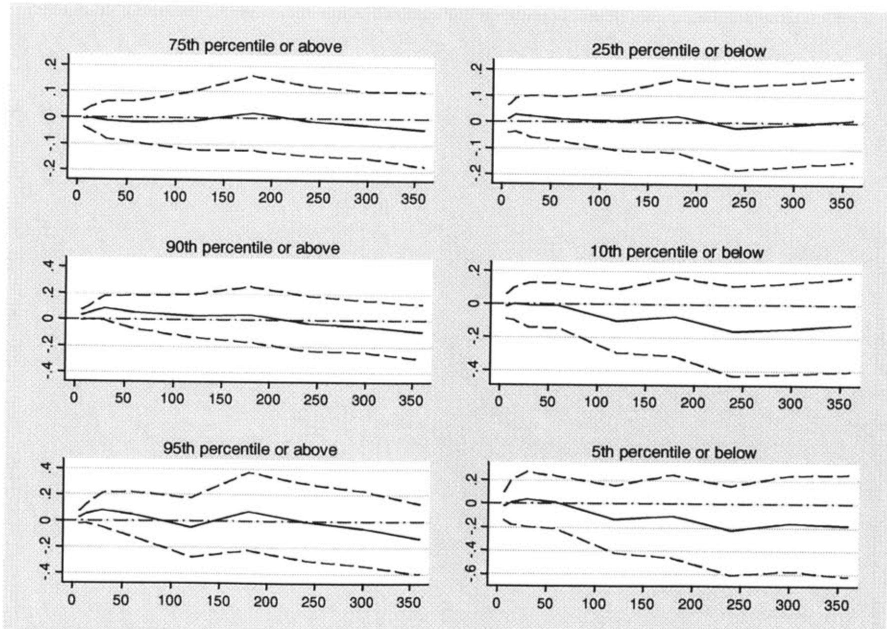

Chapter 2

The Dynamic Effects of Hurricanes in the US: The Role of Non-Disaster Transfer Payments

2.1 Introduction

Extreme weather events are a large and growing source of negative economic shocks. Larger population densities, ecosystem alteration, and population movements to hazardous areas are causing real damages from natural disasters to rise (Board on Natural Disasters, 1999). World insured losses have exceeded \$11 billion per year every year since 1987, reaching \$53 billion in 2004 (Kunreuther and Michel-Kerjan, 2007).[1] Economic losses between 1992 and 2001 averaged \$49 billion a year (Freeman et al., 2003). Damages are likely to continue growing as climate change is expected to increase the number and intensity of extreme events and to change their spatial distribution (Meehl et al., 2007; Schneider et al., 2007). One estimate is that damages will reach \$367 billion a year by 2050, a 750 percent increase in real terms (Freeman et al., 2003). However, the economic impacts of extreme weather are neither predetermined nor random: they depend not only on the meteorological strength of the event, but also on the policies and infrastructure in place (e.g. Zeckhauser, 1996). The exogenous cause of a natural catastrophe is weather, but the difference between an extreme weather event and a disaster is partly man-made. To date, we know very little about the economic impacts of natural disasters over time or the role of institutions and policy in mitigating them.

Governments spend billions of dollars annually on disaster relief and mitigation programs. And, although this is rarely discussed in relation to disaster policy, they also fund transfer programs designed for general economic downturns, such as unemployment insurance, welfare, and food stamps.

[1] Unless stated otherwise, all monetary amounts have been converted to 2008 dollars using the Consumer Price Index. Uninsured losses are difficult to estimate, but a rule of thumb is that they are at least as large as the insured losses in developed countries and at least ten times larger in developing ones.

These may in fact act as a buffer when an extreme weather event occurs, even in absence of direct disaster aid. Ignoring traditional transfer programs would then attribute too much of the resilience of a developed economy to its wealth or disaster-specific response policies. In addition, the fiscal cost of disasters will appear smaller than it actually is.

I study the county-level empirical economic effects of hurricanes, which are one of the most damaging weather events in the US. Specifically, I look at the effects of hurricanes in the 1980's and 1990's from zero to ten years after landfall. I use a simple difference-in-differences framework and focus on changes in population, earnings, employment, and various transfer payments. In addition, I semi-parametrically estimate the post-hurricane economic dynamics, which paints a richer picture of how a county adjusts to this negative shock. My goal is to identify the economic margins along which adjustment takes place (e.g., population movements versus labor market changes) and to understand the role of government spending in post-disaster economics within US counties. I interpret my estimates using a simple spatial equilibrium framework, which suggests that transfers prevent relocation and generally act as a buffer against both disaster and non-disaster negative capital shocks. Some of the results in this paper may apply to capital shocks more generally. The main advantages to using hurricane incidence as an indicator for a capital shock are that hurricanes are exogenous and their onset is known precisely. This is typically not the case with other types of capital shocks.

My results suggest that the potential negative economic consequences of the hurricane may be substantially mitigated through non-disaster social safety net programs. I find that per capita unemployment insurance payments are on average 22 percent higher in the eleven years following the hurricane while overall transfer payments are 2.1 percent higher. Correspondingly, there is no change in population, the employment rate, or wages. In addition to the funds provided through an official disaster declaration, which average \$356 (2008 dollars) per capita per hurricane during my study period, I estimate that in the eleven years following a hurricane, an affected area receives transfers from the government to individuals averaging \$67 per capita per year or about \$640 per capita in present discounted value. Transfers from businesses to individuals (mostly insurance payments) increase temporarily as well, but add only an estimated \$23 to per capita transfers over the eleven years. Together, the transfers represent a large fraction of the immediate damages, which FEMA estimates to be \$1,278 per capita for the major hurricanes during my study period.[2] This suggests that non-disaster policy, in addition to disaster aid and wealth, may be an important factor in explaining the relative resilience to natural disasters in the United States.

My estimates imply that the fiscal impact of natural disasters is nearly twice as large if non-disaster transfers are also considered. Although in the simplest public finance framework transfers are welfare-neutral, in practice the deadweight loss of taxation is estimated to be 12-30% of revenue (Ballard et al., 1985; Feldstein, 1999). Finally, because transfers are not paid for by the people receiving them, they may create moral hazard problems, leading individuals to live in riskier places

[2] Minor hurricanes, which are in my data but not in FEMA's estimates, are generally less damaging.

than they would with actuarially fair insurance. Transfers may be welfare-improving once the hurricane has occurred but their welfare implications are much less clear in the long run.

I consider the effects of hurricanes on the construction sector because it is a proxy for how post-hurricane capital adjustment takes place. Although I do find positive effects on construction wages immediately after the hurricane, employment shrinks three to eight years after the hurricane. Nine to ten years later, there is a sign of an upward movement in employment, suggesting that the decline in the construction sector may be temporary. The decline corresponds to a decrease in new single family home construction. I find suggestive evidence that over time part of the construction sector activity moves to the neighboring unaffected counties. These results suggest that longer-run effects should be an important point of focus when studying the effects of idiosyncratic regional shocks.

I also find evidence of changes in the age structure of the county, but no change in its racial composition. In particular, there is an increase in the fraction of population under 20 years of age and a decrease in the fraction of population 65 and older. However, the pattern of these changes is inconsistent with the transfer increases, implying that the change in transfers is not being driven by changes in the age structure.

Finally, I look at heterogeneity in the impact of hurricanes by the pre-hurricane median income and housing value of a county. I find quantitative as well as qualitative differences between counties in the top and bottom quartiles. Ten years after a hurricane, the increases in per capita unemployment payments and overall transfers are substantially higher in low-housing value counties than in high-housing value counties. In addition, trend break and mean shift tests reveal that, although ten years after a hurricane there is no significant difference in per capita earnings changes between the bottom and top quartiles, there are differences in their post-hurricane paths.

I contribute to two main strands of the natural disaster literature. The first focuses on the economic impacts of natural disasters, typically considering a single outcome or single event (Leiter et al., 2008; Brown et al., 2006) and looking at effects from one to four quarters (Strobl and Walsh, 2008) to three to four years after the event (Murphy and Strobl, 2009). In one of the few studies to consider long-run effects, Hornbeck (2009) finds that the US Dustbowl had persistent effects on land values and land use practices. Belasen and Polachek (2008) estimate that earnings in Florida counties affected by a hurricane increase sharply and remain higher two years after the hurricane. Brown, Mason, and Tiller (2006) estimate that hurricane Katrina had a negative but temporary effect on local employment zero to six months after. Strobl (2008) estimates that coastal counties affected by major hurricanes subsequently experience lower per capita income growth. I add to this literature by looking at a comprehensive set of outcomes for a large sample of disasters over a longer time period and connecting the outcomes together in a cohesive framework.

The second related strand of literature examines the importance of area characteristics, institutions and wealth in determining disaster-related losses and deaths (Kahn, 2005; Skidmore and Toya, 2005; Nordhaus, 2006). Skidmore and Toya (2002) find that a higher frequency of climatic disasters is correlated with higher rates of human capital accumulation. Kahn (2005) finds that a country's institutional quality is inversely related to the number of disaster-related deaths. I contribute to this literature by looking at the economic effects of disasters rather than the damages they cause and by considering the role of transfer payments and within-country heterogeneity.

In the most closely related study, Yang (2008) estimates the effect of hurricanes on international financial flows and finds that four-fifths of the estimated damages are replaced in poorer countries by both international aid and remittances. In richer countries, the increase in lending by multilateral institutions is offset by similar declines in private financial flows. I contribute to this strand of literature by focusing on the role of non-disaster transfer programs in post-disaster economics. Like Yang, I consider the impact of hurricanes on monetary transfers but focus on within-country flows related to social and private insurance.

In addition, there is a literature considering the short-run economic effects of temperature fluctuations (e.g. Dell et al., 2009; Deschenes and Greenstone, 2007a and 2007b; Jones and Olken, 2010) and relating these to climate change. Climate change is forecast to increase the intensity of hurricanes, which, all else equal, will raise the damages they cause by more than one-for-one. In this paper, I underscore the additional fiscal and long-term economic impacts which are currently not incorporated by simple measures of initial damages. My results suggest that climate-induced hurricane intensification may have larger negative consequences than previously thought. This in turn implies that the benefits of mitigation, including policies that diminish the effect of climate change and removal of insurance subsidies, are larger.

The rest of the paper is organized as follows. Section 2 presents the conceptual framework. Section 3 provides background information on hurricanes and US federal disaster aid. Section 4 describes the setting, data and empirical strategy. Sections 5 and 6 present and discuss the results, respectively. Section 7 concludes and contains suggestions for further research.

2.2 Conceptual Framework

Hurricanes in the US can be thought of as negative capital shocks; except for Hurricane Katrina, they do not cause substantial loss of life in the modern US. Thus, I use a simple production function framework to guide the discussion of the results. I describe how economic outcomes evolve following a capital shock under various assumptions about moving costs, capital adjustment costs and the ability of individuals to receive transfer payments instead of working.[3]

Suppose that there are many identical locations, so that changes in one location will not have substantive effects on other locations. Representative firms in each location produce a homogenous good with some standard production function $F(K, L)$, where K is capital and L is labor. Capital and labor are complements. Now suppose that one location experiences a negative capital shock.

[3]For a simple formal model and simulation results, see the online "Model Appendix": http://econ-www.mit.edu/files/6350

Generally, what happens to population, labor supply, and wages depends on capital and individual mobility costs, as well as the presence of unemployment insurance or other transfer programs. If capital is perfectly mobile, a capital shock will have no effect on the equilibrium population or any other economic indicators because adjustment will be immediate. This is regardless of whether there are individual moving costs or transfer programs.

If capital is not perfectly mobile, there will be observed changes in the local economy. If individuals face zero moving costs, there will be no change in the wage, but a decline in the population. This is intuitive: without moving costs, individuals will only stay in the area if they are at least as well off as before. Because the destruction of capital lowers the wage rate, all else equal, individuals will respond by decreasing their labor supply until the wage rate is equal to the pre-shock wage. Because of zero moving costs, decreasing labor supply will be equivalent to moving, as individuals who were choosing to work before will simply costlessly switch to another location. Thus, in the case where capital is not perfectly mobile but individuals are, transfers will play no role in the post-hurricane dynamics. The degree to which population falls depends on how immobile or slow-adjusting capital is.

When both capital and individuals are not perfectly mobile, we expect to see a decline in the wage rate. As long as some of the individuals have negligible moving costs, the population will also fall. Unlike in the previous case, individuals may also decrease their labor supply without moving away, so there may be a decline in the employment rate. The relative decline of population and labor supply depends on the relationship between moving costs and disutility of labor supply. For example, if both moving costs and disutility of labor supply are high, the fall in the employment rate relative to the fall in population will be larger than if moving costs are low.

If, in addition to imperfectly mobile capital and imperfectly mobile individuals, there are transfer payments, the population decline will be weakly smaller than without transfers, while the change in total labor supply and the wage rate relative to the no transfer case is ambiguous. Per capita labor supply should fall more as some individuals take the outside option of transfers instead of working. This will counteract the decrease in wages due to the lower capital. Likewise, some individuals will chose to take transfers and remain in the area instead of moving away.[4] This implies that the net effect on total labor supply is ambiguous: although labor supply per capita is lower than in the no transfer case, there are more people remaining in the area relative to the no transfer case.

In Table 1, I summarize the predictions of this framework following a negative capital shock under various assumptions about the mobility of capital and individuals, as well as the availability of transfer payments. If capital is perfectly mobile (Columns 1 and 2), a negative capital shock will have no effect on any economic indicators, regardless of individuals' mobility costs and transfer availability. If capital is not perfectly mobile, there are no transfers, and moving is costless (first

[4]Transfer payments can be either a decreasing function of the wage (i.e., compensate individuals living in an area for lower wages, as in Notowidigdo, 2010) or unemployment insurance payments that the individual can choose instead of working.

row of Column 3) wages will remain unchanged, but population will fall. In the presence of moving costs but no transfer payments (first row of Column 4), the fall in the population will be smaller, while the decline in wages will be larger. When there are employment-related transfer payments but no moving costs (second row of Column 3), the fall in capital will have the exact same effect as in the no transfer case and the total amount of transfers going to an area will remain unchanged. Finally, when there are both transfer payments and moving costs (second row of Column 4), the fall in utility resulting from a negative capital shock will be buffered by transfers. The presence of transfers will lead some individuals to cease working while remaining in the area, lowering the number of people who leave and causing the drop in labor supply to be larger than the drop in population. The fall in wages will be smaller than in the no-transfer case.

To summarize, if capital is perfectly mobile (or close to it), I expect to find no change in the economy following a hurricane. If capital is somewhat immobile but individual mobility costs are negligible, I expect to find decreases in population but no changes in transfer payments. Finally, if capital adjustment costs and individual moving costs are both non-trivial, transfer payments flowing into the area should increase. The degree to which population falls will reflect both the magnitude of the moving costs and the capital adjustment costs.

The presence of transfer payments weakly increases welfare for individuals living in the area relative to the no transfer case. However, as I discuss later, whether transfer payments increase social welfare is unclear.

2.3 Hurricanes and Federal Disaster Aid

2.3.1 Hurricanes in the United States

Hurricanes that affect the US form in the Atlantic Ocean. The Atlantic hurricane season lasts from June through November, with most hurricanes forming in August and September. Warm humid air over the ocean creates storms known as "tropical disturbances". If circulating winds develop, the disturbance becomes a tropical cyclone. Prevailing winds and currents move the cyclone across the ocean, where it gains and loses strength based on the favorability of conditions. When cyclones encounter cold water or land, they lose strength quickly and dissipate. Sometimes a circular area with low internal wind speeds, called the "eye", develops in the system's center. Although the entire storm system can span a few hundred miles, the perimeter of the eye (the "eyewall") is where the strongest winds are found. Wind intensity declines quickly as one moves away from the eyewall (or the center of the storm, if there is no eye). The outer parts of the hurricane are called "spiral bands"; these are characterized by heavy rains but typically do not have hurricane-force winds. Hurricanes that make it to land create widespread wind and flood damage: physical damages from hurricanes in the US have averaged $4.4 billion per hurricane (2008 dollars) or $7.4 billion per year between 1970 and 2005 and $2.2 billion per hurricane or $3.7 billion per year if 2005 is excluded.[5]

[5] Author calculations using data from Nordhaus (2006). I use 2008 dollars throughout the paper.

For hurricane data, I use the Best Tracks (HURDAT) dataset from the National Oceanic and Atmospheric Administration (NOAA).[6] It contains the location of the storm center and wind speed (in six hour intervals) for each North Atlantic cyclone since 1851. To determine which counties the storm passed through, I assume that the storm path is linear between the given points. Data on storm width are unfortunately not collected; this adds some measurement error. But because the eye of the hurricane is typically not very large, and counties through which the eye passes suffer much more extensive damage (as I show later), this should not be a problem for the estimation.[7] Although the hurricane data span a long time period, annual county-level economic data are only available for 1970-2006. Because the main econometric specification has ten balanced leads and lags (i.e. each lead and lag is estimated using the same set of hurricanes), I estimate the economic effects of hurricanes that occurred between 1980 and 1996.

North Atlantic cyclones are classified by maximum 1-minute sustained wind speeds using the Saffir-Simpson Hurricane Scale. A storm is considered a hurricane if maximum 1-minute sustained wind speeds exceed 74 miles per hour. Category 3 and higher hurricanes have wind speeds greater than 111 mph and are called "major hurricanes". Category 1 and 2 hurricanes are "minor hurricanes", characterized by maximum wind speeds of $74 - 110$ mph. A tropical storm is a cyclone with wind speeds of 39 - 73 miles per hour. Cyclones with lower wind speeds are called "tropical depressions". Between 1980 and 1996, there were on average 5.6 North Atlantic hurricanes per year, with at least two hurricanes each year and three years with ten or more hurricanes. About a third (1.9 out of 5.6) of hurricanes are major hurricanes. Less than a third (1.5 out of 5.6) of all hurricanes that form make landfall, and about half of the landfalling hurricanes (0.7 out of 1.5) are major hurricanes.

US hurricanes are geographically concentrated. Most of the landfalling hurricanes over this time period occur in Alabama, Georgia, Florida, Louisiana, Mississippi, North Carolina, South Carolina, Texas, and Virginia (hereafter the "hurricane region"). Figure 1 shows the geographic distribution of hurricane hits that occurred between 1980 and 1996, as well as the control counties used in subsequent analysis (selected using propensity score matching). Out of the hurricane region counties, 127 experience one or more hurricanes between 1980 and 1996 (119 experience only one hurricane). Only 19 counties outside the hurricane region experience any hurricanes during this time and virtually all the major hurricanes occur within the 9 states listed above. I therefore limit my analysis to this region. Although it may be preferable to focus on the major hurricanes, they are relatively rare (there are only 8 between 1980 and 1996). For this reason, I focus on the 21 minor and major hurricanes that affected the hurricane region during that time.

[6] Available from http://www.nhc.noaa.gov/pastall.shtml#hurdat

[7] See Appendix A for a discussion of the distribution of eye diameters.

2.3.2 Destructiveness of Hurricanes

In order to gauge the potential economic impact of hurricanes, it is helpful to look at the damages they cause in absolute terms and relative to other US disasters.

To provide evidence on the absolute level of damages caused by hurricanes, I use estimates of direct damages from HAZUS-MH, published by FEMA.[8] Table 2 shows the summary statistics of the effects of the eight major hurricanes that affected the hurricane region between 1980 and 1996. HAZUS-MH is software meant to help state, local, and Federal government officials prepare for disasters and to help the private sector estimate risk exposure. The software combines scientific and engineering knowledge with detailed historic data to produce damage estimates that are likely to be more accurate than those made using simpler estimates or reports. In addition to simulating hypothetical damages, HAZUS contains highly detailed damage estimates of past major hurricanes. These damage estimates are shown in Table 2.

Panel A summarizes the estimated effects in the counties which, according to the Best Tracks data, were in the path of the hurricane's center (I refer to these as "centrally affected" counties). On average, these counties suffered \$406 million in damages to buildings (with a standard deviation of about \$2 billion) or about 1.46% (with a standard deviation of 3.85%) of the total building value. The maximum county-level building damage was \$20 billion while the maximum loss as a percent of total building value was 23.6%.

HAZUS-MH also provides estimates of non-structural losses, such as building content and inventory losses, as well as estimates of the number of households displaced by the disaster. Total losses (including building damages) average \$571 million per county with a standard deviation of \$3.7 billion. The largest total loss on a county level over this period was \$35.4 billion. On average, about 1,500 households (with a standard deviation of 10,700) are displaced as a result of a central hit by a major hurricane and 450 people require temporary shelter. Per capita total damages average \$1,280 with a standard deviation of about \$3,340.

Panel B shows the estimated effects of the hurricane on counties that are listed as affected in the FEMA simulations but do not have the center of the storm passing through them ("peripherally affected" counties). The damage estimates are much smaller. For example, the average damage to buildings is only \$8.6 million or about 65 times smaller than the average damage in a centrally affected county. The *maximum* damage in peripherally affected counties is \$390 million, which is smaller than the *mean* damage in centrally affected counties. The average loss ratio is 0.15%, which is about 10 times smaller than the loss ratio in centrally affected counties. Per capita total losses are also about 10 times smaller, averaging \$113 per capita, and total losses are about 50 times smaller. Only 12 households are estimated to be displaced, on average, and only 3 people require temporary shelter. Thus, although the omission of these counties from the analysis may introduce some measurement error, it should not affect the estimates much.

The above estimates provide evidence both on the level of a hurricane's damage and on the

[8] Available by request from http://www.fema.gov/plan/prevent/hazus/index.shtm

likely importance of including counties not directly in the storm's path. It should be noted that the damage estimates are an upper bound on the average destructiveness of the hurricanes in my sample because my sample includes minor as well as major hurricanes. Unfortunately, FEMA does not provide detailed damage estimates for minor hurricanes. A theoretical result is that the energy carried by the wind increases with the third power of wind speed. The average maximum wind speed in a county that was centrally affected by a major hurricane between 1980 and 1996 is 124 miles per hour, while the average maximum wind speed in a county centrally affected by a minor hurricane is 86 miles per hour. If the power carried by the wind translates directly into destructiveness, a back of the envelope calculation implies that a 124 miles per hour hurricane would cause about three times more damage than an 86 miles per hour hurricane. This, in turn, would imply that the average minor hurricane in my sample caused about \$190 million in total damages per centrally affected county. Although this is not as large as the damage caused by major hurricanes, it is a non-trivial amount for a local economy and may affect subsequent economic outcomes.

I now address the relative damages caused by hurricanes. I regress three different damage statistics on measures of hurricane strength and other natural event indicators. The regression specifications are as follows:

$$\begin{aligned} D_{ct} &= a_c + a_t + \beta_1 Major\_hurricane_{ct} + \beta_2 Minor\_hurricane_{ct} \\ &\quad + \gamma_1 Flood_{ct} + \gamma_2 Tornado_{ct} + \gamma_3 Severe\_storm_{ct} + \varepsilon_{ct} \end{aligned}$$

and

$$\begin{aligned} D_{ct} &= a_c + a_t + \sum_{k=1}^{5} \beta_k \mathbf{1}\left[Category_{ct} = k\right] + \gamma_1 Flood_{ct} \\ &\quad + \gamma_2 Tornado_{ct} + \gamma_3 Severe\_storm_{ct} + \varepsilon_{ct} \end{aligned}$$

$$c = \text{ county; } t = \text{ year}$$

D_{ct} is log of property damages, property damages per capita or the log of flood insurance payments in that county. [9] All damage measures are in 2008 dollars. $Major\_hurricane_{ct}$ is an indicator for Category 3, 4, and 5 storms, while $Minor\_hurricane_{ct}$ is an indicator for Category 1 and 2 storms. $\mathbf{1}\left[Category_{ct} = k\right]$ is an indicator variable equal to 1 if the hurricane is classified as a Category k hurricane. Because there are very few Category 4 and 5 hurricanes, I combine

[9]Data on damages and extreme weather events other than hurricanes are from the Hazards & Vulnerability Research Institute (2009) and are based on weather service reports by local government officials. Data on flood claims and liabilities are from the Consolidated Federal Funds Report (CFFR).

them in the second equation. The *Flood*, *Tornado*, and *Severe_storm* indicators are equal to 1 if the county was reported as having at least one of these events over the year. These, along with hurricanes, are the most common and damaging meteorological events in the US. Other rarer events in the region include droughts, wildfires, and heat. Thus, the reference category is a combination of these extreme events and no reported extreme events. Finally, a_c and a_t are county and year fixed effects.

I estimate these two equations for the nine states in the hurricane region.[10] The results are shown in Table 3. Column 1 compares the log of damages for different disasters. A major hurricane increases the reported property damages by 4.2 log points or over 400%. In levels, this implies that a major hurricane increases the total damages in a county by about $\$760,000$ (2008 dollars). The next most damaging event is a minor hurricane, which increases property damages by 2.4 log points or about $\$110,000$. In contrast, tornadoes, floods, and severe storms increase property damages by 2.1 ($\$76,000$), 0.9 ($\$15,000$), and 1.0 ($\$18,000$) log points (dollars), respectively. A similar pattern holds when the dependent variable is property damages per capita, although some of the point estimates become statistically insignificant. This is possibly because hurricane-prone counties are more populous. Column 4 shows the effect of hurricanes broken down by category. As expected, Category 1 hurricanes are the least damaging, causing an extra 2.2 log points of damage ($\$84,000$), while Category 4 and 5 storms are the most damaging, increasing property damages by 4.6 log points ($\$1,100,000$). The least damaging hurricane is about as damaging as a tornado, and more damaging than a flood or severe storm.

An important caveat is that the damage measures are estimates made by local officials soon after the occurrence of the event. Using hurricane-level damage data from Nordhaus (2006), I estimate the direct damages from hurricanes to be about $3.7 billion per year between 1970 and 2004, in 2008 dollars. Given that there are on average 1.5 landfalling hurricanes per year, the estimates in this section appear to understate the per-county damage of hurricanes (and possibly of other disasters as well) by at least a factor of ten. However, as long as the damage measurements do not exhibit differential bias for hurricanes, floods, storms, and tornadoes, these numbers are valid for comparing the *relative* magnitudes of the different events.

Column 3 shows the effect of various extreme weather events on flood payments. Major hurricanes increase flood claims by about 3.1 log points or about $1.1 million, while minor hurricanes increase them by 1.5 log points or about $\$190,000$. The mean insurance liability in the sample is $538 million. Tornadoes have no significant impact on flood claims and the estimated effect of a severe storm is significantly negative.[11] Floods increase claims by only about 0.5 percentage points.

When the effect of a hurricane is broken down further, Category 3 storms are estimated to have the largest effect, raising flood insurance payments by about 3.1 log points. Category 1 and 2

[10] The results for all US counties are similar.

[11] The comparison category is not "no extreme weather event", but a combination of this indicator and other, rarer, weather events. Some of these, such as heat waves, may be more damaging than the average severe storm.

hurricanes raise flood-related insurance payments by 1.1 and 2.8 log points, respectively. Category 4 and 5 storms increase them by 3 log points.

The flood insurance payments are likely to be a lower bound on total insurance payments for two reasons. First, in addition to flood damage, the wind associated with hurricanes creates massive damage, which is covered by homeowner's insurance. Second, the fiscal year of the US government ends on September 30th. Some flood insurance claims originating in August and September (the peak hurricane time) may be settled in the same fiscal year, while some may not appear until the following year. Despite all the caveats, these estimates imply that hurricanes are the most destructive of the common US disasters, which makes them an important phenomenon to study.

2.3.3 Federal Disaster Aid

This section summarizes US federal disaster spending between 1980 and 1996. Federal disaster aid is given to a county if the state's governor files a request and provides evidence that the state cannot handle the disaster on its own. The final decision about whether to declare a disaster is made by the US President. If the request is approved, federal money can be used to repair public structures and to make individual and business grants and loans. Grants to individuals are made only up to the amount of uninsured damages. The Federal Emergency Management Agency (FEMA) also provides personnel, legal help, counseling, and special unemployment insurance for people unemployed due to the disaster. Although there is some long-term recovery spending in extreme cases, most of the transfers to individuals occur within six months of the declaration and most of the public infrastructure spending occurs within two-three years (FEMA, personal communication).

Between 1980 and 1996, the federal government spent \$6.4 billion (2008 dollars) on hurricane-related disaster aid and \$23.1 billion on other disasters.[12] The bulk of the non-hurricane disaster spending (\$10.1 billion) was due to the Northridge earthquake in 1994. Excluding the Northridge earthquake implies that hurricane-related spending accounts for about a third of all disaster aid. This number includes all declaration-related spending by FEMA, including assistance given for infrastructure repair, individual grants, as well as mitigation spending. The Small Business Administration also offers subsidized loans to affected individuals and businesses, which are not included here. Spending by the state and local governments is also excluded. By law, the state pays some of the cost of disaster aid, but its share cannot exceed 25%. Thus, state spending comprises at most a third of the federal spending. Unfortunately, annual county data on disaster spending over time is not available, so I cannot incorporate disaster spending into my main empirical framework. However, there are data that allow me to directly compute the county-by-hurricane amount of disaster transfers.

Table 4 shows the summary statistics for federal aid related to hurricanes between 1980 – 1996.[13] Because data on federal disaster aid is provided on the level of a declaration, which includes

[12]Data on spending are from the PERI Presidential Disaster Declarations database (Sylves and Racca, 2010).

[13]Summary statistics for other times periods are similar, with the caveat that real spending on hurricane-related

multiple counties in a state, an assumption about how the money is divided among counties is necessary. As I show in the previous section, counties through which the center of the storm passes experience much more damage than peripherally affected counties. Therefore, a natural assumption is that the money is split among only those counties and the rest can be ignored. Another natural assumption is that the money is divided among the included counties in proportion to the population in each county. Panel A shows the total and per capita federal aid transfers assuming that only centrally affected counties are given aid. The average amount of aid given to counties experiencing hurricanes was \$58.7 million. Counties experiencing major hurricanes received about 2.5 times as much on average, \$128 – 133 million. The standard deviations of aid for counties that experienced hurricanes are all larger than the mean, ranging from \$187 to over \$460 million. Note that this period excludes Hurricane Katrina and the 2004 hurricane season, in which four hurricanes affected Florida. Thus, even "business as usual" hurricane seasons are associated with non-trivial amounts of federal spending.

Per capita spending in 1980-1996 averaged \$356 per hurricane and \$412 per major hurricane (2008 dollars). An extreme assumption of a uniform split across counties (which is unlikely to be true) leads to a larger per-capita average of \$1,137 per hurricane and \$2,018 per major hurricane.

Panel B shows the same statistics assuming that the money is divided among all counties included in the declaration, not just centrally affected ones. This implies spending of \$8.4 – 8.9 million per county, \$24.6 – 30.0 million per centrally affected county, and \$59.2 – 73.4 million per county centrally affected by a major hurricane. Per capita spending estimates range from \$52 to \$187 in the proportional split case and from \$160 to \$954 in the uniform split case.

In the results section, I use the preferred number of \$356 per capita as a benchmark to compare spending by disaster relief agencies to the extra spending associated with the hurricane triggering other transfer programs.

2.4 Empirical Strategy

2.4.1 Sample of Analysis

Ideally, one would estimate the effect of hurricanes by looking at differences over time between counties in the hurricane region that do and do not experience a hurricane between 1980 and 1996. However, finding a valid control group is not straightforward. In Table 5, I compare characteristics and trends of counties that do not experience any hurricanes between 1980 and 1996 with counties that experience one hurricane.[14]

Columns 1 and 2 of Panel A show the 1970 characteristics of hurricane region counties that experience a hurricane between 1980 and 1996 and the difference from counties with no hurricanes.

declarations is rising over time.

[14]I omit the few counties that experience more than one hurricane between 1980 and 1996. Results are similar if counties with more than one hurricane are included.

Nearly fifty percent of 119 counties that experience one hurricane are coastal, compared to forty-one percent of 811 counties that have not had hurricanes over this period. Counties that experience hurricanes are about fifty percent more populous than non-hurricane counties and have lower population densities. These differences are statistically significant (as shown in Column 3). Counties with hurricanes also have larger per unemployment insurance payments, but smaller per capita transfers from the federal government. However, differences in levels are not problematic because county fixed effects are included in every specification. However, differences in levels may be indicative of differences in trends. In Panel A, I test for differential trends between 1970 and 1979 (before any hurricanes in the sample occur) for the time-varying characteristics. In this case, Columns 1 and 2 show the trend in the hurricane counties and their difference from the trend in non-hurricane counties. Only two variables show differential trends for these two groups of counties: per capita transfers from government and per capita family assistance, both significant at the 5% level.

Another way to construct the control group is by requiring balance in pre-hurricane covariates and hurricane risk.[15] I construct a hurricane risk variable using historic (1981-1970) hurricane data. I predict counties' propensity to be hit by hurricanes by spatially smoothing observed hurricane hits. I then use two-nearest neighbor propensity score matching to select a control group from the no-hurricane sample.[16] Column 4 shows the difference between the hurricane counties and the propensity matched control group, while Column 5 shows the p-value of this difference. In general, propensity score matching eliminates differences in levels and trends for all variables except population, whose trend and level differences continue to be significant at the 5% level. Because the sample in Column 4 is more similar to the treatment group than the sample in Column 2, I use the former as my preferred control group.

I discuss results using other samples in the robustness section (including using only the counties that experience a hurricane between 1980 and 1996). I show that these do not affect the estimates qualitatively and have only a moderate quantitative effect. I also address the problem of potentially different time trends by relying on mean shift and trend break tests.

2.4.2 Economic and Demographic Data

Annual county-level outcomes such as unemployment payments, population, and earnings come from either the Regional Economic Information System (REIS), while sector-specific employment, wages and number of establishments come from County Business Patterns (CBP). County-level population by race and age are from Surveillance Epidemiology and End Results (SEER) population database. Data on single family housing starts are from McGraw-Hill. All four series span the years 1970-2006.

---

[15]Matching is based on all outcome variables, although some are not shown due to space constraints.

[16]Using two-nearest neighbor rather than nearest neighbor matching ensures that the number of counties in the control group is approximately equal to the number of counties in the treatment group. When nearest neighbor matching is used, some non-hurricane counties are assigned as nearest neighbors multiple times, resulting in a control group that's much smaller than the treatment group.

I define the employment rate as the ratio of total employment to the number of people aged fifteen and older.[17] An establishment is defined as a single *physical* location of a firm with paid employees. Net earnings by place of residence (which I later refer to as simply "net earnings") include wage and salary disbursements, supplements to wages and salaries, and proprietors' income, less contributions for government social insurance. Earnings do *not* include transfer payments. Earnings by place of work are converted to earnings by residence by the Bureau of Economic Analysis (BEA) using a statistical adjustment. On average, the construction sector represents slightly over 10% of all establishments, employees, and wages.

Unemployment insurance compensation consists primarily of standard state-administered unemployment insurance schemes, but also includes unemployment compensation for federal employees, railroad workers, and veterans. Total transfers from government to individuals include unemployment insurance. In addition, the category includes income maintenance (e.g., Supplemental Security Income or SSI), family assistance, retirement and disability insurance benefits, medical benefits (Medicare and Medicaid), veterans' benefits, and federal education and training assistance. Transfers from businesses to individuals consist primarily of net insurance settlements and personal injury liability payments to non-employees.

Disaster-related transfers are technically included in the measure of total transfers from the government. However, these are computed by assuming that national estimates are distributed in proportion to population to all the counties in the US. Thus, these will not affect the estimation once year fixed effects are included.

2.4.3 Event Study Regression Framework

In this section, I outline the procedure used to estimate the economic effects of a hurricane. I first employ an event study framework. Specifically, I regress outcomes on hurricane indicators 10 years before and after a hurricane, controlling for county, year, and coastal-by-year fixed effects. It would be ideal to estimate the effects of major and minor hurricanes separately, but there are too few major hurricanes for a precise estimation of their effect.[18] Thus, I focus on the effect of all hurricanes. The identifying assumption is that, conditional on the location and the year, the occurrence of a hurricane is uncorrelated with unobservables. This is reasonable because even forecasting the severity of the hurricane season as a whole is difficult, much less the paths those hurricanes will take. Although there is no cause to believe that hurricanes are endogenous when proper controls are included, I estimate the leads of a hurricane to test for the presence of differential trends.

The basic event study framework for estimating the year-by-year effect of a hurricane up to ten years after its occurrence is:

[17] Annual county-level unemployment rates are not available until 1990.

[18] If I restrict the sample to estimate ten leads and lags using the same county-hurricane-year observations, I end up with less than 30 counties that experience major hurricanes. In contrast, there are 119 counties that experience a major or minor hurricane when the same restrictions are imposed.

$$O_{ct} = \sum_{\tau=-10}^{10} \beta_\tau H_{c,t-\tau} + \beta_{ct}^{-11} + \beta_{ct}^{11} + \alpha_c + \alpha_t + 1\,[coastal]\,\alpha_t + \varepsilon_{ct} \tag{2.1}$$

$$c = \text{county};\ t = \text{ year};\ \tau = \text{ lag}$$

O_{ct} is some economic outcome, as described in the data section. H_{ct} is a hurricane indicator, equal to 1 if the county is reported to have experienced any hurricane in year t, according to the NOAA Best Tracks data. I normalize the effect the year before the hurricane, $\tau = -1$, to zero. $\{\beta_{ct}^{-11}, \beta_{ct}^{11}\}$ are indicators for hurricanes outside the estimation window. α_c and α_t are county and year fixed effects.

$1\,[coastal]\,\alpha_t$ is a set of year fixed effects for coastal counties, as defined by the NOAA's Strategic Environmental Assessments Division. Including this interaction term is necessary because coastal counties are more likely to experience hurricanes *and* may experience different growth trajectories. For example, the population data show that the coastal population has grown disproportionately in the past 30 years.

I combine hurricane indicators into two-year bins to increase the power of the estimation.[19] The combined lags are years 1 and 2, 3 and 4, 5 and 6, 7 and 8, 9 and 10. The combined leads are -1 and -2, -3 and -4, -5 and -6, -7 and -8, -9 and -10. The assumption needed for this estimation procedure to be valid is that the effects of a hurricane for the years that are grouped together have the same sign and distribution. Year 0, which is the year that the hurricane makes landfall in a county, is not combined because the assumption that the effects in year 0 and year 1 are similar may not hold.

$$\begin{aligned} O_{ct} \;=\;& \beta_0 H_{ct} + \sum_{\tau=-5}^{-1} \beta_{-2\tau} \max\{H_{c,t-2\tau}, H_{c,t-2\tau-1}\} \\ &+ \sum_{\tau=1}^{5} \beta_{-2\tau} \max\{H_{c,t-2\tau+1}, H_{c,t-2\tau}\} \\ &+ \beta_{ct}^{-11} + \beta_{ct}^{11} + \alpha_c + \alpha_t + 1\,[coastal]\,\alpha_t + \varepsilon_{ct} \end{aligned} \tag{2.2}$$

The coefficient β_0 corresponds to year 0, which is the year in which the hurricane makes landfall in the county. For example, 1989 is year 0 for Hurricane Hugo, one of the hurricanes in my sample, and 1992 is year 0 for Hurricane Andrew.

The notation for the hurricane bins is unconventional, but straightforward. $\max\{H_{c,t-2\tau+1}, H_{c,t-2\tau}\}$ takes the maximum of the county's hurricane indicators in subsequent years, grouping them as described above. The set $\sum_{\tau=1}^{5} \beta_{-2\tau} \max\{H_{c,t-2\tau+1}, H_{c,t-2\tau}\}$ thus represents the causal effects of a

[19] Results using year-by-year hurricane indicators are qualitatively similar, but noisier. The full set of results is available upon request.

hurricane 1-10 years following its occurrence. It can be written out as $\beta_{-2} \max\{H_{c,t-1}, H_{c,t-2}\} + \beta_{-4} \max\{H_{c,t-3}, H_{c,t-4}\} + ... + \beta_{-10} \max\{H_{c,t-9}, H_{c,t-10}\}$. The reference category is "hurricane one or two years from now", corresponding to $\max\{H_{c,t+1}, H_{c,t+2}\}$. The coefficients of interest is the set of hurricane lags $\{\beta_{-2\tau}\}_{\tau=1}^{\tau=5}$ and the estimated immediate impact of a hurricane, β_0. The average effect of combined years -1 and -2 is assumed to be 0, so the estimated coefficients should be interpreted as the change relative to the two years before the hurricane.

I do not use damages estimates as the independent variable for several reasons. County-level property damage estimates between 1960 and 2009 are available from the Spatial Hazard Events and Losses Database (SHELDUS).[20] To my knowledge, this is the only database that contains county-level damage estimates for all hurricanes over this period of time. However, the data are estimates made by local emergency officials fairly close to the time of occurrence. At best, they appear to be very imprecise, as discussed in Section 3. Second, damage is not only a function of the hurricane's strength, but of local characteristics such as construction practices and population density, which may be correlated with economic trajectories. Finally, damages may be endogenous with respect to the variable of interests. For example, communities with lower chances of recovery may be damaged relatively more because of poor construction. The county with heavier damages, all else equal, may be in decline or may be less prepared to deal with the disaster overall. Alternatively, the county with larger absolute damages may be more affluent and able to recover more quickly (for example, because of better access to credit, coordination, or governance).

Because there may be unobserved heterogeneity across hurricanes, I also restrict the sample of hurricanes to those for which I can estimate the full set of leads and lags. In practice, this means I am estimating the effects using hurricanes that occurred between 1980 and 1996. To maximize my sample size, I create indicator variables for the county 10 years before and after it experienced a hurricane that was taken out of the sample (i.e., counties that were affected between 1960-1979 and 1997-2006). This allows me to exclude certain county-year observations from the estimation without excluding the county completely. I also restrict my sample to counties that have a continuous record for each outcome variable in order to avoid biasing my results.

Many of the outcome variables are autocorrelated as well as correlated with each other. Appendix B shows the empirical auto- and cross-correlation in the outcome variables. The autocorrelation creates multicollinearity concerns, which is why it is useful to rely on joint tests of significance to determine whether there are significant effects.

2.4.4 Differences in Differences Regression Framework

The basic results suggest that hurricanes may have an effect on the mean of the economic variable, its trend, or both. In addition to estimating the effect for each time period, I also test for trend breaks and mean shifts in the outcome variable. The trend break specification tests for a change in the slope of the economic outcome after the hurricane, while the mean shift specification tests for

[20] Hazards & Vulnerability Research Institute (2009). Available from http://webra.cas.sc.edu/hvri/products/sheldus.aspx

a change in the mean, assuming that there is no change in trend. These specifications summarize the net effect of a hurricane more concisely and are more powerful when the assumption of linear trends holds. In addition, if the assumption of parallel trends does not hold, the trend break test is useful for determining whether the hurricane has a significant impact on the economy.

The regression equation for testing for a mean shift controlling for an overall time trend is:

$$\begin{aligned} O_{ct} \quad = \quad & \theta_1 * \mathbf{1}[\text{Hurr in past 10 years}]_{ct} + \beta^{11} post11_{ct} + \beta^{-11} pre11_{ct} \\ & + \gamma_1 \mathbf{1}[\text{Hurr within 10 years}]_{ct} t + \gamma_2 \mathbf{1}[\text{Hurr outside 10 years}]_{ct} t \\ & + \alpha_c + \alpha_t + 1\left[coastal\right]\alpha_t + \varepsilon_{ct} \end{aligned} \tag{2.3}$$

O_{ct} is some economic outcome, such as population or the employment rate. $\mathbf{1}$[Hurr in past 10 years]$_{ct}$ is an indicator variable equal to 1 if county c has experienced a hurricane in the ten years prior to and including t. Thus, θ_1 is the variable of interest, representing the average change in the outcome in the eleven years after the hurricane (the year of the hurricane and ten subsequent years).

Because my data span a large time period, including a single linear trend variable may be overly restrictive. Thus, I separately control for the trend in the ten years before and eleven years following and including the hurricane year with the variable $\mathbf{1}$[Hurr within 10 years]$_{ct}t$. $\mathbf{1}$[Hurr within 10 years]$_{ct}$ is an indicator equal to 1 if county c experienced a hurricane in the ten years before or in the ten years after time t. γ_2, the coefficient on $\mathbf{1}$[Hurr outside 10 years]$_{ct}$, controls for the overall trend in hurricane counties outside of the twenty-one year window of interest.

I include indicator variables $post11_{ct}$ and $pre11_{ct}$ to ensure that I am comparing the eleven-year post-hurricane mean to the ten-year pre-hurricane mean. These are equal to 1 if county c in year t experienced a hurricane eleven or more years ago or will experience a hurricane eleven or more years in the future. As before, I control for county, year, and coastal-county-by-year fixed effects with α_c, α_t, and $1\left[coastal\right]\alpha_t$.

The growth rate in outcomes may also be affected by a hurricane. To test for a change in the linear trend following a hurricane (i.e., a trend break model), I add an additional variable to the equation above:

$$\begin{aligned} O_{ct} \quad = \quad & \theta_1 * \mathbf{1}[\text{Hurr in past 10 years}]_{ct} + \theta_2 * \mathbf{1}[\text{Hurr in past 10 years}]_{ct} t \\ & + \gamma_1 \mathbf{1}[\text{Hurr within 10 years}]_{ct} t + \gamma_2 \mathbf{1}[\text{Hurr outside 10 years}]_{ct} t \\ & + \beta^{11} post11_{ct} + \beta^{-11} pre11_{ct} + \alpha_c + \alpha_t + 1\left[coastal\right]\alpha_t + \varepsilon_{ct} \end{aligned} \tag{2.4}$$

$\mathbf{1}$[Hurr in past 10 years]$_{ct}t$ is the interaction of the eleven-year post-hurricane indicator with year. As above, $\mathbf{1}$[Hurr within 10 years]$_{ct}t$ controls for the average trend in the ten years before and ten years after the hurricane. Because I want to compare trends ten years before the hurricane to eleven years after, I include indicators for hurricanes ($post11_{ct}$ and $pre11_{ct}$) as well as linear hurricane-specific trends ($\mathbf{1}[\text{Hurr outside 10 years}]_{ct}t$) outside of this window of interest.

The test for a mean shift without a trend break amounts to testing $\theta_1 = 0$ in equation (2.3), while the test for a mean shift with a trend break amounts to testing $\theta_1 = 0$ and $\theta_2 = 0$ in equation (2.4). For the trend break test, I also calculate the hurricane-driven change in the outcome five years after the hurricane (the year of the hurricane and four subsequent years) and eleven years after the hurricane (the year of the hurricane and ten subsequent levels). This is equivalent to calculating $\theta_1 + 5 * \theta_2$ for the five-year change and $\theta_1 + 10 * \theta_2$ for the eleven-year change. Note that the mean shift test restricts the hurricane-driven change to be identical in each year following the hurricane.

2.4.5 Heterogeneity of Effects by Wealth

Understanding the determinants of the post-disaster economic trajectory is important for policy design. The wealth of an area, such as income and house values, is likely to be important in determining how its economics are affected by a hurricane. The poor typically have lower access to credit. If they cannot borrow, their labor supply or mobility response following a capital shock may differ from richer individuals. Specifically, credit constraints can cause the poor to supply labor inefficiently or prevent them from moving, which exacerbates the negative welfare effect of the initial shock. Other factors can also be at play: for example, Masozera et al. (2007) find that poor neighborhoods are less likely to have flood insurance and vehicles, suggesting that they may have a harder time dealing with the disaster's aftermath.

Whether wealth is measured by house values or median income may matter for the estimated heterogeneity in post-hurricane economics because hurricanes destroy housing. The median home value may be a good proxy for the absolute level of the wealth shock experienced by an area's residents. Income could be an important predictor of post-disaster economics because it may proxy for borrowing constraints, among other things.

To look at the effects of wealth on post-hurricane dynamics, I interact the county's quartile for (a) 1970 median housing value and (b) 1970 median income with the hurricane indicator ten years before and after its occurrence. The data on income and housing values are from the Census. As in the main trend break and mean shift specifications, I compare the means and trends of low and high-income counties before and after the hurricane. First, I estimate a mean shift model that allows for an overall time trend, but has no differential time trends after the hurricane:

$$\begin{aligned} O_{ct} = \; & \theta_1^{TOP} * \mathbf{1}[\text{Hurr in past 10 years}]_{ct} * TOP_{1970}^c \\ & + \theta_1^{BOT} * \mathbf{1}[\text{Hurr in past 10 years}]_{ct} * BOT_{1970}^c \\ & + Controls_{ct} + \varepsilon_{ct} \end{aligned} \tag{2.5}$$

O_{ct} is some economic outcome, as before. TOP^c_{1970} is an indicator equal to 1 if the county was in the top quartile in 1970 while BOT^c_{1970} is the corresponding indicator for being in the bottom quartile. Thus, the changes in the mean are relative to counties that are in the two middle quartiles. To test for a differential change in the mean, I compare the estimated mean shift for the counties in the upper quartile of income (θ_1^{TOP}) to the mean shift in the bottom quartile (θ_1^{BOT}). Specifically, I compute $\theta_1^{BOT} - \theta_1^{TOP}$ and whether this is statistically different from 0. $Controls_{ct}$ ensures that I am comparing the ten year pre-hurricane means to the eleven year post-hurricane means. It includes a quartile-specific trend variable for the 21-year window around the hurricane (minus ten years to plus ten years), as well as a set of year, county, and coastal-by-year fixed effects, indicators for hurricanes outside the time window of interest, and trends outside the window of interest.

In order to test for a trend break, it is necessary to add two more variables which capture post-hurricane changes in the trend by quartile:

$$\begin{aligned} O_{ct} &= \theta_1^{TOP} * \mathbf{1}[\text{Hurr in past 10 years}]_{ct} * TOP^c_{1970} \\ &\quad + \theta_1^{BOT} * \mathbf{1}[\text{Hurr in past 10 years}]_{ct} * BOT^c_{1970} \\ &\quad + \theta_2^{TOP} \mathbf{1}[\text{Hurr in past 10 years}]_{ct} * TOP^c_{1970} t \\ &\quad + \theta_2^{BOT} \mathbf{1}[\text{Hurr in past 10 years}]_{ct} * BOT^c_{1970} t \\ &\quad + Controls_{ct} + \varepsilon_{ct} \end{aligned} \tag{2.6}$$

To test for a differential change in the top and bottom quartile counties, I compare the 10-year change in the top quartile ($\theta_1^{TOP} + 10 * \theta_2^{TOP}$) to the 10-year change in the bottom quartile $\left(\theta_1^{BOT} + 10 * \theta_2^{BOT}\right)$. As in trend break equation (2.4), $Controls_{ct}$ ensures that I am estimating trend changes in the eleven years after the hurricane relative to the ten years before and includes the same set of fixed effects and quartile-specific trends. In addition to the variables included for equation (2.5), it includes quartile-specific post-hurricane indicators to allow for a mean shift.

Due to the small sample of hurricanes and affected counties, it is difficult to estimate the importance of these variables precisely, so these results should be taken as suggestive. Note that the quartile indicators also capture other differences between areas, such as race and other demographics. This means that the estimated coefficient should not be interpreted as the marginal effect of having more expensive housing, but as the effect on the average high housing value county in the hurricane region.

The average median family income in a county that experienced one hurricane between 1980 and 1996 is $40, 000 (2008 dollars), with a standard deviation of $9, 595. The bottom ten percent of counties has median incomes of $30, 991 and lower, while the top ten percent has median incomes of $51, 954 and higher. The variation in median housing values is similar, with a mean of $59, 297, a standard deviation of $17, 585, and tenth and ninetieth percentiles of $37, 894 and $82, 580, respectively. The distribution of median family income and housing values for the hurricane region as a whole is similar.

2.5 Empirical Results

2.5.1 Dynamic Effects of Hurricanes

In this section, I present the estimated effects of a hurricane. I graph the coefficients from equation (2.2) in Section 4. Because the results suggest that hurricanes have lasting effects and that there may be some differential pre-trends, following each figure is a table with the results of the trend break and mean shift tests described in equations (2.3) and (2.4). All monetary figures are in 2008 dollars. Standard errors are clustered by county. Each regression includes year, county, and year-by-coastal fixed effects, as well as indicators for hurricanes occurring outside of the estimation window of interest. The point estimates from the figures are shown in Appendix Tables A1-A4.

The disaggregated results and the trend break/mean shift estimates are complementary. The trend break and mean shift tests may pick up effects that are not detectable in a single year. However, they may miss non-monotonic dynamic effects. Thus, I view both as important in understanding post-hurricane economics.

Figure 2 shows the impact of a hurricane on the construction sector, measured in terms of employment, wages, and the number of firm locations. In addition, I present estimates of changes in per capita single family home construction. The y-axis shows the estimated coefficient and the 95% confidence interval. The x-axis represents the number of years since the hurricane; thus, negative numbers refer to leads of the hurricane variable. Because the coefficients are estimated from two-year bin variables, they are plotted at the midpoint of the two years (e.g., the point estimate for 1 and 2 years post-hurricane is plotted at 1.5 years). The coefficient for the two-year bin grouping years -1 and -2 (one and two years before the hurricane) is assumed to be 0.

Looking at the effects of the hurricane on the construction sector is useful for determining whether there are any effects of a hurricane that are observable a year or more after the event. Overall, these estimates clearly show that there are significant effects of a hurricane years after its occurrence. After remaining unchanged in the year of the hurricane, employment falls to $7-19\%$ below pre-hurricane levels (implying $70-170$ fewer construction workers).[21] The number of establishments is about 3.2% higher in the year of the hurricane and 3.7% higher the subsequent year (implying $2-3$ more construction establishments. They subsequently return to their pre-hurricane levels. Construction wages increase in years $1-4$, by $5.4-7.7\%$ ($\$1,400-2,000$), suggesting there may be a change in the composition of labor demand (e.g., more demand for specialized workers) or lower labor supply. The overall decline indicates a drop in construction demand three to eight years later: either less housing is being built or existing housing is being repaired less. This is possibly due to repairs being moved up temporally because of the hurricane. However, it is not

[21] I estimate this by computing $e^{(\beta_l+\mu)} - e^{(\mu)}$, where μ is the mean of the outcome and β_l is the estimated effect of the hurricane l years ago. This gives the approximate hurricane-driven change for logged variables.

clear whether the decline is temporary; nine to ten years later, the construction sector employment is still significantly lower than the year before the hurricane, but appears to be slowly increasing. Per capita single family housing starts are 5.7% lower in the year of the hurricane and 9.7% lower 3-4 years later (implying $0.2 - 0.4$ fewer housing units per 1,000 people), with no significant changes in other years. The hurricane lags are jointly significant at the 1% level for all the outcomes.

Table 6 shows the mean shift and trend break test results corresponding to Figure 2. There is a significant trend break in construction employment, establishments, and per worker wages, and the estimated coefficients follow the pattern seen in Figure 2. Specifically, the number of construction firm locations (establishments) declines by 1.3% each year. Construction employment is on average 9.0% lower in the ten years following the hurricane, and declines by 2.0% per year. Wages increase by an average of 7.0%, but then fall by an additional 1.1% each year. From the trend break specification, I estimate that construction employment is 19% lower five years after the hurricane and 29% lower at the end of my estimation sample, ten years after landfall.

One possible interpretation of the decline in the local construction sector is spatial; the construction activity may have simply shifted to nearby counties without any aggregate effect. The implications of spatial changes, while non-trivial for the local economy, are different than if there's a general downturn in the housing market. However, I also estimate that per capita housing starts fall by about 7.6% on average, which indicates a substantial decrease in construction demand. Thus, the downturn in the local construction sector is not solely driven by spatial shifts in construction activity.

Figure 3 shows the estimated effect of a hurricane on population and demographics. Population does not change significantly in any given year and the effects of a hurricane zero to ten years after are not jointly significant. The fraction of black residents is significantly lower in the years after the hurricane, but pre-trends suggest that further testing is necessary. The fraction of residents who are 65 and older falls steadily following the hurricane, while the fraction of those under 20 years of age steadily grows.

Trend break and mean shift tests in Table 8 indicate that there are no significant changes in the mean or trend of population or the fraction of residents who are black. There is indeed evidence of a change in the age structure of the county. In particular, the fraction of population under 20 is 0.0036 higher 10 years after the hurricane, a 1% increase. The fraction over 65 is 0.0058 lower, a 4.7% decrease relative to the mean. These changes are significant at the 10% and 5% levels, respectively.

Figure 4 shows the effect of a hurricane on the employment rate, earnings, and transfers. There is no change in the employment rate. Per capita net earnings by residents and the employment rate show a significant pre-hurricane trend, as evidenced by the significance of the joint test of hurricane leads, but no change following the hurricane. Overall per capita transfers from the government to individuals increase by 1.2% in the year of the hurricane and are $1.8 - 2.5\%$ larger in subsequent years. Per capita transfers to individuals from businesses immediately increase by 11.6% following a hurricane.

In Table 8, I show the results of the mean shift and trend break tests for the outcomes shown in Figure 4. The mean shift test indicates a 2% average increase in per capita government to individual transfers, equivalent to about \$67 per person per year. Per capita business to individual transfers in the eleven years following the hurricane are estimated to be 3% higher than the pre-hurricane transfers, or about \$2.4 per year. There are no significant changes in the trends of any of these variables. Assuming a 3% discount rate, the present discounted value (PDV) of all government transfers is about \$640 per capita, and the PDV of transfers from businesses is \$23 per capita. Thus, post-hurricane transfers from general social programs are larger than transfers from disaster-specific programs and much larger than insurance payments.

Figure 6 looks at specific types of government transfers: namely, family assistance, public medical benefits (i.e., Medicaid and Medicare), SSI payments, and unemployment insurance. Per capita unemployment insurance payments increase immediately by 14.2% and are $18 - 30\%$ higher in years one through ten after a hurricane. Family assistance payments are 2.9% lower in the year of the hurricane, but subsequently rise to $7 - 10\%$ above their pre-hurricane average. SSI is estimated to fall, but the variable clearly exhibits a significant pre-trend. All of the increases appear to be temporary: per capita UI is the only variable that's significantly higher ten years later and appears to be coming back down to pre-hurricane levels.

Table 9 shows the corresponding mean shift and trend break tests. On average, per capita unemployment benefits increase by 22%, equivalent to about \$19 per person per year. Assuming a 3% discount rate, the present discounted value (PDV) of the unemployment payments is about \$180 per capita. Per capita medical benefits increase by about 4.6% on average, but the dollar equivalent of the increase is very small. The overall increase in per capita family assistance payments is only marginally significant and the dollar equivalent of the increase is likewise small.

2.5.2 Heterogeneity of Effects by Wealth

In this section, I focus on the heterogeneity in the post-hurricane employment rate, per capita earnings, population, and transfer payments. In Table 10, I show the results of the mean shift and trend break tests by the quartile of median housing value. There are no significant differences in population or employment rate changes. Per capita earnings increase on average in low housing value counties and subsequently decrease (relative to counties in the two middle quartiles). The reverse pattern holds in high housing values counties, so that ten years after the hurricane, per capita earnings changes are not estimated to be significantly different between high and low housing value counties.

Changes in per capita overall transfers from the government and per capita unemployment insurance are also qualitatively different for bottom and top quartiles of housing value (relative to counties in the two middle quartiles). Per capita transfers from the government are substantially higher in low-value counties while in high-value counties they are substantially lower on average.

Per capita unemployment insurance increases by 0.21 log points more in low-value counties than in medium-value counties and shows an upward trend while remaining unchanged in high-value counties. These results highlight interesting qualitative differences between counties of different housing values and suggest that government transfers may play a larger role in low housing value counties in the aftermath of a hurricane. Appendix Table A5 shows the corresponding estimates for high-income and low-income counties. The estimates generally follow the pattern in Table 10.

Overall, hurricanes appear to produce differences (some lasting and some temporary) in areas that differ in incomes and housing values, but the mechanism for how and why this occurs cannot be determined with the current data. The differential increase in per capita transfers reinforces the idea that these may also play an important role in absorbing the impact of the shock. Because heterogeneity in the post-hurricane economic dynamics should be an important factor for policy design, potential explanations such as differential credit constraints and moving costs deserve further detailed study.

2.5.3 Robustness

In this section, I report the results of various checks to verify that the results in the previous section are robust and to examine the variation in the magnitude of estimated effects. Overall, the qualitative result of higher transfer payments with no corresponding change in other variables is robust across different samples, and the magnitude of the estimated increase is relatively stable.

Joint tests of the lead hurricane indicators in Appendix Tables A1-A4 suggest that there are pre-trends in some of the hurricane variables. One explanation for the significance of these lead coefficients is that hurricane-prone areas have a different time trend. Combined with the fact that outcomes are autocorrelated, this implies that leads of the hurricane variable are likely to be significant spuriously, due to the omitted variable bias. Unfortunately, the paucity of hurricanes does not allow me to estimate a county-specific trend and include year and county fixed effects at the same time. In Appendix C, I use a Monte Carlo simulation to demonstrate that the pre-trends are not likely to affect the qualitative estimates in this case. In addition, the overall time trend is estimated to be insignificant in all trend break tests except for per capita unemployment and SSI payments and the fraction of population that is black (which are all estimated to be decreasing in the hurricane counties).

In the first robustness test, I restrict the sample to only counties affected once by a hurricane between 1980 and 1996 (in other words, only the treated group). Although the sample is smaller, the basic results still hold. The estimated amount of extra government transfers is somewhat larger, but comparable to the original estimates. The only substantial difference is that, in addition to the mean increase, unemployment payments show an upward trend of about 2.7% per year. Per capita UI payments are estimated to be 39% higher five years after the hurricane and 53% higher eleven years after the hurricane.

Another robustness sample includes a control group that's constructed using only historic hurricane data and propensity score matching (recall that the main control group was also matched by 1970 covariates). The simple risk-based matching also yields results that are very similar to the main sample and in some cases produces more precise estimates.

One other concern with the basic specification and sample is that there may be spatial effects. In other words, a neighbor of a centrally affected county may also be affected. This could be either due to unmeasured hurricane destruction, as discussed in Section 3, or because of spatial economic spillovers. The spillovers can be positive or negative, so the sign of the bias created by spatial effects is ambiguous. To see if spatial spillovers are a concern, I omit unaffected neighbors of counties that experience hurricanes for eleven years after the hurricane. There is no significant fall in construction employment but a 1.2% decline in the number of establishments. As before, the average construction wage increases by an average of 8.5% following the hurricane but shows no downward trend in this sample. This suggests that some of the hurricane county's construction activity moves to the neighboring counties, implying that hurricanes may permanently affect the business patterns in centrally affected and neighboring counties.

One other potential confounder is that those likely to receive government transfers may be moving into the counties affected by hurricanes from nearby counties so that there is no aggregate impact on transfers, only a compositional change. One way to test for this is to look at changes in transfers on the state level. Unfortunately, the affected population represents 11% of the state population on average. Thus, the power to detect an aggregate affect is low. Instead, I look at the changes in transfers in counties whose center is within 50 miles from the center of the affected county (including the affected county itself). This distance should be large enough to capture potential compositional changes, but not so large that the power to detect a change in transfers is reduced.

The results are generally very similar. In the ten years following a hurricane, employment in the 50-mile radius is unaffected. Per capita transfers from the government increase by 2.2% on average and show an increasing trend of 0.4% per year following the hurricane. Per capita unemployment insurance payments increase by 26.6% on average and show an increasing trend of 4% per year. The only substantive difference between this and the other samples is that per capita earnings are estimated to decline by an extra 0.57% per year following the hurricane. Finally, including all the counties in the hurricane region as controls also leads to similar results.

Adding state-by-year fixed effects to the basic specification generally makes the results insignificant. This is not surprising given the autocorrelation of the outcomes and relatively few counties in the affected sample. In particular, transfers from businesses to individuals are no longer estimated to be significantly higher. As these represent insurance payment, which should increase following a hurricane, this suggests that including state-by-year fixed effects is overly conservative.

2.6 Interpretation and Discussion

The construction estimates show that hurricanes have non-monotonic medium-run effects on that sector, with an initial increase in wages followed by a gradual return to pre-hurricane levels three to eight years after the hurricane. At the same time, there is a fall in single family housing starts. However, the effects on general economic variables, such as population, employment, and wages, are insignificant. Although the US has a developed disaster response system, my estimates suggest that traditional social safety nets, such as unemployment insurance, also play an important role in post-disaster economics. The largest empirical effect of a hurricane is on non-disaster transfer programs, the transfers from which increase substantially after a hurricane. For a county with the average population of $78,000$, the estimated increase of \$640 per capita in non-disaster government transfers translates to a total of \$50 million. This is much larger than the estimated disaster aid, which contributes \$356 per capita, and could have non-trivial fiscal implications in the future if climate change intensifies the strength of hurricanes. Together, disaster and non-disaster transfers represent a large fraction of the direct damages caused by the hurricane. These estimates also imply that the fiscal impact of natural disasters is more than twice as large if non-disaster transfers are also considered. Although in the simplest public finance framework transfers are welfare-neutral, in practice the deadweight loss of taxation is estimated to be $12 - 30\%$ of revenue (Ballard et al., 1985; Feldstein, 1999). Assuming a 15% deadweight loss implies a real cost of \$53 per capita per hurricane for disaster transfers (\$4.1 million for a county with a population of $78,000$) and \$96 (\$7.5 million) per capita per hurricane for non-disaster transfers. Taking the upper estimate of 30% doubles these estimates. Moreover, the *marginal* deadweight loss of taxation, which is the relevant cost if we're considering mitigating the effect of hurricanes, is thought to be much larger. Feldstein (1999) estimates it to be \$1 – \$2 per dollar of revenue. Thus, this additional cost of hurricanes to society is not trivial.

Of course, the estimate that non-disaster transfers are $1.5 - 2$ times larger than disaster transfers does not imply that they are $1.5 - 2$ times as important. The designs of the disaster and non-disaster government programs suggest that they may be complementary. Social insurance programs may fill an important gap left by current disaster policy and private insurance markets. Disaster transfers target individuals immediately impacted by the disaster and provide funds to restore public infrastructure. Disaster aid to individuals makes up only about 39% of total disaster aid; the rest is allocated to activities such as debris cleanup and restoration of public buildings and roads (FEMA, personal communication). Private insurance targets individuals who sustain disaster losses in the form of property damage. Non-disaster social insurance programs, such as unemployment insurance, are able to target individuals who are affected indirectly.

Although the US has a disaster-related unemployment insurance program (it is included in the figure for disaster-related transfers), it provides benefits only to those who can show that they lost their jobs directly as a result of the disaster. Individuals who lose their jobs as a result of an economic downturn months to years later would be unable to claim these benefits. If there are lasting economic effects (as seems to be the case with US hurricanes), people may be affected months to years following the disaster. In that case, disaster aid and property insurance are not helpful, but standard social safety net programs will be automatically triggered. The presence of these programs can thus serve as insurance against delayed effects of natural disasters. As discussed in the conceptual framework, non-disaster transfers may buffer the economic shock of a hurricane and also explain why there are no large changes in population, employment or wages. According to the World Labour Report 2000, seventy-five percent of the world's unemployed are not receiving any benefit payments (International Labour Organization, 2000). In addition to making individuals vulnerable to economic shocks, my analysis suggests that a lack of social safety nets also has implications for the economic recovery of an area following a natural disaster.

One possible explanation for the increase in unemployment payments and overall government transfers is changes in the demographic composition of an area. This change in the age composition is inconsistent with the changes in non-disaster transfers. Total government transfers include social security and disability payments. There is no a priori reason to think that a larger number of young people and a decline in the number of elderly would increase the total transfers. Young people are more likely to be unemployed than the elderly, but most of the people in the "under 20 years old" category are unlikely to be receiving unemployment insurance payments. When I separate the category "20 to 64 years old" into ten-year age categories, I find that there is no change in the fraction or log of population that is between 20 and 29, 30 and 39, or 40 and 49 years old. There is a slight increase in the fraction of population that is between 50 and 64, but it is not large enough to explain the increases in transfers. This age group makes up about 14% of the total population. To explain a non-trivial part of the increase in government transfers, each person in this age group would have to be receiving an implausibly large amount of them.

This demographic change does raise concerns about other unobserved changes in the population. However, to the extent that the changes in unobservable characteristics are correlated with the changes in observable ones, this is not likely to be an issue. Disaggregated estimates indicate that the compositional change is gradual, while the increases in the unemployment insurance and overall transfers are immediate and non-monotonic. If the non-disaster transfers were driven by demographic changes, the change in the age profile should correspond to the change in transfers. As the two differ, it's likely that the demographic change is another effect of the hurricane that is unrelated to the change in transfers.

One possible explanation for the demographic change is a change in the composition of job opportunities that makes the county a relatively less attractive place to retire and a relatively more attractive place to raise children. Another is different risk preferences across different ages combined with updated beliefs following the hurricane. If the elderly are more risk averse, they will be more reluctant to live in a hurricane-prone area. Population as a whole may not decrease if housing prices adjust to compensate for the increase in the perceived risk of living in the county.

Whether the presence of unemployment insurance for those living in disaster areas is welfare-improving on a national level is not straightforward. On one hand, the presence of insurance against economic losses not covered by homeowner's and flood insurance is a benefit when individuals are risk averse or credit constrained. Theoretically, they may allow credit constrained individuals to avoid moving costs during rebuilding. However, disaster risk is not currently accounted for in unemployment insurance premiums. This subsidizes business activity in disaster-prone areas, which decreases social welfare. Thus, disaster and non-disaster transfers may be creating a moral hazard problem. In addition, there are many other distortions in insurance and aid policy that discourage insurance and encourage people to live in disaster-prone areas. This makes even a theoretical welfare analysis of unemployment insurance difficult.

2.7 Conclusion

If current demographic and economic trends continue, damages from natural disasters will increase, both in absolute terms and as a percentage of GDP. In addition, climate change is projected to increase the frequency and intensity of extreme weather events. Projections for future increases in disaster damages due to climate change are highly uncertain but thought to be large. A country's infrastructure and institutions have been identified as important determinants of the impact of extreme weather events, both theoretically and empirically. Thus, informed policy has the potential to reduce the damages caused by extreme weather and mitigate its economic impacts. However, the economic impacts of extreme weather are understudied. Most of the literature to date has focused on studying damages or very short-run impacts on isolated variables; a comprehensive picture of post-disaster economic dynamics is lacking.

I estimate the medium-run economic effects of hurricanes on US counties, focusing on population, employment, wages, and transfers to individuals. Population, the employment rate and wages are largely unaffected in the ten years following the hurricane, while construction employment and new housing starts decline substantially.

I find that hurricanes have large and persistent effects on non-disaster transfer payments. Real transfers from traditional safety net programs over the eleven years following the hurricane (including the year of the hurricane) are estimated to total \$640 per capita, much larger than the disaster-related transfers of \$356 per capita. Insurance payments increase temporarily in the year of the hurricane but add only an estimated \$23 per capita in present discounted value.

Most of the transfers from traditional safety net programs are estimated to occur later than government disaster transfers and insurance payments typically occur, suggesting that traditional safety net programs are filling in a gap in public and private disaster insurance. Private insurance in this case is best suited to targeting those who lose their homes, but traditional social insurance may target those who are affected by the dynamic economic effects of the disaster.

Transfer programs designed for general economic downturns, such as unemployment insurance and food stamps, can act as buffers against adverse economic impacts following destruction. First, this has implications for the actual costs of a disaster: in addition to money spent on disaster relief, extreme weather has fiscal effects on other government transfer programs. Second, ignoring traditional transfer programs attributes too much of a developed economy's resilience to its wealth or disaster response policies and not enough to general social policies.

My results also show gradual demographic changes in the affected county. The fraction of the population under 20 increases while the fraction of the population 65 and older falls. These changes could be caused by differential economic opportunities for different demographic groups or by updated beliefs about risk and risk preferences that vary across groups. Finally, a county's wealth and housing stock value also seems to matter for earnings and transfer trajectories. Whether this is because of different decisions or constraints of individuals or because of differential hurricane impacts is an area for future research.

My findings have several suggestive policy implications. First, policymakers should consider the potential role of non-disaster programs in recovery. Second, they may want to incorporate disaster-related unemployment risk into the design of social safety net programs to avoid moral hazard issues. Third, as the fiscal costs of disasters are larger than previously thought, implementing mitigation programs is correspondingly more beneficial. Admittedly, I cannot estimate what the effects of a US hurricane would be without social insurance programs using the current data. Given that much of the world's population does not have access to social or disaster insurance and is at an increasing risk of natural disasters, the causal effect of social insurance on disaster impacts and whether it creates moral hazard are two other areas that deserve further study. Estimating the impact of federal disaster aid is another important area for future research.

2.8 References

Ballard, C., Shoven, J. and J. Whalley. 1985. "The Welfare Cost of Distortions in the United States Tax System: A General Equilibrium Approach", NBER Working Paper 1043.

Belasen, A. and S. Polachek. 2008. "How Hurricanes Affect Employment and Wages in Local Labor Markets", IZA Discussions Paper No. 3407.

Board on Natural Disasters. 1999. "Mitigation emerges as major strategy for reducing losses caused by natural disasters", *Science* 284: 1943-1947.

Brown, S., Mason, S. and R. Tiller. 2006. "The effect of Hurricane Katrina on employment and unemployment", *Monthly Labor Review*, August: 52-69.

Meehl, G.A., T.F. Stocker, W.D. Collins, P. Friedlingstein, A.T. Gaye, J.M. Gregory, A. Kitoh, R. Knutti, J.M. Murphy, A. Noda, S.C.B. Raper, I.G. Watterson, A.J. Weaver and Z.-C. Zhao. 2007. "Global Climate Projections." In: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L.

Miller (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Dell, M., Jones, B., and B. Olken. 2009. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century", Working Paper.

Deschenes, O., and M. Greenstone. 2007a. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US", NBER Working Paper 13178.

Deschenes, O., and M. Greenstone. 2007b. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather", *American Economic Review*, 97 (1): 354-385.

Feldstein, M. 1999. "Tax Avoidance And The Deadweight Loss Of The Income Tax", *The Review of Economics and Statistics*, 81(4): 674-680.

Freeman, P., Keen, M., and M. Mani. 2003. "Dealing with Increased Risk of Natural Disasters: Challenges and Options", IMF Working Paper 03/197.

Hazards & Vulnerability Research Institute (2009). The Spatial Hazard Events and Losses Database for the United States, Version 7.0 [Online Database]. Columbia, SC: University of South Carolina. Available from http://www.sheldus.org.

Hornbeck, R. 2009. "The Enduring Impact of the American Dust Bowl: Short and Long-run Adjustments to Environmental Catastrophe", NBER WP 15605.

International Labour Office. 2000. *World Labour Report 2000: Income Security and Social Protection in a Changing World*, Geneva International Labour Office, 2000.

Jones, B., and B. Olken. 2010. "Climate Shocks and Exports", *American Economic Review Papers and Proceedings*, forthcoming.

Kahn, M. 2005. "The Death Toll from Natural Disasters: The Role of Income, Geography, and Institutions", *Review of Economics and Statistics*, 87(2): 271-284.

Leiter, A., Oberhofer, H., and P. Raschky. 2008. "Productive disasters? – Evidence from European firm level data". Wharton WP 2008-04-14.

Masozera, M., Bailey, M. and C. Kerchner. 2007. "Distribution of impacts of natural disasters across income groups: A case study of New Orleans", *Ecological Economics*, 63: 299-306.

Murphy, A. and E. Strobl. 2009. "The Impact of Hurricanes on Housing Prices: Evidence from US Coastal Cities", Working Paper.

Nordhaus, W. 2006. "The Economics of Hurricanes in the United States", NBER WP 12813.

Notowidigdo, M. 2010. "The Incidence of Local Labor Demand Shocks", Working Paper.

Okazaki, T. Ito, K., and A. Imaizumi. 2009. "Impact of Natural Disasters on Industrial Agglomeration: The Case of the 1923 Great Kanto Earthquake", WP CIRJE-F-602.

Schneider, S.H., S. Semenov, A. Patwardhan, I. Burton, C.H.D. Magadza, M. Oppenheimer, A.B. Pittock, A. Rahman, J.B. Smith, A. Suarez and F. Yamin. 2007. Assessing key vulnerabilities and the risk from climate change." Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, M.L. Parry, O.F. Canziani, J.P. Palutikof, P.J. van der Linden and C.E. Hanson, Eds., Cambridge University Press, Cambridge, UK, 779-810.

Skidmore, M. and H. Toya. 2002. "Do Natural Disasters Promote Long-Run Growth?", *Economic Inquiry*, 40 (4): 664-687.

Skidmore, M. and H. Toya. 2005. "Economic Development and the Impacts of Natural Disasters", Working Paper.

Strobl, E. 2009. "The Economic Growth Impact of Hurricanes: Evidence from US Coastal Counties", *The Review of Economics and Statistics*, forthcoming.

Strobl, E. and F. Walsh. 2008. "The Re-Building Effect of Hurricanes: Evidence from Employment in the US Construction Industry", IZA Discussion Paper No. 3544.

Sylves, R. and D. Racca. 2010. "All About Presidential Disaster Declarations," sponsored by Public Entity Risk Institute. Newark, DE: University of Delaware, www.peripresdecusa.org, last accessed March 2009.

Weatherford, C.L. and W.M. Gray. 1988. "Typhoon Structure as Revealed by Aircraft Reconnaissance. Part II: Structural Variability", *Monthly Weather Review*, 116 (May): 1044-1056.

Yang, D. 2008. "Coping with Disaster: The Impact of Hurricanes on International Financial Flows, 1970-2002," *The B.E. Journal of Economic Analysis & Policy*, 8(1), Article 13.

Zeckhauser, R. 1996. "The Economics of Catastrophes", *Journal of Risk and Uncertainty*, 12: 113-140.

Appendix A. Distribution of Eye Diameters

Because eye diameters are related to the size of the affected area but the relevant data are not systematically collected, I have to consider how not observing the diameter might impact estimates of a hurricane's effect. The most comprehensive information on the characteristics of North Atlantic storms come from air reconnaissance data analyzed by Weatherford and Gray (1988). About 15% of cyclones in the data have eyes less than 18.6 miles in diameter, 25% have eyes 18.6 – 37.3 miles in diameter (average is 27.3 miles), and 8% had larger eyes (the largest was 149 miles). Most (53%) had no discernible eye and very few had eyes larger than 49.7 miles in diameter. Thus, nearly 70% of all hurricane eyes appear to not be very large.

Appendix B. Cross- and Auto-correlation in Outcome Variables

In Appendix Tables A6-A8, I show some of the cross-correlation and autocorrelation coefficients of the variables used in the analysis (in logs), taking out year and county fixed effects.

Table A6 shows that total unemployment payments are significantly and positively correlated with total wages and employment, although the magnitudes are not large (8.6 and 6.3%, respectively), while unemployment payments per capita are negatively correlated with these variables (−23.1% and −32.9%, respectively). Earnings per capita are negatively correlated with income maintenance payments (the magnitude is −6.4%) while total earnings are positively correlated (21%). State and local tax receipts are highly positively correlated with both total earnings (82.3%) and earnings per capita (58.8%).

Table A7 shows the correlation between the construction sector variables and other economic outcomes. The number of construction firms, employees, wages, and payroll are heavily correlated with overall earnings and earnings per capita and the magnitudes are large. A 10% increase in earnings per capita is associated with a 7.3% increase in the number of construction firms, a 10.6% increase in the number of construction employees, and 12.6% and 2.3% increases in total construction payroll and wages, respectively.

Table A8 focuses on the autocorrelation between the outcome variables. The autocorrelation coefficients are all significant at the 1% level and range from 0.095 for construction wages per worker to 0.986 for overall business receipts. Total population and government transfers to individuals and non-profits are also heavily correlated, while earnings per capita, business transfers to individuals, and payroll wages per worker are among the least correlated (although the correlation coefficients are still between 0.275 and 0.693). This also leads to an R-squared that is nearly equal to 1.

These cross- and auto-correlations present some challenges for the estimation. In particular, it may be more difficult to estimate the effect of a hurricane in a given time period precisely. For this reason, I rely on both joint and individual significance testing when interpreting the results. Because I can identify the precise onset of a hurricane, I can still estimate its duration and net impact fairly accurately.

Appendix C. Risk-related Trends, Autocorrelation and Lead Significance: a Monte Carlo Analysis

The F-tests of hurricane lead indicators indicate that they are significant for many economic outcomes. In this section, I use a Monte Carlo simulation to show how this can arise when, in addition to autocorrelation, there is an unobserved time trend that is correlated with hurricane risk and discuss how this affects my estimates. Both heterogeneous time trends and autocorrelation appear to be present in my sample.

I generate a sample of 1000 "counties" that are observed for 30 time periods. I randomly assign 5% of the observations to experience an "event". In addition, each county is assigned an unobserved risk variable which is correlated with the occurrence of the event. The outcome of interest is determined as follows:

$$Outcome_{ct} = \beta Event_{ct} + \gamma Outcome_{c,t-1} + \theta y * Risk_c + \varepsilon_{ct}$$

$Event_{ct}$ is the event indicator for county c in year t. $Risk_c$ is a uniform variable between zero and one, multiplied by the mean of the event indicator for the county (this implies that $Risk_c$ will be correlated with $E_c\left[Event\right]$). I assume that $\beta = 10$, $\gamma = 0.9$, and $\theta = 0.001$. ε_{cy} is standard normal and is identically and independently distributed across counties and time periods.

The risk variable captures the possibility that time trends are related to a county's propensity to be affected by hurricanes. This could be because the county is investing increasing amounts of mitigation over time, insurance is becoming more widely available and adopted or because people are slowly leaving the area as they realize the hazard they face. Alternatively, as the economy becomes wealthier, people may disproportionately prefer to live in hurricane-prone places if there is risk aversion or if wealthier people are more able to weather the shock of a hurricane. All of these factors could produce unobserved heterogeneity in the time trend.

Following the generation of the Monte Carlo sample, I estimate two regression specifications similar to those in the paper. One of them includes leads and the other considers lags only. These are specified as follows:

$$Outcome_{ct} = \sum_{\tau=-5}^{5} \beta_\tau Event_{c,t-\tau} + \beta_{ct}^{6} + \beta_{ct}^{-6} + \alpha_c + \varepsilon_{ct}$$

and

$$Outcome_{ct} = \sum_{\tau=0}^{5} \beta_\tau Event_{c,t-\tau} + \beta_{ct}^{6} + \alpha_c + \varepsilon_{ct}$$

$\{\beta_{ct}^{-6}, \beta_{ct}^{6}\}$ are indicator variables for the "event" outside of the estimation window of interest. Thus, this estimation is analogous to the estimation of the effect of hurricanes on the US, with $\{\beta_\tau\}_{\tau=0}^{\tau=5}$ being the estimated effect of the event 0-5 years after relative to the pre-event outcome. Note that the estimated effect of the "event" in years other than 0 is entirely due to the autocorrelation in the outcome variable and is also possibly affected by the unobserved time trend heterogeneity.

Appendix Table A9 shows the theoretical and estimated effects of the event 0 - 5 time periods after and 2 - 5 periods before its occurrence. Although the theoretical effects of lead variables are zero when no time trends are present, separate and joint tests indicate that all leads negatively affect the outcome. Including the lagged outcome variable does not change the significance of the leads. The inclusion of leads also appears to bias the estimated effect of an event down.

Despite the presence of risk-related time trends, the estimated coefficients of lags appear to be fairly close to the theoretical effect, although they are biased downwards. A priori, it is not clear whether the inclusion of leads will decrease or increase the bias; this depends on whether the time trend is positive or negative.

I do not include the lagged outcome variable in my analysis of hurricanes because the large number of fixed effects combined with high autocorrelation makes the estimation ill-behaved. County and time fixed effects are likely to be more important to include. Moreover, I am interested in the overall effects of hurricanes, including through autocorrelation. The lagged hurricane indicators implicitly capture the autocorrelation *and* any non-standard dynamic that may occur following a hurricane (for example, the non-monotonicity of the construction sector response).

Tables and Figures

Table 1: Predicted changes in the economy following a negative capital shock

	(1)	(2)	(3)	(4)
	Mobile capital		Immobile capital	
	No moving costs	Moving costs	No moving costs - m	Moving costs - c
No transfers, inelastic labor supply - n	No change		dwnm, drnm < 0, dLnm=dNnm<0, dV(s)=0	0>dwnm>dwnc, 0>drnc>drnm, dLnc=dNnc=dNnm<0, dV(s)≤0
Transfers, elastic labor supply - t			dwtm=dwnm<0, drtm=dwnm < 0, dLtm=dNtm=dNnm<0, dV(s)=0, dT(s)=0	dwtc, drtc < 0, dLtc < dNtc < dNnm<0, dV(s)≤0, dT(s)>0

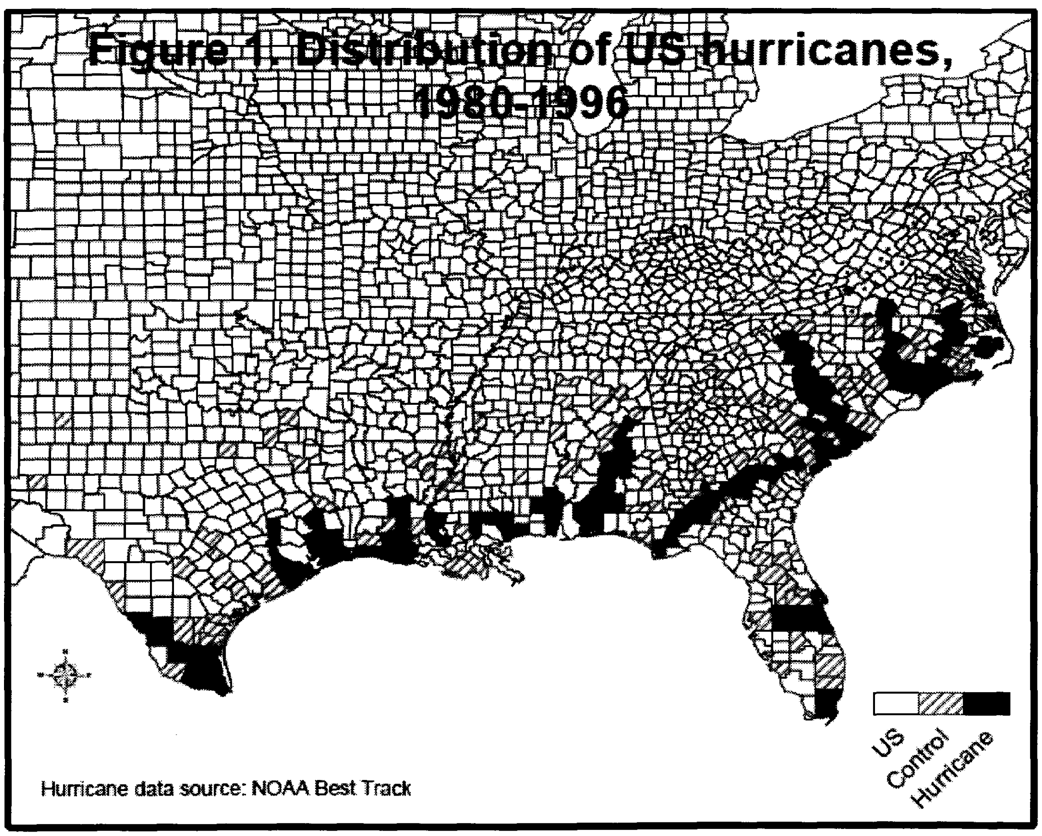

Figure 1. Distribution of US hurricanes, 1980-1996

Table 2: Damages caused by major US hurricanes, 1980-1996

	(1) Mean	(2) Standard deviation	(3) Maximum	(4) Obs.
	Panel A: centrally affected counties			
Total building value ($1000's)	10,704,091	32,552,770	268,081,632	99
Building damage ($1000's)	405,555	2,101,017	20,300,000	97
Loss ratio (percent)	1.46	3.85	23.62	97
Total losses ($1000's)	570,558	3,662,500	35,400,000	97
Total per capita loss ($)	1,278	3,336	16,238	97
Displaced households	1,546	10,702	104,559	99
People requiring shelter	449	3,078	29,945	99
	Panel B: peripherally affected counties			
Total building value ($1000's)	7,464,867	26,808,163	464,355,684	400
Building damage ($1000's)	8,635	40,294	388,928	390
Loss ratio (percent)	0.15	0.51	5.20	390
Total losses ($1000's)	11,462	57,071	632,972	390
Total per capita loss ($)	113	430	4,816	385
Displaced households	12	85	1,193	403
People requiring shelter	3	23	331	403

Source: HAZUS-MH simulation software published by FEMA. All monetary figures are in 2008 dollars.

Table 3: Determinants of property damages in hurricane region[1]

	(1) Log damages	(2) Per capita damages	(3) Flood insurance payments (log)	(4) Log damages	(5) Per capita damages	(6) Flood insurance payments (log)
Major hurricane	4.236	678.279	3.090			
	(0.426)***	(309.795)*	(0.410)***			
Minor hurricane	2.417	65.184	1.515			
	(0.250)***	(46.066)	(0.247)***			
Category = 1				2.151	72.735	1.131
				(0.297)***	(57.567)	(0.263)***
Category = 2				3.253	39.028	2.769
				(0.411)***	(19.173)*	(0.492)***
Category = 3				4.049	710.213	3.100
				(0.311)***	(399.246)	(0.471)***
Category = 4 or 5				4.642	607.060	3.019
				(0.915)***	(316.198)*	(0.850)***
Tornado	2.061	12.507	-0.008	2.061	12.441	-0.011
	(0.197)***	(6.847)	(0.070)	(0.196)***	(7.174)	-0.070
Flood	0.862	0.380	0.762	0.864	0.299	0.758
	(0.102)***	(5.690)	(0.065)***	(0.102)***	(5.757)	(0.065)***
Severe storm	0.958	8.952	-0.205	0.956	9.075	-0.201
	(0.180)***	(3.915)*	(0.078)***	(0.181)***	(4.213)*	(0.079)**
Depvar mean (median)	9.31 (9.66)	11.00 (0.09)	10.87 (10.80)	9.31 (9.66)	11.00 (0.09)	10.87 (10.80)
Observations	18,592	24,331	7,029	18,592	24,331	7,029
R-squared	0.45	0.08	0.42	0.45	0.08	0.42

Standard errors (clustered by county) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Damages and flood claims are in current dollars. Includes county and year fixed effects. Property damage data is from SHELDUS. Flood insurance payments data is from the Consolidated Federal Funds Report (CFFR). Time period is 1980-1996 for damages, 1983-1996 for flood claims.

[1]Alabama, Florida, Georgia, Louisiana, Mississippi, North Carolina, South Carolina, Texas, and Virginia.

Table 4: Descriptive statistics for hurricane aid, 1980 - 1996

	(1) Uniform split[1]	(2) Proportional split[2]	(3) Per capita - uniform split[1]	(4) Per capita - proportional split[2]
	Panel A: centrally affected counties only			
Centrally affected, all hurricanes (N = 89)	58,700,000	58,700,000	1,137	356
	(187,000,000)	(260,000,000)	(3,193)	(307)
Centrally affected, major hurricanes (N = 27)	128,000,000	133,000,000	2,018	412
	(332,000,000)	(467,000,000)	(5,623)	(343)
	Panel B: all counties listed in declaration			
All observations (N = 568)	8,982,356	8,417,279	160	52
	(48,400,000)	(65,200,000)	(631)	(91)
Centrally affected, all hurricanes (N = 89)	24,600,000	30,100,000	460	131
	(94,100,000)	(152,000,000)	(1,594)	(140)
Centrally affected, major hurricanes (N = 27)	59,200,000	73,400,000	954	187
	(167,000,000)	(273,000,000)	(2,824)	(184)

[1]Assumes aid money is split evenly among all counties in sample

[2]Assumes aid money is split in proportion to the population of counties in sample

Source: NOAA Best Tracks data, PERI disaster declarations. Standard errors in parentheses. All amounts are in 2008 dollars.

Table 5: Comparison of hurricane region[1] by 1980-1996 hurricane experience.

	(1)	(2)	(3)	(4)	(5)
	One hurricane	Difference from no hurricanes	p-value of difference	Difference from matching	p-value of difference
			Panel A: 1970 characteristics		
Coastal indicator	0.49	0.08	0.205	0.02	0.826
Land area, sq. mi	666.33	27.76	0.578	-54.02	0.363
Population (log)	10.32	0.47	0.000	-0.34	0.022
Population density (person/sq. mile)	87.29	-99.47	0.006	-24.52	0.181
Employment rate (fraction)	0.57	0.02	0.184	0.01	0.524
Net earnings per capita (log)	9.37	0.00	0.952	-0.02	0.508
Per capita transfers from government (log)	7.30	-0.09	0.002	-0.03	0.372
Per capita transfers from businesses (log)	4.05	0.00	0.918	-0.01	0.471
Per capita family assistance (log)	-2.81	0.05	0.590	0.04	0.689
Per capita public medical benefits (log)	-1.42	-0.05	0.231	0.00	0.955
Per capita UI payments (log)	3.81	0.23	0.001	-0.02	0.784
Per capita new single family housing (log)	-5.07	0.04	0.391	0.05	0.362
			Panel B: 1970-1979 trend		
Population (log)	0.0120	-0.0025	0.184	-0.0043	0.066
Population density (person/sq. mile)	0.3774	-0.9444	0.288	-1.5121	0.022
Employment rate (fraction)	-0.0002	0.0003	0.710	-0.0003	0.791
Net earnings per capita (log)	0.0190	0.0004	0.861	-0.0003	0.893
Per capita transfers from government (log)	0.0631	0.0043	0.042	0.0003	0.910
Per capita transfers from businesses (log)	0.0151	-0.0001	0.718	-0.0002	0.526
Per capita family assistance (log)	0.0386	0.0177	0.047	0.0161	0.116
Per capita public medical benefits (log)	0.0889	-0.0013	0.673	-0.0011	0.755
Per capita unemployment insurance payments (log)	0.0864	-0.0065	0.518	-0.0047	0.643
Per capita new single family housing (log)	0.0075	-0.0100	0.177	-0.0099	0.296
Number of counties[2]	119	811		195	

[1]Alabama, Georgia, Florida, Louisiana, Mississippi, North Carolina, South Carolina, Texas, Virginia

[2]Number may be smaller for some variables because of missing values.

Source: 1970 REIS, 1970 CBP and 1970 Census. Standard errors in parentheses. Monetary values are in 2008 dollars.

Figure 2. The effect of a hurricane on the construction sector

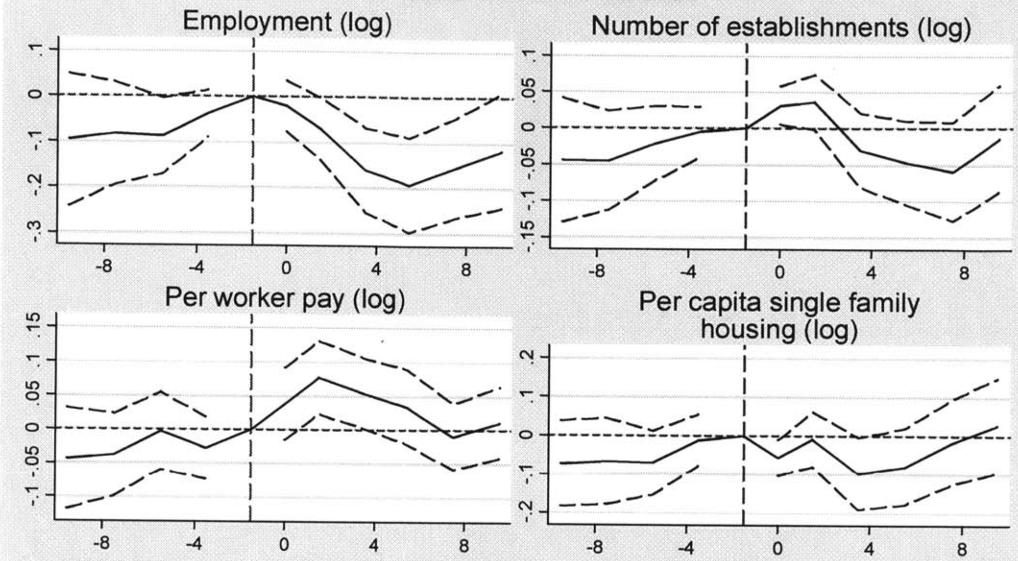

Sample: counties in Alabama, South Carolina, North Carolina, Florida, Georgia, Mississippi, Louisiana, Texas, and Virginia. 95% confidence intervals shown. Standard errors clustered by county. Controls include county, year, and coastal by year fixed effects.

Table 6: Mean shift and trend break tests for Figure 2

	Construction employment (log)		Construction establishments (log)		Construction per worker wage (log)		Per capita single family housing construction (log)	
Post hurricane	-0.0698	-0.0899	0.0188	0.0055	0.0804	0.0697	-0.0744	-0.0763
	(0.0375)*	(0.0432)**	(0.0232)	(0.0230)	(0.0263)***	(0.0276)**	(0.0392)*	(0.0397)*
Post hurricane time trend		-0.0204		-0.0133		-0.0108		-0.0019
		(0.0119)*		(0.0065)**		(0.0052)**		(0.0100)
Overall time trend	0.0002	0.0118	-0.0012	0.0064	-0.0023	0.0038	0.0078	0.0088
	(0.0044)	(0.0093)	(0.0033)	(0.0054)	(0.0023)	(0.0044)	(0.0048)	(0.0074)
Mean of dep. var.	6.90		4.33		10.16		-5.40	
Observations	4,744	4,744	6,166	6,166	4,744	4,744	6,630	6,630
R-squared	0.95	0.95	0.97	0.97	0.87	0.87	0.63	0.63
Estimated 5-year change		-0.1916		-0.0611		0.0158		-0.0856
		(0.0917)**		(0.0417)		(0.0441)		(0.0697)
Estimated 11-year change		-0.2934		-0.1276		-0.0381		-0.0949
		(0.1483)**		(0.0710)*		(0.0672)		(0.1152)

Standard errors (clustered by county) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes indicators and time trends for hurricane events outside of the 21 year window of interest, year, county, and year-by-coastal fixed effects.

Figure 3. The effect of a hurricane on population and demographics

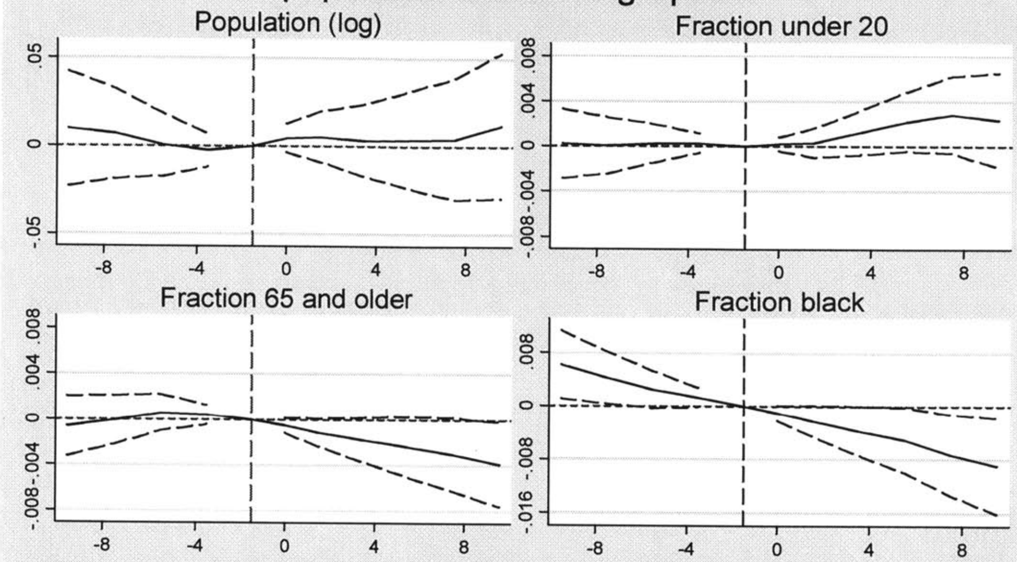

Sample: counties in Alabama, South Carolina, North Carolina, Florida, Georgia, Mississippi, Louisiana, Texas, and Virginia. 95% confidence intervals shown. Standard errors clustered by county. Controls include county, year, and coastal by year fixed effects.

Table 7: Mean shift and trend break tests for Figure 3

	Fraction black residents		Fraction 65 and older		Fraction 20 and younger		Population (log)	
Post hurricane	0.0007	0.0006	-0.0005	-0.0010	-0.0001	0.0002	0.0040	0.0068
	(0.0007)	(0.0007)	(0.0005)	(0.0006)*	(0.0005)	(0.0004)	(0.0048)	(0.0046)
Post hurricane time trend		-0.0001		-0.0005		0.0003		0.0028
		(0.0003)		(0.0002)**		(0.0002)*		(0.0016)*
Overall time trend	-0.0008	-0.0008	-0.0002	0.0001	0.0001	0.0000	-0.0001	-0.0017
	(0.0003)**	(0.0003)**	(0.0001)	(0.0002)	(0.0002)	(0.0002)	(0.0019)	(0.0020)
Mean of dep. var.	0.28		0.12		0.31		10.56	
Observations	6,860	6,860	6,899	6,899	6,899	6,899	6,899	6,899
R-squared	0.98	0.98	0.92	0.92	0.95	0.95	0.99	0.99
Estimated 5-year change		0.0000		-0.0034		0.0019		0.0209
		(0.0016)		(0.0013)***		(0.0011)*		(0.0096)**
Estimated 11-year change		-0.0006		-0.0058		0.0036		0.0349
		(0.0027)		(0.0023)**		(0.0020)*		(0.0167)**

Standard errors (clustered by county) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes indicators and time trends for hurricane events outside of the 21 year window of interest, year, county, and year-by-coastal fixed effects.

Figure 4. The effect of a hurricane on employment, earnings, and transfers

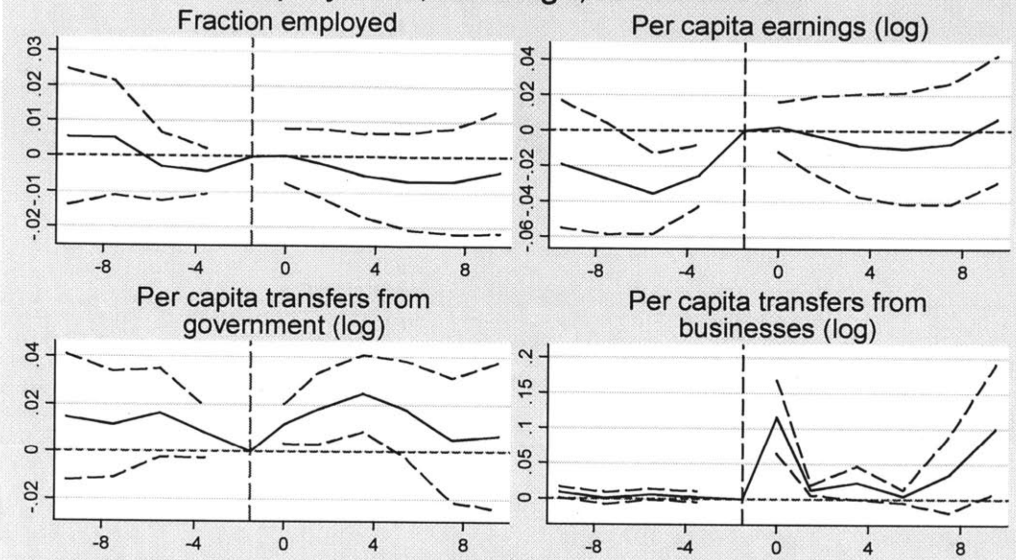

Sample: counties in Alabama, South Carolina, North Carolina, Florida, Georgia, Mississippi, Louisiana, Texas, and Virginia. 95% confidence intervals shown. Standard errors clustered by county. Controls include county, year, and coastal by year fixed effects.

Table 8: Mean shift and trend break tests for Figure 4

	Employment rate (fraction)		Per capita transfer from government (logs)		Per capita transfer from businesses (logs)		Per capita net earnings (log)	
Post hurricane	0.0019	0.0025	0.0205	0.0200	0.0302	0.0328	0.0061	0.0044
	(0.0057)	(0.0063)	(0.0071)***	(0.0068)***	(0.0152)**	(0.0120)***	(0.0136)	(0.0151)
Post hurricane time trend		0.0006		-0.0005		0.0027		-0.0017
		(0.0013)		(0.0023)		(0.0037)		(0.0031)
Overall time trend	-0.0007	-0.0010	-0.0016	-0.0013	0.0007	-0.0008	0.0011	0.0021
	(0.0009)	(0.0013)	(0.0012)	(0.0017)	(0.0022)	(0.0005)	(0.0012)	(0.0024)
Mean of dep. var.	0.58	0.58	8.09	8.09	4.37	4.37	9.61	9.61
Estimated mean change (levels)	3.40E-03	4.46E-03	67.40	65.84	2.43	2.65	91.55	65.68
Observations	6,860	6,860	6,860	6,860	6,706	6,706	6,860	6,860
R-squared	0.87	0.87	0.96	0.96	0.87	0.87	0.90	0.90
Estimated 5-year change		0.0054		0.0177		0.0463		-0.0043
		(0.0113)		(0.0137)		(0.0110)***		(0.0275)
Estimated 11-year change		0.0084		0.0153		0.0598		-0.0129
		(0.0174)		(0.0244)		(0.0280)**		(0.0427)

Standard errors (clustered by county) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes indicators and time trends for hurricane events outside of the 21 year window of interest, year, county, and year-by-coastal fixed effects.

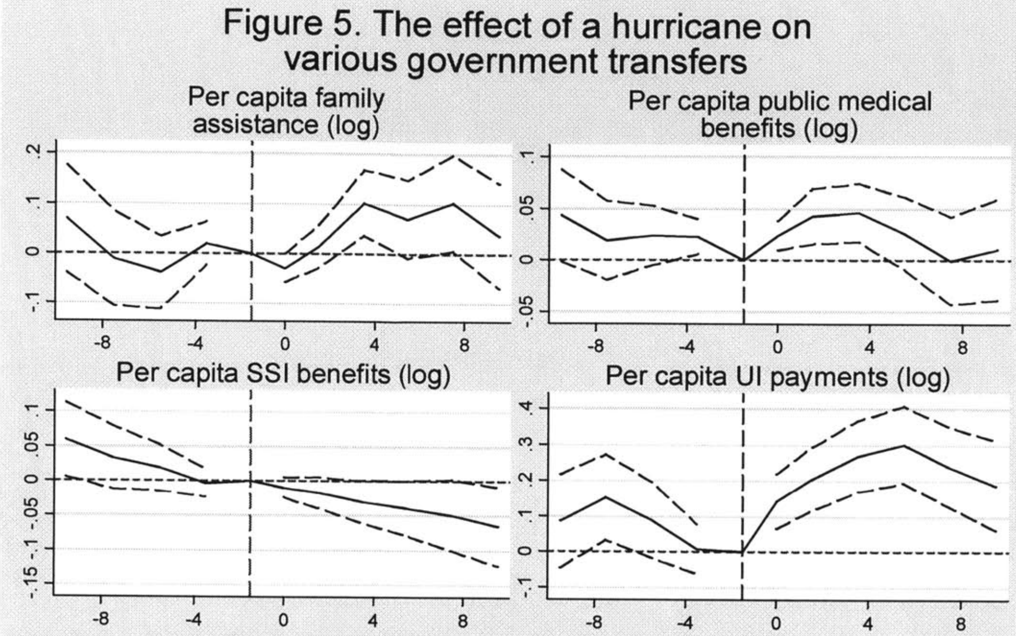

Sample: counties in Alabama, South Carolina, North Carolina, Florida, Georgia, Mississippi, Louisiana, Texas, and Virginia. 95% confidence intervals shown. Standard errors clustered by county. Controls include county, year, and coastal by year fixed effects.

Table 9: Mean shift and trend break tests for Figure 5

	Per capita family assistance (logs)		Per capita public medical benefits (log)		Per capita SSI benefits (log)		Per capita unemployment insurance (log)	
Post hurricane	0.0350	0.0459	0.0470	0.0461	0.0109	0.0115	0.2178	0.2346
	(0.0271)	(0.0269)*	(0.0137)***	(0.0127)***	(0.0109)	(0.0111)	(0.0538)***	(0.0559)***
Post hurricane time trend		0.0109		-0.0009		0.0006		0.0168
		(0.0097)		(0.0035)		(0.0038)		(0.0107)
Overall time trend	0.0013	-0.0049	-0.0042	-0.0037	-0.0066	-0.0070	-0.0051	-0.0147
	(0.0041)	(0.0070)	(0.0022)*	(0.0029)	(0.0024)***	(0.0035)**	(0.0048)	(0.0085)*
Mean of dep. var.	-2.87	-2.87	-0.02	-0.02	-1.93	-1.93	4.37	4.37
Estimated mean change (levels)	2.01E-03	2.66E-03	4.70E-02	4.61E-02	1.60E-03	1.69E-03	19.17	20.83
Observations	6,709	6,709	6,860	6,860	6,860	6,860	6,821	6,821
R-squared	0.82	0.82	0.97	0.97	0.86	0.86	0.65	0.65
Estimated 5-year change		0.1002		0.0418		0.0145		0.3185
		(0.0593)*		(0.0200)**		(0.0245)		(0.0877)***
Estimated 11-year change		0.1545		0.0374		0.0175		0.4025
		(0.1052)		(0.0355)		(0.0428)		(0.1342)***

Standard errors (clustered by county) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes indicators and time trends for hurricane events outside of the 21 year window of interest, year, county, and year-by-coastal fixed effects.

Table 10: Post-hurricane mean shift and trend breaks by 1970 median housing value

	Population (log)		Per capita earnings (log)		Employment rate		Per capita government transfers (log)		Per capita unemployment payments (log)	
Mean change in bottom quartile	-0.0059	-0.0010	0.0495	0.0911	0.0109	0.0168	0.0628	0.0602	0.2499	0.2177
	(0.0277)	(0.0270)	(0.0289)*	(0.0377)**	(0.0110)	(0.0124)	(0.0239)***	(0.0315)*	(0.0808)***	(0.1054)**
Mean change in top quartile	0.0186	0.0169	-0.0493	-0.0989	-0.0012	-0.0079	-0.0814	-0.0809	-0.0223	0.1041
	(0.0414)	(0.0390)	(0.0255)*	(0.0373)***	(0.0132)	(0.0154)	(0.0298)***	(0.0408)**	(0.0822)	(0.1213)
Trend change in bottom quartile		0.0028		-0.0112		-0.0003		-0.0002		0.0286
		(0.0020)		(0.0057)*		(0.0018)		(0.0040)		(0.0162)*
Trend change in top quartile		0.0003		0.0102		0.0014		-0.0001		-0.0263
		(0.0023)		(0.0052)*		(0.0018)		(0.0037)		(0.0161)
Estimated 11-year change in bottom quartile		0.0273		-0.0209		0.0138		0.0582		0.5040
		(0.0357)		(0.0573)		(0.0210)		(0.0317)*		(0.1711)***
Estimated 11-year change in top quartile		0.0199		0.0030		0.0057		-0.0816		-0.1593
		(0.0467)		(0.0368)		(0.0170)		(0.0284)***		(0.1089)
Difference in means/11-year changes (bottom - top)	-0.0245	0.0074	0.0987	-0.0239	0.0121	0.0081	0.1442	0.1398	0.2722	0.6632
p-value of difference	0.714	0.901	0.062	0.634	0.595	0.863	**0.006**	**0.012**	0.066	**0.008**
Observations	6,899	6,899	6,860	6,860	6,860	6,860	6,860	6,860	6,821	6,821
R-squared	0.99	0.99	0.90	0.90	0.88	0.88	0.97	0.97	0.66	0.66

Bold denotes p-values less than 0.05. Standard errors (clustered by county) in parentheses. * significant at 5%; ** significant at 1%. Includes year, county, year-by-coastal fixed effects, as well as quartile-specific trends.

Appendix Tables

Table A1: The effect of a hurricane on the construction sector

	Construction employment (log)	Construction establishments (log)	Construction per worker wage (log)	Per capita single family housing construction (log)
T = - 9 or - 10	-0.096	-0.044	-0.044	-0.073
	(0.073)	(0.043)	(0.037)	(0.056)
T = - 7 or - 8	-0.082	-0.045	-0.038	-0.066
	(0.056)	(0.034)	(0.030)	(0.056)
T = - 5 or - 6	-0.088	-0.021	-0.002	-0.071
	(0.041)**	(0.026)	(0.028)	(0.042)*
T = - 3 or - 4	-0.038	-0.005	-0.028	-0.012
	(0.026)	(0.017)	(0.023)	(0.033)
T = 0	-0.021	0.032	0.038	-0.057
	(0.028)	(0.013)**	(0.026)	(0.023)**
T = 1 or 2	-0.071	0.037	0.077	-0.009
	(0.034)**	(0.019)*	(0.027)***	(0.035)
T = 3 or 4	-0.160	-0.030	0.054	-0.097
	(0.046)***	(0.025)	(0.026)**	(0.047)**
T = 5 or 6	-0.193	-0.047	0.034	-0.079
	(0.052)***	(0.029)	(0.027)	(0.049)
T = 7 or 8	-0.155	-0.059	-0.009	-0.017
	(0.055)***	(0.034)*	(0.024)	(0.055)
T = 9 or 10	-0.117	-0.014	0.011	0.028
	(0.063)*	(0.037)	(0.026)	(0.061)
Mean of dep. var.	6.903	4.330	10.155	-5.396
Observations	4,744	6,166	4,744	6,630
R-squared	0.95	0.97	0.87	0.63
p-value of all leads F-test	0.279	0.655	0.345	0.305
p-value of all lags F-test	0.000	0.000	0.005	0.000
p-value of T=0 to T=4 lags F-test	0.001	0.000	0.020	0.003

Standard errors (clustered by county) in parentheses. * significant at 5%; ** significant at 1%. Includes year, county, and year-by-coastal fixed effects.

Table A2: The effect of a hurricane on population and demographics

	Fraction black residents	Fraction 65 and older	Fraction 19 and younger	Population (log)
T = - 9 or - 10	0.006	-0.001	0.000	0.010
	(0.002)**	(0.001)	(0.001)	(0.016)
T = - 7 or - 8	0.004	0.000	0.000	0.007
	(0.001)**	(0.001)	(0.001)	(0.012)
T = - 5 or - 6	0.003	0.001	0.000	0.001
	(0.001)*	0.000	0.000	(0.009)
T = - 3 or - 4	0.001	0.000	0.000	-0.003
	(0.000)*	0.000	0.000	(0.004)
T = 0	-0.001	-0.001	0.000	0.005
	(0.000)*	0.000	0.000	(0.004)
T = 1 or 2	-0.002	-0.001	0.000	0.005
	(0.001)*	(0.000)*	0.000	(0.007)
T = 3 or 4	-0.004	-0.002	0.001	0.003
	(0.001)*	(0.001)*	0.000	(0.010)
T = 5 or 6	-0.005	-0.002	0.002	0.004
	(0.002)**	(0.001)*	(0.001)	(0.014)
T = 7 or 8	-0.007	-0.003	0.003	0.004
	(0.003)**	(0.001)*	(0.001)	(0.017)
T = 9 or 10	-0.009	-0.004	0.002	0.012
	(0.003)**	(0.001)**	(0.002)	(0.020)
Mean of dep. var.	0.285	0.123	0.312	10.558
Observations	6,860	6,899	6,899	6,899
R-squared	0.98	0.92	0.95	0.99
p-value of all leads F-test	0.066	0.095	0.271	0.245
p-value of all lags F-test	0.091	0.378	0.011	0.124
p-value of T=0 to T=4 lags F-test	0.125	0.199	0.040	0.150

Standard errors (clustered by county) in parentheses. * significant at 5%; ** significant at 1%.
Includes year, county, and year-by-coastal fixed effects.

Table A3: The effect of a hurricane on employment, wages, and transfers

	Employment rate (fraction)	Per capita transfer from government (logs)	Per capita transfer from businesses (logs)	Per capita net earnings (log)
T = - 9 or - 10	0.005	0.014	0.009	-0.019
	(0.009)	(0.013)	(0.003)**	(0.018)
T = - 7 or - 8	0.005	0.011	0.001	-0.027
	(0.008)	(0.011)	(0.004)	(0.015)*
T = - 5 or - 6	-0.003	0.016	0.006	-0.036
	(0.004)	(0.009)*	(0.003)	(0.011)***
T = - 3 or - 4	-0.004	0.008	0.003	-0.025
	(0.003)	(0.005)	(0.003)	(0.008)***
T = 0	0.000	0.012	0.116	0.002
	(0.003)	(0.004)***	(0.026)***	(0.007)
T = 1 or 2	-0.002	0.018	0.013	-0.003
	(0.004)	(0.007)**	(0.003)***	(0.011)
T = 3 or 4	-0.005	0.025	0.023	-0.008
	(0.006)	(0.008)***	(0.011)*	(0.014)
T = 5 or 6	-0.007	0.018	0.004	-0.010
	(0.006)	(0.010)*	(0.004)	(0.015)
T = 7 or 8	-0.007	0.005	0.035	-0.007
	(0.007)	(0.013)	(0.027)	(0.017)
T = 9 or 10	-0.004	0.007	0.100	0.007
	(0.008)	(0.015)	(0.046)**	(0.018)
Mean of dep. var.	0.585	8.088	4.373	9.605
Observations	6,860	6,860	6,706	6,860
R-squared	0.87	0.96	0.87	0.90
p-value of all leads F-test	0.456	0.340	0.055	0.020
p-value of all lags F-test	0.757	0.003	0.000	0.107
p-value of T=0 to T=4 lags F-test	0.491	0.011	0.001	0.719

Standard errors (clustered by county) in parentheses. * significant at 5%; ** significant at 1%.
Includes year, county, and year-by-coastal fixed effects.

Table A4: The effect of a hurricane on various government transfers

	Per capita family assistance (logs)	Per capita public medical benefits (log)	Per capita SSI benefits (log)	Per capita unemployment insurance (log)
T = - 9 or - 10	0.068	0.043	0.059	0.085
	(0.054)	(0.022)*	(0.026)**	(0.066)
T = - 7 or - 8	-0.011	0.019	0.033	0.152
	(0.048)	(0.019)	(0.023)	(0.061)**
T = - 5 or - 6	-0.039	0.024	0.019	0.089
	(0.037)	(0.014)	(0.016)	(0.054)
T = - 3 or - 4	0.020	0.023	-0.003	0.008
	(0.022)	(0.008)***	(0.010)	(0.035)
T = 0	-0.029	0.024	-0.009	0.142
	(0.014)**	(0.007)***	(0.007)	(0.039)***
T = 1 or 2	0.015	0.043	-0.017	0.204
	(0.021)	(0.013)***	(0.011)	(0.044)***
T = 3 or 4	0.103	0.046	-0.030	0.268
	(0.033)***	(0.014)***	(0.015)*	(0.050)***
T = 5 or 6	0.070	0.026	-0.039	0.301
	(0.039)*	(0.017)	(0.020)*	(0.054)***
T = 7 or 8	0.103	0.000	-0.049	0.239
	(0.049)**	(0.021)	(0.026)*	(0.056)***
T = 9 or 10	0.037	0.011	-0.065	0.187
	(0.053)	(0.024)	(0.029)**	(0.063)***
Mean of dep. var.	-2.873	-0.023	-1.926	4.367
Observations	6,709	6,860	6,860	6,821
R-squared	0.82	0.97	0.86	0.65
p-value of all leads F-test	0.000	0.002	0.002	0.004
p-value of all lags F-test	0.000	0.000	0.278	0.000
p-value of T=0 to T=4 lags F-test	0.000	0.005	0.113	0.000

Standard errors (clustered by county) in parentheses. * significant at 5%; ** significant at 1%.
Includes year, county, and year-by-coastal fixed effects.

Table A5: Post-hurricane mean shift and trend breaks by county income

	Population (log)		Per capita earnings (log)		Employment rate		Per capita government transfers (log)		Per capita unemployment payments (log)	
Mean change in bottom quartile	-0.0224	-0.0235	0.0534	0.0696	0.0133	0.0186	0.0571	0.0724	0.1898	0.1445
	(0.0207)	(0.0201)	(0.0236)**	(0.0300)**	(0.0114)	(0.0130)	(0.0186)***	(0.0230)***	(0.0851)**	(0.0992)
Mean change in top quartile	0.0144	0.0152	-0.0365	-0.0638	0.0046	-0.0004	-0.0798	-0.1095	-0.0040	0.0334
	(0.0344)	(0.0333)	(0.0241)	(0.0346)*	(0.0160)	(0.0175)	(0.0237)***	(0.0311)***	(0.1023)	(0.1328)
Trend change in bottom quartile		0.0047		-0.0066		-0.0002		-0.0034		0.0386
		(0.0018)***		(0.0054)		(0.0016)		(0.0031)		(0.0148)**
Trend change in top quartile		-0.0002		0.0056		0.0010		0.0061		-0.0075
		(0.0020)		(0.0044)		(0.0015)		(0.0026)**		(0.0135)
Estimated 11-year change in bottom quartile		0.0235		0.0033		0.0162		0.0381		0.5301
		(0.0285)		(0.0527)		(0.0189)		(0.0287)		(0.1659)***
Estimated 11-year change in top quartile		0.0137		-0.0078		0.0097		-0.0484		-0.0416
		(0.0384)		(0.0311)		(0.0177)		(0.0217)**		(0.1102)
Difference in means/11-year changes (bottom - top)	-0.0368	0.0098	0.0898	0.0111	0.0087	0.0065	0.1369	0.0865	0.1939	0.5717
p-value of difference	0.482	0.814	**0.044**	0.989	0.734	0.868	**0.001**	0.092	0.253	**0.017**
Observations	6,899	6,899	6,860	6,860	6,860	6,860	6,860	6,860	6,821	6,821
R-squared	0.99	0.99	0.90	0.90	0.88	0.88	0.97	0.97	0.66	0.66

Bold denotes p-values less than 0.05. Standard errors (clustered by county) in parentheses. * significant at 5%; ** significant at 1%. Includes year, county, year-by-coastal fixed effects, as well as quartile-specific trends.

Table A6: the co-movement of local economic indicators

	Total unemployment payments (log)		Per capita unemployment payments (log)		Earnings per capita (log)	Net earnings (log)	State receipts (log)	
Total employment (log)	0.063 (0.035)*		-0.329 (0.032)***					
Total wages paid (log)		0.086 (0.031)***		-0.231 (0.026)***				
Net earnings per capita (log)								0.588 (0.103)***
Net earnings (log)							0.823 (0.031)***	
Income maintenance (log)					-0.064 (0.008)***	0.210 (0.022)***		
Observations	33,844	34,767	33,822	34,745	36,809	36,809	12,168	12,168
R-squared	0.91	0.91	0.62	0.62	0.85	0.98	0.97	0.96

Standard errors (clustered by county) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes county and year fixed effects.

Table A7: the co-movement of the local construction sector with earnings

	Total establishments (log)		Total employment (log)		Total pay (log)		Per worker pay (log)	
Net earnings per capita (log)		0.733 (0.056)***		1.058 (0.078)***		1.259 (0.089)***		0.229 (0.026)***
Net earnings (log)	0.899 (0.027)***		0.961 (0.036)***		1.137 (0.042)***		0.182 (0.013)***	
Observations	34,352	34,352	30,872	30,872	31,751	31,751	29,743	29,743
R-squared	0.97	0.96	0.95	0.94	0.95	0.94	0.78	0.78

Standard errors (clustered by county) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes county and year fixed effects.

Table A8: autocorrelation in local economic indicators

	Panel A: general							
	Total population (log)	Gov't --> ind transfers	Business --> ind transfers	Income maintenance (log)	Earnings per capita (log)	UE compensation (log)	UE per capita compensation (log)	
Lagged variable	0.984	0.966	0.396	0.892	0.693	0.737	0.729	
	(0.002)***	(0.003)***	(0.029)***	(0.005)***	(0.021)***	(0.007)***	(0.007)***	
Observations	36,791	35,834	32,414	35,865	35,870	35,794	35,744	
R-squared	1.00	1.00	0.96	1.00	0.92	0.96	0.82	
	Panel B: producer-side							
	Construction employment (log)	Total employment (log)	Construction establishments (log)	Total establishments (log)	Construction total pay (log)	Total pay (log)	Construction per worker pay (log)	Total per worker pay (log)
Lagged variable	0.785	0.895	0.858	0.957	0.820	0.896	0.095	0.275
	(0.007)***	(0.008)***	(0.006)***	(0.003)***	(0.006)***	(0.007)***	(0.012)***	(0.023)***
Observations	30,730	34,844	35,233	36,002	31,630	35,828	29,871	34,799
R-squared	0.98	1.00	0.99	1.00	0.98	0.99	0.77	0.91

Standard errors (clustered by county) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes county and year fixed effects.

Table A9: Monte Carlo simulation, heterogeneous trends with autocorrelation

Time relative to event	Theoretical effect of event[1]	Estimation without including lagged variable		Estimation including lagged variable	
T= - 5	0	-0.809 (0.067)***		-0.084 (0.036)**	
T= - 4	0	-0.793 (0.066)***		-0.081 (0.036)**	
T= - 3	0	-0.767 (0.066)***		-0.046 (0.036)	
T= - 2	0	-0.877 (0.066)***		-0.131 (0.036)***	
T= - 1	0	reference category		reference category	
T= 0	10.00	9.084 (0.067)***	9.423 (0.061)***	9.889 (0.036)***	9.947 (0.030)***
T= 1	9.00	8.074 (0.066)***	8.441 (0.061)***	0.698 (0.050)***	0.554 (0.041)***
T= 2	8.10	7.234 (0.067)***	7.632 (0.061)***	0.684 (0.047)***	0.538 (0.039)***
T= 3	7.29	6.372 (0.067)***	6.827 (0.062)***	0.557 (0.045)***	0.464 (0.038)***
T= 4	6.56	5.626 (0.067)***	6.083 (0.062)***	0.508 (0.043)***	0.403 (0.037)***
T= 5	5.90	4.943 (0.068)***	5.413 (0.062)***	0.398 (0.042)***	0.334 (0.035)***
Lagged outcome				0.82 (0.004)***	0.842 (0.003)***
Observations		20,000	25,000	20,000	25,000
R-squared		0.89	0.87	0.97	0.97
p-value of all leads F-test		0.000		0.001	
p-value of all lags F-test		0.000	0.000	0.000	0.000

[1]When lagged outcome is excluded from the estimation and time trends are appropriately controlled for.

Standard errors in parentheses. * significant at 5%; ** significant at 1%. Includes fixed effects for county and dummies for "event more than 5 years ago" and "event more than 5 years in the future".

Chapter 3

Selection in Area Yield Crop Insurance

3.1 Introduction

Despite a long-run decrease in developed countries' vulnerability to weather shocks, agriculture worldwide remains susceptible to weather fluctuations. According to the Intergovernmental Panel on Climate Change (IPCC), climate change will increase the frequency and intensity of extreme weather events and change their spatial distribution (Meehl et al., 2007; Schneider et al., 2007). This is estimated to be costly to the agricultural sector (Deschenes and Greenstone, 2007). Yield variability and the variance of food prices are likely to increase, at least in the short run. A functioning insurance market may be key to keeping the agricultural sector stable, especially in developing countries.

A number of developed and developing countries have crop insurance markets; most of these are public or subsidized. A private market in crop insurance has been difficult to achieve. Standard explanations for this include moral hazard and adverse selection, coupled with large monitoring costs. In addition, unlike other insurance markets, such as health or auto, there is greater potential for sudden aggregate shocks. The difficulty in establishing a private market, the potential for aggregate shocks, and impending climate change warrant devoting more attention to the crop insurance market.

In this paper, I examine the US insurance market for two common crops: corn and soybeans. I discuss the general structure of available insurance plans. I note the moral hazard and adverse selection concerns in crop insurance, how the US crop insurance market design addresses them, and which features of insurance plans leave them susceptible to adverse selection and moral hazard. Despite the attempts of market designers to minimize moral hazard and adverse selection, these plans remain susceptible to strategic behavior.

I then test for a particular form of adverse selection into area yield insurance, where indemnity payments to farmers are based on the average yield in the county and not on individual yields. These plans are also known as "group insurance" plans. Although the group yield plans are not very popular, they provide the cleanest test of adverse selection in this market. Moreover, because monitoring costs and moral hazard concerns are low in these plans, group yield or weather-based insurance is thought to be the most efficient way to provide crop insurance in many developing countries. Thus, testing for adverse selection in this context is relevant to a wide range of markets.

Similar to Finkelstein and Poterba (2006) and Finkelstein and McGarry (2006), I use the presence of information that affects outcomes but is not used in pricing to test for the presence of adverse selection. Specifically, because yields are autocorrelated but prices don't account sufficiently for recent yield shocks, recent yields are informative about future yields and should affect farmers' insurance decisions. Thus, the first test consists of checking whether last year's yield (which is predictive of the current yield but empirically does not affect the price) affects the total demand for group insurance.

The first test is "structural" in that it assumes that the mechanism through which farmers adversely select is last year's yield. I also perform a reduced-form test of adverse selection that does not require knowing the selection mechanism. In addition to last year's yields, farmers or insurers may have other information about future yields that is not reflected in prices. Even though this information may not be observable to the econometrician, its presence should lead to a correlation between contemporaneous yields and takeup, after controlling for other factors. Three features of group insurance make it possible to attribute any residual correlation between the current yield and takeup to selection. First, the pricing formula for group insurance plans is known. Second, because group insurance is typically offered in counties that have many farmers, there is no potential for moral hazard. Finally, the insurance decision is made before crops are planted (and thus before yields are known). Thus, the reduced-form test for selection consists of checking whether the number of group insurance policies is significantly correlated with the current yield. The validity of this test relies only on the fact that the presence of a selection mechanism described above should lead to a significant correlation between contemporaneous yields and takeup.

I find no evidence that last year's yields influence takeup of group insurance plans, although the estimates are not very precise. However, the second (reduced-form) test indicates that group insurance takeup is higher when average current yields in the county are higher. This suggests that the net selection into area yield plans favors providers, not buyers of insurance. Although this is surprising, it is consistent with earlier findings that insurance companies may obtain significant excess rents through reinsurance decisions based on weather, a form of supply-side adverse selection (Ker and McGowan, 2000).

It may be unlikely that providers have more information than individuals in many other insurance markets. However, it is plausible in the market for area yield insurance. Providers may observe detailed yield information for many farmers in the county (through information provided to buy individual insurance plans, which are very popular) and thus be able to form better aggregate yield forecasts than individual farmers. Because providers cannot compete on prices by law, the selection may be working through insurance agents convincing farmers to choose one plan over another or through targeting particular counties in years when yields are likely to be high.

Finally, I find that in some specifications current and last year's yields are predictive of the total number of insurance policies in the county (including individual plans). This suggests that the desirability of other plans may be changing with yields as well, affecting other options in the farmers' choice set. However, because prices in non-group plans are determined using individual yields, which I do not observe, I cannot determine whether the relationship between aggregate takeup of insurance and yields is due to selection or changes in prices.

I contribute to two strands of literature. First, I perform the first test of adverse selection into area yield insurance plans. Although there have been numerous studies of moral hazard and adverse selection in crop insurance, these have typically examined the correlation between farmer characteristics and crop insurance decisions (e.g., Makki and Somwaru, 2001) or focused on a narrow geographical area (Smith and Goodwin, 1996; Horowitz and Lichtenberg, 1993; Coble et al., 1997; Roberts et al., 2006). Several program-wide tests have been performed (e.g., Walters et al., 2007), but to my knowledge, this is the first test to use unpriced public information as the cause of adverse selection.

Second, I contribute to the body of literature testing for moral hazard and adverse selection in various insurance markets. These include health insurance (e.g., Altman et al., 1998; Cutler and Reber, 1998; Simon, 2005; Einav et al, 2011; Handel, 2011), long-term care insurance (e.g., Sloan and Norton, 1997; Finkelstein and McGarry, 2006; Brown and Finkelstein, 2007), annuities (Finkelstein and Poterba, 2002 and 2004l; Fong, 2002; McCarthy and Mitchell, 2010), nursing home use (Gruber and Grabowski, 2007), auto insurance (Dionne et al., 2004), and credit cards (Agarwal et al., 2010). For a broader overview of the evidence for adverse selection in various insurance markets, see Cohen and Siegelman (2010). Complementary studies across different markets help distinguish features common to all insurance markets from idiosyncratic ones.

The rest of the paper is organized as follows. Section 2 provides a brief overview of the US crop insurance market. Section 3 develops a model of crop insurance choice within the existing market structure and discusses how adverse selection and moral hazard could arise. In Section 4, I discuss the empirical framework used to test for selection and moral hazard. Section 5 presents the results, and Section 6 concludes.

3.2 The US Crop Insurance Market

Crop insurance plans in the US differ in the metric that determines payments and the level of the deductible, called the "coverage level".[1] Payments can be determined by (a) individual yield, (b) individual revenue, (c) mean county yield or (d) mean county revenue. Farmers cannot take out multiple insurance plans for the same plot. Within these plan types, farmers can choose from several

[1] For a more comprehensive overview of the US crop insurance market, see Babcock (2011).

coverage levels ranging from 50% to 90%.[2] The coverage level specifies the amount by which yield or revenue has to fall (relative to a baseline) before any payment is made. If a farmer chooses a 75% coverage level, for example, he does not receive payments until his yield, revenue, the county yield or the county revenue (depending on the type of plan) falls to more than 25% below the established baseline. The plan type and coverage level largely describe the space of all insurance plans available to farmers.[3] In this section, I describe the individual yield and county yield plans. Individual revenue and county revenue plans are discussed briefly in Appendix A.

To reduce the amount of moral hazard and adverse selection, farmers are required to initiate the purchase of insurance by a certain date, called the "sales closing date". This date varies by county, crop, and year and precedes the earliest allowed planting date. In some circumstances, a farmer may purchase insurance after the sales closing date, but these circumstances are limited.[4] In all cases, the insurance decision is made months before yields for that year are realized.

Because farmer productivity varies, individual baseline yields are necessary for individual-level plans to correctly determine the basis on which payment should be made. The baseline yield for an individual yield plan is established by averaging a farmer's historic certified yields (4 consecutive years is the minimum and 10 years is the maximum). Once 10 years of continuous yield history is available, the baseline yield becomes a 10-year moving average, updated every year. If less than 4 years of continuous yield history is available, average county yields are used in place of an individual yield until the farmer builds up an adequate yield history. If county yields are used to calculate the baseline yield, they are discounted by $15 - 35\%$, depending on how many years of actual yield history are available.

Per-acre payments under the individual yield plan (called Actual Production History or APH) are determined by the following formula:

$$\text{Pay\_APH}_{it} = \max(0, [\text{Yield guarantee}_{it} - \text{Yield}_{it}] \times \text{price election}_i) \tag{3.1}$$

$$\text{Yield guarantee}_{it} = X_i * \frac{1}{10}\sum_{s=1}^{10} \text{Yield}_{i,t-s}$$

where i indexes the farmer, t indexes the year, and X_i is the chosen coverage level. For individual yield plans, it ranges from 50% to 85%, in 5 percentage point increments. Price election is the payment per unit of yield shortfall, chosen by the farmer from a range set by the Federal Crop Insurance Corporation (FCIC) or the Risk Management Agency (RMA). For example, a corn farmer with a yield guarantee of 148 bushels per acre, a 75% coverage level, and a \$1.50 price election will get paid \$1.5 for every unit shortfall in yield below 111 bushels per acre.

[2]Not all coverage levels are available for all plan types and in all counties

[3]Farmers also have some choices within a plan-coverage-level combination, such as how to combine different plots and how much to get paid in the case of a shortfall. These are briefly discussed below.

[4]One example is where the farmer first failed to plant a different crop.

The group yield plan, also known as county yield or area yield plan, is based on the deviation of current county-level yields from their historic average. Because it is not based on individual yields, individual yield histories are not necessary to sign up for it. Baseline county yields are based on at least 30 years of yield history, not 10, and are trend-adjusted to reflect long-run productivity changes.[5] The payment formula is similar to the individual yield plan with the exception that the yield guarantee and actual yield are based on county, not individual, yields. Per-acre payments in the group yield plan (Group Risk Plan or GRP) are determined as follows:

$$\text{Pay\_GRP}_{it} = \max\left(0, \left[\text{Yield guarantee}_{ict} - \overline{\text{Yield}}_{ct}\right] \times \text{price election}_i\right) \tag{3.2}$$

$$\text{Yield guarantee}_{ict} = X_i * \frac{1}{30}\sum_{s=1}^{30} \widehat{\text{Yield}}_{c,t-s}$$

$\widehat{\text{Yield}}_{c,t-s}$ is the de-trended average county yield. Payments are made if county yields fall far enough below the yield guarantee. To ensure that farmers aren't simply "gambling" on yields, farmers have to plant the relevant crop in order to participate in the group yield plan.

Plans similar to GRP have been advocated in developing nations because the enforcement costs are much lower and the risk of moral hazard is completely eliminated, as long as the covered area is large enough so that no farmer is able to affect the average county yields on his own. One drawback of these plans is that they may not provide as much protection as plans based on individual yields. Another is that they may lead to some inefficient behavior, such as a farmer planting a crop in inappropriate conditions because planting is a requirement for participation. Finally, as I discuss below, they may leave open opportunities for adverse selection.

3.3 Conceptual Framework

I now outline a simple discrete choice model for crop insurance and describe the potential for moral hazard and adverse selection given the market structure explained in the previous section. For ease of exposition, I assume risk neutrality; adding risk aversion does not change the fundamental points below.

At time $t-1$, the expected utility of farmer i from having insurance plan j at time t can be written as:

$$u_{ij,t-1} = \beta E_{t-1}\left[Indem_{ijt}\right] - \gamma P_{jlt} + \psi_j + \varepsilon_{ijt}$$

$E_{t-1}\left[Indem_{ijt}\right]$ is the expectation of the indemnity payment in plan j at time t, formed at $t-1$

[5]Data on yield histories used to establish county yields are provided by the National Agricultural Statistics Service (NASS).

(when the insurance decision is made). Because indemnity payments in individual revenue and individual yield plans depend on the individual yield history, the expected indemnity payment can be farmer-specific. P_{jlt} is the time t price of plan j in county l, the location of the farmer. ψ_j is an unobservable (to the econometrician) fixed plan characteristic. The expected utility of the outside option (not having insurance) is normalized to 0 ($u_{i0,t-1} = 0$). I abstract from some within-plan choices of how to combine insurable units and which price election to choose. To the extent that the optimal choice of these options within a plan is unaffected by changes in the expected indemnity payment or price, the additional options can be thought of as a part of the fixed plan characteristic.

If ε_{ijt} are iid extreme value, the probability of an individual choosing plan j at time $t-1$ is:

$$\Pr_{ijt} = \frac{\exp\left(\beta E_{t-1}\left[Indem_{ijt}\right] - \gamma P_{jlt} + \psi_j\right)}{\sum_k \exp\left(\beta E_{t-1}\left[Indem_{ikt}\right] - \gamma P_{klt} + \psi_k\right)}$$

The expression for $E_{t-1}\left[Indem_{ijt}\right]$ varies by plan. For individual yield plans, it is:

$$E_{t-1}\left[Indem_{ijt}\right] = (1-\theta)\, C_{ji} E_{t-1}\left[X_j \overline{Yield}_{i,t-1} - Yield_{it} | X_j \overline{Yield}_{i,t-1} - Yield_{it} < 0\right]$$

where $\overline{Yield}_{i,t-1}$ is the baseline yield, established based on the available yield history and $Yield_{it}$ is the actual yield at time t. θ is the probability that $X_j\overline{Yield}_{i,t-1} - Yield_{it}$ is greater than 0 (i.e., that the indemnity payment is 0). C_{ji} is the price election (how much the farmer receives for every unit of yield shortfall) chosen by the farmer. A higher price election will also increase the price of the plan. X_j is the coverage level, expressed as a fraction between 0.5 and 1; the yield has to be lower than $1 - X_j$ of the baseline yield before a farmer receives any payment.

For group yield plans, the expected indemnity payment is:

$$E\left[Indem_{ijt}\right] = (1-\theta)\, C_{jt} E_{t-1}\left(X_j \overline{Yield}_{l,t-1} - Yield_{lt} | X_j \overline{Yield}_{l,t-1} - Yield_{lt} < 0\right)$$

This expression is similar to the individual yield plans, except that the yields are now indexed by l, the county that individual i is located in. Thus, indemnity expectations from a group yield plan do not vary by individual. The expected indemnity expressions for individual revenue and county revenue plans are similar, except that yields are replaced by revenues.

The design of the individual insurance plans appears to lend itself to moral hazard. After taking out insurance, farmers may reduce the use of costly inputs, such as fertilizer. Alternatively, they may put in less effort during harvest time. However, because baseline yields are calculated using a 10-year running average, the incentive for yield-reducing moral hazard is attenuated. Although lower-than-baseline yields one year will increase insurance payments, they will also decrease the baseline yield in the following year. The extent to which moral hazard is a problem in this case depends on the farmer's discount rate and expectations about keeping the farm and the individual insurance plan.

However, the use of the 10-year running average does create incentives for adverse selection.

Because a 10-year average is unlikely to reflect the true average yield, farmers who have worse-than-average yield histories have an increased incentive to forego individual insurance until their yield histories improve. Similarly, farmers who experienced better-than-average yields in the past 10 years have an increased incentive to take out insurance in the following year. In other words, if $\overline{Yield}_{i,t-1}$ is substantially below (above) the long-run yield average (assumed to be known to the farmer, but unobservable to the insurer), the farmer has an incentive to decrease (increase) his insurance coverage. Moreover, the mean is only one characteristic of yield distributions; farmers whose yield distributions differ from those used in actuarial calculations may also be able to adversely select into crop insurance plans, even if their mean yield is perfectly measured.

There is no moral hazard in the group yield (GRP) and group revenue (GRIP) plans unless a farmer is "large" enough to substantially affect group yields. There remains the potential for adverse selection based on private information about likely group yields or on public information that isn't priced into the insurance plan (e.g., the forecasted weather conditions). Farmers whose yields generally track group yields have a greater incentive to sign up. However, unless their farms are large enough to affect group yields, expected GRP and GRIP payments are not affected by who signs up for these plans.

Because there is uncertainty in the utility function above, informational asymmetries can arise in this market. This can affect plan choice in at least two ways. First, farmers may have private information about their own future yield shocks ($E_{t-1}[Yield_{it}]$) based on evolving field characteristics or some other idiosyncratic component. Second, as I show later, because yields are autocorrelated across years, farmers may have information about expected yields that is public but not used in pricing the plans. This may lead to adverse selection and undermine the functioning of the market. In the next two sections, I outline the estimation procedure and present the estimates of this type of adverse selection.

3.4 Empirical Framework

Ideally, one would observe individual farmers' yield histories and test for adverse selection across all plans. However, the available data provide only plan-by-coverage-level summaries, making a clean test of selection in individual insurance plans infeasible. Therefore, I focus on group yield insurance plans (GRP), where all the relevant pricing information is observable. In addition, group plans do not create the risk of moral hazard, making this a pure test of adverse selection. I focus on corn and soybeans because these are the most commonly grown crops and thus have the largest number of relevant observations. I perform the analysis separately for each crop.

To test for the presence of adverse selection, I follow Finkelstein and McGarry (2006). Specifically, I estimate the following equations :

$$Log(Yield_{lt}) = \beta_1 Log(Yield_{l,t-1}) + \beta_2 Log(E_{t-1}[Yield_{lt}]) + \alpha_t + \alpha_l + \varepsilon_{lt} \tag{3.3}$$

$$Policies_{lt} = \gamma_1 Log(Yield_{l,t-1}) + \gamma_2 Log(E_{t-1}[Yield_{lt}]) + \alpha_t + \alpha_l + \varepsilon_{lt} \tag{3.4}$$

where l is the county and t is the year. $E_{t-1}[Yield_{lt}]$ is the expected corn or soybean yield in county l and year t, as reported by the Risk Management Agency (RMA).[6] $Yield_{l,t-1}$ is the yield in the previous year, also reported by RMA. $Policies_{lt}$ is the number of corn GRP policies held. α_t and α_l are year and county fixed effects, respectively. The specification for equation (3.3) is OLS, while the specification for equation (3.4) is negative binomial. Because there are relatively few states in the sample, standard errors are clustered by state-year in this and all following specifications. Clustering by county generally decreases the standard errors.

Equation (3.3) tests whether last year's yield is predictive of current yield, controlling for the expected yield (which is used to set the price of the insurance policies). Although last year's yield is used in computing the expected yield for group insurance plans as part of 30 or more years of yield data, it may have an independent predictive value if there is significant autocorrelation in yields.[7] If β_1 is significantly different from 0, then there is potential for adverse selection into the GRP plans.

Equation (3.4) tests for the presence of such adverse selection. In particular, I estimate whether last year's yield affects the number of policies chosen, conditional on the expected yield. The coefficient of interest is thus γ_1. Depending on the relationship between current and last year's yield, it may be positive or negative. All else equal, farmers have an incentive to take out insurance if they know that yields will be lower than expected by the insurance company. Thus, if β_1 is positive (negative), γ_1 should be negative (positive) when adverse selection is present.

There is another test of selection that can be performed in this setting because of the specific features of group insurance plans. First, there is no moral hazard in GRP plans because they are based on the average county yield and are only available in counties with many farms. Second, the insurance choice in a given year is made months *before* yields in that year are realized. Finally, as I show in Appendix B, contemporaneous yields and prices are not correlated.

Suppose there is some information accessible to farmers (e.g., weather or changes in soil quality) that is not reflected in the price but is predictive of yield in the upcoming season. If farmers act on this information, there should be a significant relationship between current yields and the number of policies, even though current yields are not known when the insurance decisions are made. Therefore, to perform this reduced-form test, I estimate the relationship between the current yield and the number of GRP policies, controlling for the expected yield:

$$Policies_{lt} = \theta_1 Log(Yield_{lt}) + \theta_2 Log(E_{t-1}[Yield_{lt}]) + \alpha_t + \alpha_l + \varepsilon_{lt} \tag{3.5}$$

Instead of looking at selection on past observable information, this specification tests for selection on future yield realizations, controlling for the expected yield. Because farmers must purchase

[6] Available on http://www.rma.usda.gov/data/grpfinal/

[7] The RMA calculation of the expected yield does not allow for possible autocorrelation.

insurance long before the current yield is known, θ_1 implicitly captures the effect of other (un-priced) information that farmers may have about yield realizations without having to measure this information explicitly.

Finally, I test that prices (as reported in RMA actuarial documents) are not affected directly by last year's or current yield:

$$Log\,(P_{glt}) = \beta_1 Log(Yield_{l,t-1}) + \beta_2 Log(E_{t-1}\,[Yield_{lt}]) + \alpha_{gt} + \alpha_{gl} + \varepsilon_{glt} \tag{3.6}$$

$$Log\,(P_{glt}) = \delta_1 Log(Yield_{lt}) + \delta_2 Log(E_{t-1}\,[Yield_{lt}]) + \alpha_{gt} + \alpha_{gl} + \varepsilon_{glt} \tag{3.7}$$

g = coverage level; t = year; l = location (county)

I perform this test at a more disaggregated level than the adverse selection test. If the relative proportion of people in each coverage level changes because of adverse selection, the average price across coverage levels might be statistically but not causally related to last year's yields. I use the unsubsidized prices set by the insurance providers. Because premium subsidy rates vary only across years and not across counties, this does not affect the results. The results of this regression are shown in Tables A1 and A2 and confirm the hypothesis that last year's and current yields do not affect prices. [CHECK] The pricing and expected yield procedure used by the RMA is similar to the procedure described in Skees, Black, and Barnett (1997). Note that they highlight the tradeoff of more actuarily fair v. implementable prices. Moreover, Harri et al. (2011) note that the current pricing algorithm does not account for heteroskedasticity in yields observed in some US counties.

I use crop insurance data from 1995-2009, published by the Risk Management Agency (RMA).[8] The data are annual and broken down by county, crop, insurance plan type, and coverage level. They contain information about the number of acres insured, total liabilities, indemnity payments, premiums, and premium subsidies. These data include the premium and subsidy rate for each plan, crop, and county over time.[9]

Table 1 shows the summary statistics for the key variables used in the analysis, including yields, number of group insurance policies, and the total number of insurance policies. Columns 1 and 2 show the mean and standard deviations for corn and soybeans, respectively. The average number of insurance policies is similar for corn (169 policies) and soybeans (179 policies). Group Risk Plans (GRP) are available in 1,076 counties for corn and 961 counties for soybeans at least once during the time period of interest. GRP plans represent, on average, about 2-3 percent of all insurance policies. Thus, they are fairly unpopular. A likely explanation for this is that individual plans offer more tailored risk protection and generally have higher indemnity payment-to-premium ratios.

[8] Available from http://www.rma.usda.gov/data/sob/scc/index.html

[9] Available from ftp://ftp.rma.usda.gov/pub/Publications/FCI-35/

3.5 Results

Before presenting the formal regression results, I illustrate the raw relationships between actual, expected, and last year's yields. Later, I re-examine this relationship taking fixed effects into account. Figure 1 shows the relationships between actual and expected yields for corn and soybeans. Although there is a substantial amount of variance, the two are strongly and positively related for both crops. However, a fitted quadratic relationship lies above the 45 degree line for corn, indicating that expected yields are generally below actual yields. For soybeans, the fitted line is somewhat flatter than the 45 degree line, but does cross it. Overall, the expected yield is most biased at high and low yields, but the magnitude does not appear to be large.

Figure 2 shows the relationships between actual yield and last year's yield for the two crops. Again, the two are strongly and positively related. The plots for soybeans and corn look very similar. A quadratic fit line indicates that low (high) yields last year are indicative of higher (lower) yields this year. This is true for both crops and suggests that there is some mean reversion in yields. If this is not accounted for in the price, there is potential for exploiting this information.

Table 2 shows the relationship between last year's yield and current yield for each crop, controlling for expected yield, county fixed effects and year fixed effects. Columns 1 and 2 have the log of the current yield as the dependent variable, while Columns 3 and 4 have the level of the current yield as the dependent variable. In Columns 2 and 4, I restrict the sample to counties in which GRP policies are ever taken up. The results are similar in both samples. Panel A shows these results for corn. There is no level-level relationship between this year's and last year's yield and a very weak one between the levels of this year's yield and expected yield. For soybeans (Panel B), a one unit increase in yields in a given year is associated with a 0.08-0.09 unit decrease in yields the next year, and this relationship is statistically significant. The relationship between logs of yield and lagged yield is strong and significant for both crops, however.

According to Columns 1 and 2, a 1% increase in corn yields is associated with a $0.09 - 0.12\%$ decrease in yields in the following year, holding expected yield constant. Thus, yields appear to exhibit some mean reversion. This process is not necessarily random: the lower yields following higher yields may be due to patterns of nutrient depletion in the soil or of crop rotation. The expected yield is also negatively correlated with the actual yield: a 1% increase in the expected yield is associated with a $0.32 - 0.38\%$ reduction in the actual yield. This is conditional on county and year fixed effects, as well as last year's yield; regressing current corn yield on the expected yield with no controls results in a partial correlation coefficient of 0.97 in levels and 0.53 in logs. The results for soybeans are very similar.

Overall, current yield is predictive of future yield and, if not taken into account in prices, may result in adverse selection. If last year's yield was average, the current yield is expected to be lower than average. Holding price and expected yields constant, this should increase the demand for GRP insurance because the likelihood of a payout is higher. Ceteris paribus, if farmers have private or unpriced information that this year's yield will be low (high), they should increase (decrease) their demand for insurance. Thus, the relationship between the number of GRP policies and the current yield should be negative. Because last year's and current yields are negatively correlated, the relationship between the number of GRP policies and last year's yield should be positive.

Table 3 shows the relationship between (a) log of last year's yield and the number of GRP policies and (b) log of this year's yield and the number of GRP policies, controlling for the log of expected yield, county fixed effects and time fixed effects. The regression specification is a negative binomial. There is no significant relationship between last year's yield and the number of corn or soybean area yield policies, although the large standard errors make the estimate imprecise.

Surprisingly, Table 3 also shows that current yields are positively correlated with insurance takeup for both corn and soybeans. This suggests that, on net, selection into GRP plans acts in favor of insurers: in years when yields will be high (controlling for expected yield) more corn and soybean farmers will take out county yield insurance.

Appendix B shows the results of an OLS regression of the log of GRP policies on current and past yields, controlling for expected yield and fixed effects. In general, the results for corn are insignificant and those for soybeans are marginally significant. However, due to the count nature of the data, a negative binomial specification is more appropriate.

One issue with equation (3.4) is that it may not adequately control for other changes that may be correlated with last year's yields. In particular, changes in last year's yields may affect the desirability of individual yield or individual revenue plans. This does not matter for the validity of the test of adverse selection into group yield plans. However, in order to understand the behavior of market participants and whether there is the possibility of adverse selection on a more aggregate level, it is useful to look at the aggregate relationship between yields and takeup.

To see whether yield fluctuations are related to the aggregate demand for insurance, I regress total insurance takeup on last year's and current yields, controlling for the expected yield. The results are shown in Table 4. Columns 1-3 show the results of an OLS regression of log policies on yields, while Columns 4-6 show the corresponding results using a negative binomial specification. A 1% increase in last year's yield is associated with a $0.48 - 0.53\%$ decrease in the total number of insurance policies for corn and a $0.46 - 0.54\%$ decrease for soybeans, while a 1% increase in this year's yield is associated with a $0.43 - 0.44\%$ increase in the number of insurance policies for corn and a 0.48% increase for soybeans. The same pattern holds for past and current yields in the negative binomial specifications in Columns 3-6, although the coefficients are generally less significant.

One explanation for this result is that something correlated with current yields (e.g., the prices farmers face) also makes crop insurance more attractive to farmers. Unfortunately, because of the reduced form nature of the test and lack of individual data, it is difficult to determine the precise channel through which this type of selection is taking place. Another possibility is that it is actually insurance providers who have superior information about future yields and are able to exploit that information to increase the takeup of insurance in years when yields will be high (and thus payouts will be low).

In most other markets where there is potential for adverse selection, the unobserved or unpriced information is assumed to be known by the individual. In the case of county yields, however, it is plausible that providers have superior information, especially if they observe individual information for many other farmers in the same county and are thus better informed about the data generating process. If it is the case that providers are better informed, why don't they change their prices? First, it may be more profitable for providers to practice selection of farmers into insurance plans rather than change prices. Second, because crop insurance is reinsured by the government, providers are not as free to set their prices as they would be in a perfectly competitive market. In particular, once a price is approved by the Risk Management Agency, providers are prohibited from competing on prices. Thus, much of the competition may be taking place on the margin of who to sign up for a particular insurance plan and when.

3.6 Conclusion

Worldwide, the question of how to efficiently provide farmers with adequate protection against weather shocks remains unanswered. A way to create sustainable unsubsidized markets in crop insurance is yet to be developed. The Federal Crop Insurance Act of 1980 was a major step taken by the US government to replace a standing crop disaster assistance program for which farmers did not pay with (subsidized) crop insurance. Initial hopes of insuring 50% of eligible acreage by 1988 proved overly optimistic; only 25% was insured by this time. Participation rates remained low into the 1990's, despite increasing subsidization of premiums. Although the current participation rate is high, it comes at a high subsidy cost: the federal government pays nearly 60% of the premiums, on average. Because of the large costs of the current insurance program, there have been proposals to eliminate the crop insurance program and return to a standing disaster assistance program (Glauber, 2007).

Standard explanations for the reluctance of the private sector to enter this market include the non-idiosyncratic nature of shocks, moral hazard, and adverse selection. Although the US crop insurance market was designed with the awareness of adverse selection and moral hazard in this sector, numerous ways in which farmers can exploit the design of the system remain. In this paper, I test for the presence of adverse selection in group risk insurance plans, where farmers are paid based on yield shortfalls in the county as a whole, regardless of their own yields. I use last year's yield (which is predictive of the current yield but does not affect the price of the plan) as information on which farmers could adversely select in and out of group insurance. I also test whether the current yield (which is realized after the insurance purchase decision is made) is predictive of group insurance takeup.

I find no evidence that last year's yields influence takeup of group insurance plans. However, the reduced-form test indicates that group insurance takeup is higher when average current yields are higher. This suggests that the net selection into area yield plans favors providers, not buyers of insurance. This is consistent with earlier findings that insurance companies may practice weather-based adverse selection into reinsurance plans (Ker and McGowan, 2000). In this case, the selection may be working through insurance agents convincing farmers to choose one plan over another or through targeting particular counties in years when yields are likely to be high. Unfortunately, I cannot determine the exact mechanism with the current data. However, this is an important avenue of future research. If providers of area yield or weather-based insurance are able to better predict outcomes than the individuals they insure and are able to use non-price mechanisms to increase takeup in years when insurance is least necessary, this undermines the potential of such plans to provide a cheap and moral hazard free insurance mechanism.

Finally, I find that yields are significant predictors of total insurance demand in OLS regressions of total insurance takeup on current and past yields. This suggests that the desirability of non-group insurance plans is changing with yields as well. However, because prices in these plans are determined using individual yields, which I do not observe, I cannot determine whether the relationship between aggregate takeup of insurance and yields is due to selection or changes in prices. One of the shortcomings of this test is the lack of individual data, which would allow a similar test to be performed with individual yield and revenue plans. Because farmers may have much better information about their own potential yields than insurers do, adverse selection based on unpriced information may operate differently in individual plans than in group plans.

3.7 References

Agarwal, S., Chomsisengphet, S., and C. Liu. 2010. "The Importance of Adverse Selection in the Credit Card Market: Evidence from Randomized Trials of Credit Card Solicitations", *Journal of Money, Credit and Banking*, 42(4): 743–754.

Altman, D., Cutler, D. and R. Zeckhauser. 1998. "Adverse Selection and Adverse Retention," *American Economic Review*, 88(2): 122-126.

Babcock, B. 2011. "The Politics and Economics of the U.S. Crop Insurance Program", working paper.

Brown, J. and A. Finkelstein. 2007. "Why Is the Market for Long Term Care Insurance so Small", *Journal of Public Economics*, 91(10): 1967-1991.

Coble, K., Knight, T., Pope, R. and J. Williams. 1997. "An Expected-Indemnity Approach to the Measurement of Moral Hazard in Crop Insurance", *American Journal of Agricultural Economics*, 79(1): 216-226.

Cohen, A., and P. Siegelman. 2010. "Testing for Adverse Selection in Insurance Markets", *Journal of Risk and Insurance*, 77(1): 39–84.

Cutler, D. and S. Reber. 1998. "Paying for Health Insurance: The Trade-off between Competition and Adverse Selection," *Quarterly Journal of Economics*, 113(2): 433-466.

Deschenes, O. and M. Greenstone. 2007. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather", *American Economic Review*, 97(1): 354-385.

Dionne, G., Michaud, P. and M. Dahchour. 2004. "Separating Moral Hazard from Adverse Selection and Learning in Automobile Insurance: Longitudinal Evidence from France", Canada Research Chair in Risk Management Working Paper No. 04-05.

Einav, L., Finkelstein, A., Ryan, S., Schrimpf, P., and M. Cullen. 2011. "Selection on moral hazard in health insurance", NBER Working Paper 16969.

Finkelstein, A. and K. McGarry. 2006. "Multiple dimensions of private information: evidence from the long-term care insurance market", *American Economic Review*, 96 (4): 938-958.

Finkelstein, A. and J. Poterba. 2002. "Selection Effects in the United Kingdom Individual Annuities Market", *The Economic Journal*, 112(476): 28-50.

Finkelstein, A. and J. Poterba. 2004. "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market", *Journal of Political Economy*, 112(1): 183-208.

Fong, W. 2002. "On the Cost of Adverse Selection in Individual Annuity Markets: Evidence From Singapore", *Journal of Risk and Insurance*, 69(2): 193–208.

Garrett, T., Marsh, T., and M. Marshall. 2006. "Political allocation of US agriculture disaster payments in the 1990s", *International Review of Law and Economics*, 26 (2): 143-161.

Glauber, J. 2007. "Double Indemnity: Crop Insurance and the Failure of U.S. Agricultural Disaster Policy", Working Paper.

Glauber, J., Collins, K., and P. Barry. 2002. "Crop insurance, disaster assistance, and the role of the federal government in providing catastrophic risk protection", *Agricultural Finance Review*, 62 (2): 81-101.

Goodwin, B. 1993. "An Empirical Analysis of the Demand for Multiple Peril Crop Insurance", *American Journal of Agricultural Economics*, 75 (2): 425-434.

Gruber, J. and D. Grabowski. 2007. "Moral Hazard in Nursing Home Use", *Journal of Health Economics*, 26: 560-577.

Handel, B. 2011. "Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts", working paper.

Harri, A., Coble, K., Ker, A. and B. Goodwin. 2011. "Relaxing Heteroskedasticity Assumptions in Area-Yield Crop Insurance Rating", *American Journal of Agricultural Economics*, 93(3): 707-717.

Horowitz, J. and E. Lichtenberg. 1993. "Insurance, Moral Hazard, and Chemical Use in Agriculture", *American Journal of Agricultural Economics*, 75(4): 926-935.

Ker, A. and P. McGowan. 2000. "Weather-Based Adverse Selection and the U.S. Crop Insurance Program: The Private Insurance Company Perspective", *Journal of Agricultural and Resource Economics*, 25(2): 386-410.

Makki, S. and A. Somwaru. 2001. "Evidence of Adverse Selection in Crop Insurance Markets", *The Journal of Risk and Insurance*, 68 (4): 685-708.

McCarthy, D. and O. Mitchell. 2010. "International Adverse Selection in Life Insurance and Annuities", *International Studies in Population*, 8(2): 119-135.

Meehl, G.A., T.F. Stocker, W.D. Collins, P. Friedlingstein, A.T. Gaye, J.M. Gregory, A. Kitoh, R. Knutti, J.M. Murphy, A. Noda, S.C.B. Raper, I.G. Watterson, A.J. Weaver and Z.-C. Zhao. 2007. "Global Climate Projections." In: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Roberts, M., Key, N., and E. O'Donoghue. 2006. "Estimating the Extent of Moral Hazard in Crop Insurance Using Administrative Data", Applied Economic Perspectives and Policy, 28(3): 381-390.

Schneider, S.H., S. Semenov, A. Patwardhan, I. Burton, C.H.D. Magadza, M. Oppenheimer, A.B. Pittock, A. Rahman, J.B. Smith, A. Suarez and F. Yamin. 2007. Assessing key vulnerabilities and the risk from climate change." Climate Change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, M.L. Parry, O.F. Canziani, J.P. Palutikof, P.J. van der Linden and C.E. Hanson, Eds., Cambridge University Press, Cambridge, UK, 779-810.

Simon, K. 2005. "Adverse selection in health insurance markets? Evidence from state small-group health insurance reforms", *Journal of Public Economics*, 89(9-10): 1865-1877.

Skees, J., Black, R., and B. Barnett. 1997. "Designing and Rating an Area Yield Crop Insurance Contract", *American Journal of Agricultural Economics*, 79(2): 430-438.

Sloan, F. and E. Norton. 1997. "Adverse Selection, Bequests, Crowding Out, and Private Demand for Insurance: Evidence from the Long-term Care Insurance Market", *Journal of Risk and Uncertainty*, 15(3): 201-219.

Smith V. and B. Goodwin. 1996. "Crop Insurance, Moral Hazard, and Agricultural Chemical Use", *American Journal of Agricultural Economics*, 78(2): 428-438.

Walters, C., Shumway, R., Chouinard, H. and P. Wandschneider. 2007. "The Effects of Transitional Yields on Adverse Selection in Crop Insurance", working paper.

3.8 Appendix A

The baseline for individual revenue insurance is also based on the farmer's $4 - 10$ year yield history but takes prices into account as well. The baseline revenue is the average of the individual's historic yields multiplied by the Chicago Board of Trade pre-growing season futures prices for that crop.[10] The actual revenue is calculated using the 1-month futures price near the harvest time for that crop (called the "harvest futures price"):

[10] The exact month varies by crop.

$$\text{Pay\_RA}_{it} = \max\left(0, \text{Revenue guarantee}_{it} \; - \text{Revenue}_{it}\right) \tag{3.8}$$

$$\begin{aligned}\text{Revenue guarantee}_{it} &= X_i * P\left(pre - season\right)_t * \frac{1}{10}\sum_{s=1}^{10}\text{Yield}_{i,t-s} \\ \text{Revenue}_{it} &= \text{Yield}_{it} * P\left(harvest\right)_t\end{aligned}$$

In some plans, farmers have the option to have the revenue guarantee based on the *maximum* of the February and the harvest futures prices. This decreases the probability of a claim as well as the price of the insurance plan.

There are four different types of individual revenue plans: Crop Revenue Coverage (CRC), Income Protection (IP), Indexed Income Protection (IIP), and Revenue Assurance (RA). The characteristics of all individual revenue plans are very similar. The only difference between CRC and RA plans is that the former has a limit on the amount of payment that is made in case of a loss. The IP plan is similar to the CRC and RA plans, but there is no possibility to base the revenue guarantee on the harvest futures price. In addition, the IP requires that all cropland in a county that belongs to the same entity and grows a particular crop be insured together. This is called an "enterprise unit". CRC and RA allow the cropland to be divided into more basic units and insured separately. Starting in 2011, the four revenue protection plans have been combined into two plans to eliminate redundancies.

The IIP plan is a variation of the IP plan that is based on individual yield histories *relative* to county yields. The baseline yield is established by subtracting the average historic difference between individual and county yields from the expected county yield, which is defined as the average county yield in the previous year. In other words, the yield guarantee for year t is:

$$X_i * \left[\overline{\text{Yield}}_{c,t-1} - \frac{1}{10}\sum_{s=1}^{10}\overline{\text{Yield}}_{c,t-s} + \frac{1}{10}\sum_{s=1}^{10}\text{Yield}_{i,t-s}\right]$$

where X_i is again the coverage level chosen by the farmer.

The final category of insurance plans is group revenue (Group Risk Income Protection or GRIP) insurance that pays farmers based on a combination of county yields and futures prices. As with the individual revenue plans, prices in the county revenue plan are based on CBT futures prices for the crop. The county revenue plan payments are determined as follows:

$$\text{Pay\_GRIP}_{it} = \text{Revenue guarantee}_{ict} \; - \text{Revenue}_{ct} \tag{3.9}$$

$$\text{Revenue guarantee}_{ict} = X_i * P\left(pre - season\right)_t * \frac{1}{30}\sum_{s=1}^{30}\widehat{\text{Yield}}_{c,t-s}$$

$$\text{Revenue}_{ct} = \overline{\text{Yield}}_{ct} * P\left(harvest\right)_t$$

3.9 Appendix B

In this section, I regress the prices of GRP crop insurance on last year's, current, and expected yields. The level of observation is county-year-coverage-level. The coverage level ranges from 65% to 90% in increments of 5 percentage points.

Table A1 shows the relationship between the price of GRP policies and last year's yield, controlling for expected yield. There is no statistically significant relationship between prices and last year's yield. The expected yield is negatively and significantly correlated with the price.

Table A2 shows the relationship between the price of GRP policies and this year's yield, controlling for expected yield. Again, there is no statistically significant relationship between the current yield and the price. Moreover, the high R-squared suggests that the variables included in the regression capture nearly all the relevant variation in prices.

In Table A3, I show OLS regressions corresponding to Table 3. Although high yields last year are strongly predictive of lower yields this year, there is no significant relationship between last year's yields and the number of currently held policies (although the estimate is not very precise). Current yields are also not significant predictors of the total takeup. For both current and past yields, the sign of the estimated coefficient is the opposite of what would be predicted by theory. Changing the dependent variable to $Log\left(Policies_{lt} + 1\right)$ to take into account observations where no GRP policies are purchased does not change the results qualitatively, although it does increase their precision.

Figures

Figure 1. The relationship between expected and actual yields.

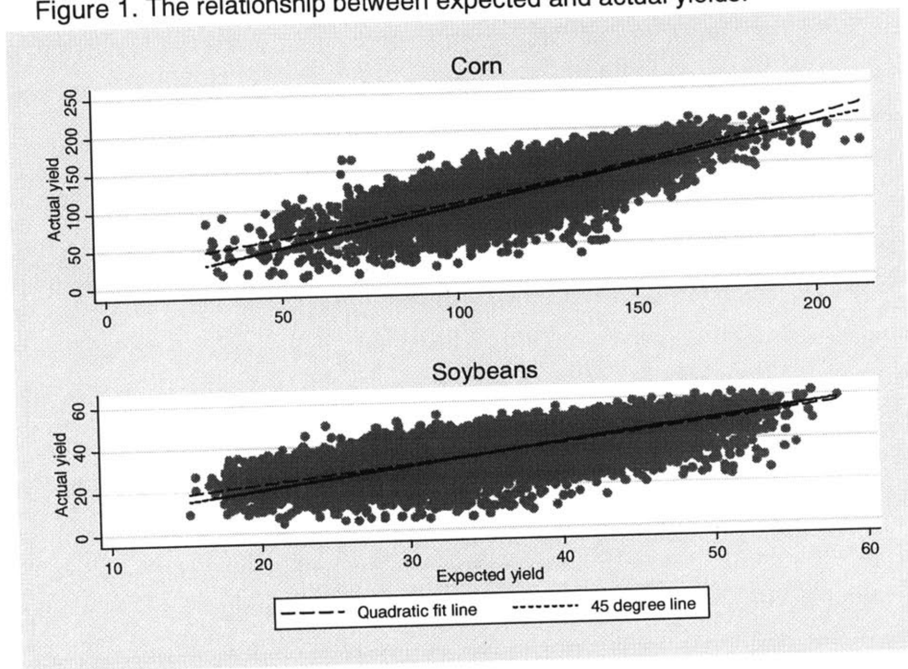

Figure 2. The relationship between last year's and actual yields.

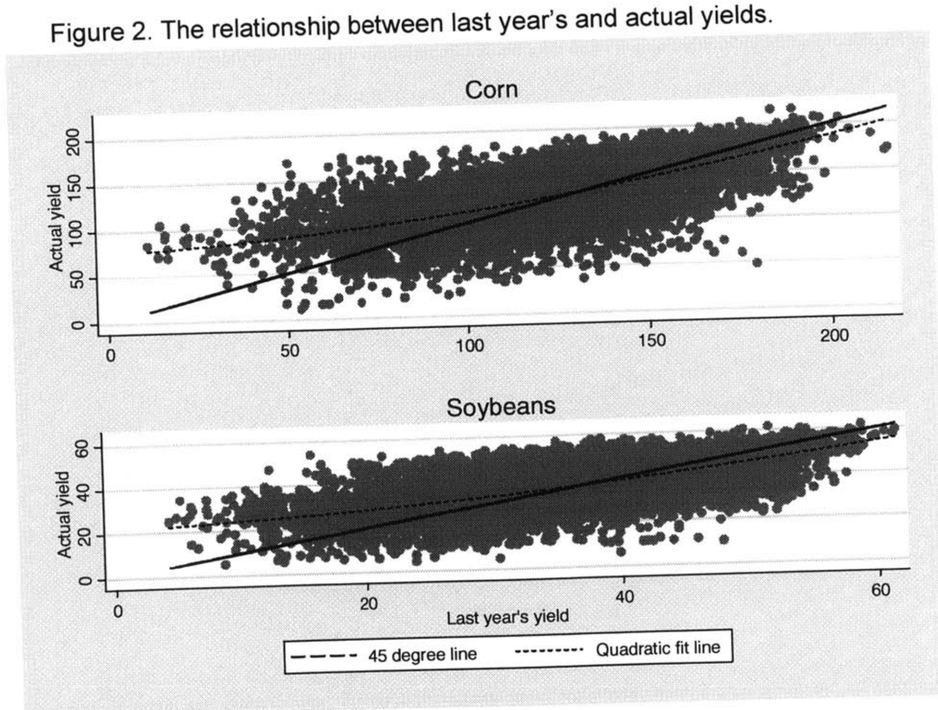

Tables

Table 1: Summary Statistics, 1994-2009

	Corn	Soybeans
Yield	128.39	37.41
	(30.52)	(9.67)
Number of insurance policies	169.35	179.17
	(265.46)	(269.09)
Number of Group Risk Plan (GRP) policies	6.60	6.89
	(14.64)	(18.73)
Share of GRP policies	0.02	0.03
	(0.05)	(0.08)
Number of counties with GRP	1,076	961

Note: standard deviations in parentheses.

Table 2. The relationship between consecutive years' yields

	Log(Yieldlt)		Yieldlt	
	All	Nonzero	All	Nonzero
Panel A: corn				
Log(Yieldl,t-1)	-0.1209	-0.0873		
	(0.0465)**	(0.0421)**		
Yieldl,t-1			-0.0672	-0.0546
			(0.0408)	(0.0454)
Log(Et-1[Yieldlt])	-0.3173	-0.3778		
	(0.1682)*	(0.1172)***		
Et-1[Yieldlt]			-0.1734	-0.2038
			(0.1129)	(0.1037)*
Mean of dep. var.	4.84	4.88	130.78	135.16
Observations	8,272	6,604	8,272	6,604
Panel B: soybeans				
Log(Yieldl,t-1)	-0.1048	-0.0860		
	(0.0379)***	(0.0384)**		
Yieldl,t-1			-0.0910	-0.0815
			(0.0366)**	(0.0401)**
Log(Et-1[Yieldlt])	-0.7145	-0.6854		
	(0.1729)***	(0.1474)***		
Et-1[Yieldlt]			-0.5858	-0.5810
			(0.1159)***	(0.1180)***
Mean of dep. var.	3.59	3.64	37.72	39.35
Observations	9,814	7,533	9,814	7,533

Note: Robust standard errors (clustered by state-year) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes year and county fixed effects.

Table 3. Past yields and current takeup, negative binomial regression

	Corn			Soybeans		
Log(YieldI,t-1)	0.386		0.378	0.015		-0.026
	(0.293)		(0.284)	(0.192)		(0.180)
Log(YieldIt)		0.739	0.753		0.438	0.492
		(0.268)**	(0.266)**		(0.221)*	(0.216)**
Log(Et-1[YieldIt])	-1.507	-1.814	-2.130	1.364	1.482	1.101
	(1.024)	(0.906)*	(0.955)**	(0.556)**	(0.527)**	(0.560)*
Mean of dep. var.	8.01	7.94	8.01	9.33	8.74	9.33
Observations	6,483	7,209	6,483	7,291	8,179	7,291

Note: Cluster bootstrap standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Dependent variable is the number of policies in county I at time t. Includes county and year fixed effects.

Table 4. Past yields and aggregate insurance demand

	Log(Policiesit)			Policiesit (negative binomial)		
	Panel A: corn					
Log(Yieldl,t-1)	-0.5341		-0.4792	-0.062		-0.055
	(0.2266)**		(0.2299)**	(0.035)		(0.037)
Log(Yieldit)		0.4441	0.4312		0.100	0.066
		(0.1993)**	(0.2090)**		(0.046)**	(0.050)
Log(Et-1[Yieldit])	-2.7741	-2.2683	-2.6207	-0.648	-0.301	-0.634
	(0.6984)***	(0.6923)***	(0.6790)***	(0.136)***	(0.153)*	(0.131)***
Mean of dep. var.	2.16	2.12	2.16	384	389	384
Observations	5,262	5,724	5,262	8,244	9,255	8,244
	Panel B: soybeans					
Log(Yieldl,t-1)	-0.5421		-0.4620	-0.039		-0.033
	(0.1606)***		(0.1578)***	(0.035)		(0.035)
Log(Yieldit)		0.4769	0.4840		0.084	0.062
		(0.1387)***	(0.1354)***		(0.034)**	(0.032)*
Log(Et-1[Yieldit])	-2.5129	-1.3108	-2.0636	0.071	0.259	0.111
	(0.6161)***	(0.6125)**	(0.6180)***	(0.152)	(0.157)	(0.158)
Mean of dep. var.	2.15	2.10	2.15	361	370	361
Observations	5,428	5,839	5,428	9,797	10,726	9,797

Note: Robust standard errors (clustered by state-year for OLS, cluster bootstrapped for binomial) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes county and year fixed effects.

Appendix Tables

Table A1. The relationship between prices and last year's yield

	Corn				Soybeans			
	Log(Pricelgt)		Pricelgt		Log(Pricelgt)		Pricelgt	
	All	Nonzero	All	Nonzero	All	Nonzero	All	Nonzero
Log(Yieldl,t-1)	0.0113	0.0196			0.0010	-0.0044		
	(0.0154)	(0.0208)			(0.0167)	(0.0225)		
Log(Et-1[Yieldlt])	-0.4062	-0.3993			-0.1797	-0.1987		
	(0.0843)***	(0.0649)***			(0.0399)***	(0.0527)***		
Yieldl,t-1			-0.0002	-0.0001			-0.0013	-0.0016
			(0.0004)	(0.0004)			(0.0012)	(0.0014)
Et-1[Yieldlt]			-0.0138	-0.0105			-0.0203	-0.0141
			(0.0027)***	(0.0012)***			(0.0038)***	(0.0031)***
Mean of dep. var.	0.78	0.66	2.80	2.38	0.45	0.34	2.03	1.74
Observations	37,567	28,925	37,567	28,925	44,947	34,061	44,947	34,061
R-squared	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99

Note: Robust standard errors (clustered by state-year) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes county-by-coverage-level and year-by-coverage-level fixed effects.

Table A2. The relationship between prices and current yields

	Corn				Soybeans			
	Log(Pricelgt)		Pricelgt		Log(Pricelgt)		Pricelgt	
	All	Nonzero	All	Nonzero	All	Nonzero	All	Nonzero
Log(Yieldlt)	0.0249	0.0307			-0.0118	-0.0254		
	(0.0163)	(0.0192)			(0.0162)	(0.0214)		
Log(Et-1[Yieldlt])	-0.3993	-0.3949			-0.1886	-0.2133		
	(0.0825)***	(0.0640)***			(0.0428)***	(0.0562)***		
Yieldlt			0.0004	0.0003			-0.0017	-0.0025
			(0.0006)	(0.0004)			(0.0011)	(0.0012)**
Et-1[Yieldlt]			-0.0136	-0.0103			-0.0204	-0.0145
			(0.0027)***	(0.0012)***			(0.0040)***	(0.0031)***
Mean of dep. var.	0.80	0.66	2.88	2.40	0.46	0.34	2.05	1.75
Observations	40,103	29,337	40,103	29,337	46,251	34,591	46,251	34,591
R-squared	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99

Note: Robust standard errors (clustered by state-year) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes county-by-coverage-level and year-by-coverage-level fixed effects.

Table A3. Past yields and current takeup, OLS

	Log(PoliciesIt)			Log(PoliciesIt+1)		
	Panel A: corn					
Log(YieldI,t-1)	-0.0966		-0.0651	-0.0784		-0.0655
	(0.1994)		(0.1952)	(0.1024)		(0.1018)
Log(YieldIt)		0.1975	0.2437		0.0966	0.1070
		(0.1694)	(0.1695)		(0.1027)	(0.1123)
Log(Et-1[YieldIt])	-2.4695	-2.1521	-2.3682	-2.0288	-1.7190	-1.9949
	(0.5608)***	(0.5507)***	(0.5579)***	(0.3531)***	(0.3129)***	(0.3515)***
Mean of dep. var.	1.71	1.71	1.71	1.08	1.05	1.08
Observations	4,550	4,937	4,550	8,272	9,303	8,272
	Panel B: soybeans					
Log(YieldI,t-1)	-0.2804		-0.2479	-0.2039		-0.1961
	(0.1372)**		(0.1370)*	(0.0882)**		(0.0876)**
Log(YieldIt)		0.2424	0.2397		0.0300	0.0746
		(0.1280)*	(0.1274)*		(0.0887)	(0.0886)
Log(Et-1[YieldIt])	-0.7349	0.0302	-0.5101	-0.7992	-0.4891	-0.7459
	(0.5501)	(0.5472)	(0.5565)	(0.3253)**	(0.3161)	(0.3194)**
Mean of dep. var.	1.81	1.77	1.81	1.00	0.98	1.00
Observations	4,785	5,163	4,785	9,814	10,749	9,814

Note: Robust standard errors (clustered by state-year) in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%. Includes county and year fixed effects.