Title
Weather, Expectations, and Complex Incentives

Author
Miller, Benjamin Michael

Publication Date
2015

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Weather, Expectations, and Complex Incentives

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Economics

by

Benjamin Michael Miller

Committee in charge:

Professor Gordon B. Dahl, Chair
Professor Jennifer Burney
Professor Julie Berry Cullen
Professor Craig McIntosh
Professor Paul Niehaus

2015

The dissertation of Benjamin Michael Miller is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

____

____

____

____

____

Chair

University of California, San Diego

2015

DEDICATION

To my family, friends, and mentors for supporting me during this process. A special thanks to my parents, Ken and Bethia Miller, for their advice, empathy, and unwavering support. You are the best role models I've ever met. Thanks to my wife, Kristina Thorsell, for being my teammate through too many years of a coast-to-coast relationship. And thanks to my brother, Nathan Miller, for always being on the same page as me.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

First and foremost, much gratitude is due to my dissertation committee members: Gordon Dahl, Jen Burney, Julie Cullen, Craig McIntosh, and Paul Niehaus for invaluable guidance. Special thanks are also due to Kevin Mumford and the late Mike Watts, who have each been a wonderful combination of coauthor, friend, and mentor.

I would also like to thank the participants at the following seminars and conferences for comments on this research: American Meteorological Society 95th Annual Meeting's Third Symposium on Building a Weather Ready Nation, Belpasso International Summer School on Environmental and Resource Economics, Bureau of Economic Analysis, Bureau of Labor Statistics, Cornerstone Research, East Carolina University, Economic Research Service (USDA), European Association of Agricultural Economists' PhD Workshop, Federal Trade Commission, National Tax Association Annual Meeting, Pacific Conference for Development Economics, Purdue University, RAND, Studies of Precipitation, flooding, and Rainfall Extremes Across Disciplines (SPREAD) Workshop, U.S. Census Bureau, U.S. Treasury, University of California San Diego, University of New Hampshire, University of South Carolina, Western Economics Association International Graduate Student Dissertation Workshop, and World Weather Open Science Conference.

Additionally, I would like to thank anonymous reviewers and the following people to whom I am indebted for their insight and various assistance above and beyond the call of duty: Prashant Bharadwaj, Graton Gathright, Laura Gee, Marina Kutyavina, and Sam Young. Finally, this acknowledgment would be woefully incomplete without thanking Rafael Acevedo and Suzi Harlow for their constant support.

Chapter 1, in part is currently being prepared for submission for publication of the material. Miller, Benjamin M. The dissertation author was the primary investigator and author of this material.

Chapter 2, in full, has been submitted for publication of the material as it may appear in the National Tax Journal, 2015, Miller, Benjamin M.; Mumford, Kevin. The dissertation author was a primary investigator and author of this paper.

Chapter 3, in part is currently being prepared for submission for publication of the material. Miller, Benjamin M. The dissertation author was the primary investigator and author of this material.

VITA

2010	B.A. in Economics Honors with Distinction & University Honors Program Honors Diploma, Purdue University
2011-2012	Statistician, U.S. Census Bureau, Social, Economic, and Housing Statistics Division
2015	Ph.D. in Economics, University of California, San Diego

PUBLICATIONS

Miller, Watts "Oh, the Economics You'll Find in Dr. Seuss!", *The Journal of Economic Education*, 2011

ABSTRACT OF THE DISSERTATION

Weather, Expectations, and Complex Incentives

by

Benjamin Michael Miller

Doctor of Philosophy in Economics

University of California, San Diego, 2015

Professor Gordon B. Dahl, Chair

Uncertain outcomes or complex incentives can cause individuals to form and respond to expectations. The first chapter examines the formation of expectations over future rainfall among Indian farmers, and the implications of anticipatory adaptation on standard identification strategies using rainfall. The second chapter examines how complexity caused U.S. taxpayers to responded to a naive proxy for changes in the value of the Child and Dependent Care Credit rather than a more nuanced measure. The third chapter the impact of information provision on

the prevention of tornado fatalities and injuries, through the expansion of radio transmitters for the National Oceanic and Atmospheric Administration's Weather Radio All Hazards.

Chapter 1

Does Validity Fall from the Sky? Observant Farmers and the Endogeneity of Rainfall

1.1 Introduction

Weather, particularly rainfall, is a popular source of identifying variation in many areas of empirical economics. With few exceptions, researchers have treated deviations from mean rainfall as an exogenous and unpredictable shock. Yet the prevalence of 10-day weather forecasts and longer-range seasonal forecasts suggest plenty of information about future weather events is available. Once information about future rainfall deviations becomes available, rainfall's impacts on income, prices, health, and other factors also become predictable. When possible, utility maximizing agents should use this predictability to adapt to rainfall deviations before they occur in order to optimize future consumption, production, and other outcomes. Anticipatory adaptation makes the eventual impact of rainfall deviations on consumption and other outcomes of interest an endogenous functions of agents' behavioral choices.

The use of rainfall deviations as a source of exogenous and unpredictable shocks is neither new nor outdated. Explicit description of rainfall as "exogenous" goes at least as far back as Koopmans (1949). At present rainfall is commonly used as an exogenous shock to household consumption, aggregate consumption, and income. That said, other concerns about rainfall-based identification strategies are not new. Rosenzweig and Wolpin (2000) criticize treating rainfall as a exogenous shock to income through crop yields, citing evidence that rainfall also alters relative prices. Kochar (1999) and Rose (2001) discuss rainfall impacting the productivity of own-farm production relative to labor market production. Rainfall may have other direct effects as well. Residual soil moisture may impact future yields, or standing water may spread disease. Because these types of direct effects violate the exclusion restriction necessary in an instrumental variables setting, recent papers more commonly pursue reduced form estimation.

The main contribution of this paper is to empirically document that rainfall deviations are not an unexpected shock at the time they occur, as evidenced by farmers adapting their crop choices at the beginning of the agricultural season in advance of seasonal weather realizations. The resulting endogeneity of income provides another reason why the exclusion restriction for IV is unlikely hold. This also means reduced form estimates using rainfall deviations require further reinterpretation. While the extent to which adaptation alters estimates depends on the outcome of interest, a variety of cases are examined to show the role of adaptation is far from trivial. In some cases, removing as much adaptive behavior as possible causes standard estimates of income elasticities to be cut in half.

If economists could control for agents' expectations at a given period of time, then it would be possible to remove the endogenous impacts of earlier adaptation by using rainfall which was unexpected at the given period. Subsequent changes in expectations were by definition unexpected at this given period. Adaptations driven by these updates to expectations were also unexpected and exogenous at this period. Deviations from expected rainfall should be interpreted as an unexpected shock which occurs not at the time of precipitation, but before more nuanced expectations are formed. After this point, the rainfall deviation is no longer unexpected. Reduced form estimates should be interpreted as reflecting not only the impact of the physical rainfall event through income and relative prices, but also all impacts of anticipatory adaptation. Deviations from mean rainfall should then be interpreted as a shock which occurs before agents are able to form expectations more nuanced than the mean. This paper empirically shows that at least as early as the start of the agricultural season, farmers have accurate rainfall expectations more nuanced than a long-run mean.

India is a popular location among papers using rainfall for identifying variation. This paper examines multiple data sets on Indian agriculture which are popular sources of data both in general and among papers using rainfall for identifying variation.[1] The crop selections of Indian farmers are found to be strongly correlated with the season's upcoming rainfall in an agronomically efficient manner. In years where rainfall is one standard deviation below the mean, the district-wide acreage of sorghum (a relatively drought resistant crop) increases by almost 3%, while the district-wide acreage of rice (a relatively water-intensive crops) decreases by over 1%. The response of average farming households from a popular survey of rural villages are larger than the average response of aggregate acreage. These

[1]Prominent examples using the ICRISAT data examined here include Jacoby and Skoufias (1998), Kochar (1999), Mazzocco and Saini (2012), and Rosenzweig and Udry (2013). Jayachandran (2006) uses IAC data. Other papers using Indian rainfall for identifying variation include Wolpin (1982), Rose (1999), Rose (2001), and Taraz (2012).

farmers alter acreage devoted to drought-resistant and water-intensive crops by as much as 10% per crop, consistent with poorer, smaller-scale farmers being more risk averse. The temporal order for these correlations is incorrect for identifying a causal relationship. Farmers must make decisions about how much acreage to sow before observing rainfall outcomes. The response of crop acreage to the upcoming season's rainfall is not a delayed response to early-season rain, nor is it a response to medium-run regime shifts in average rainfall as suggested in Taraz (2012). Such apparent anticipation of weather variation should not be possible when the variation is unpredictable at the time of planting.

This paper suggests farmers observe signals about future rainfall outcomes and respond appropriately. While the idea of short-run weather expectations influencing crop choices has been occasionally speculated, this paper's broad-based empirical evidence of such adaptation is new. Indeed, repeated adaptation to annual expectations offers an important mechanism through which cropping patterns may adapt to long-run climate change. Beyond agriculture, such predictability would be of obvious value in any situation where returns on investments depend in part on weather outcomes, including fishing, tourism, and retail. When adaptive behaviors are available, agents should use available information to form and respond to expectations over future weather events in economically meaningful ways. Concurrent and complementary research by Rosenzweig and Udry (2013) finds early-season investments respond to seasonal government forecasts, and respond more strongly in areas where forecasts are more accurate. Rosenzweig and Udry (2013) focus on the returns to forecast accuracy, while this paper focuses on how the predictability of rainfall deviations impacts the interpretation of rainfall-based identification and whether the econometrician's information set can serve as a sufficient statistic for expectations.

As noted above, controlling for agents' expectations at a given time would make it possible to remove some of the endogenous impacts of rainfall-specific adaptation. The response of aggregate crop acreage within a district appears to be entirely driven by rainfall variation which is predictable at the time of planting given the econometrician's information set. This supports the interpretation of crop selection as anticipatory adaptation, and suggests researchers may be able to isolate unexpected rainfall shocks. However, uncertainty over the information and adaptations available to the agent makes isolating an unexpected rainfall shock rather complex. First, the agent may engage in adaptation using important location-specific information which is unobserved by the econometrician. Crop selections of individual farmers exhibit strong anticipatory adaptation to future demeaned rainfall, and continue to anticipate rainfall variation which was unpredicted by the econometrician. Second, failure to find evidence of adaptation in one behavior does not imply the rainfall deviation was unexpected, or that there is no adaptation in other unobserved behaviors. While the acreage individual farmers devote to sorghum and rice does not respond to unpredicted rainfall variation, the acreages of many other crops do respond to unpredicted rainfall variation. Nevertheless, this paper examines the extent to which removing as much adaptive behavior as possible alters popular estimates.

Taken together, these results make it clear that the ability to predict rainfall drives a much broader set of adaptive behaviors than would be observed if farmers were simply responding to the shocks that occur during the course of the planting season. These adaptive behaviors alter income and other outcomes which result from weather events. Anticipatory adaptation further limits the extend to which reduced form estimates are informative about more general responses to unexpected income shocks. Of first order concern is that adaptations may be specific to rainfall. Further, if individuals non-randomly face heterogeneous restrictions in their ability to engage in adaptive behavior, the resulting incomes and other outcomes influenced by rainfall are not randomly assigned. Even outcomes between agents with homogeneous choice sets of adaptive behavior will suffer from selection bias if different types of agents choose to engage in different adaptive behaviors. These issues may make the impacts of rainfall shocks difficult to compare across regions. They also mean reduced form estimates may in part reflect capital constraints, information constraints, or even risk preferences.

The rest of the paper is organized as follows. Section 2 discusses the standard use of rainfall as a source of identifying variation, presents the two-period optimization model, and discusses how information availability impacts empirical estimation. Section 3 covers background information on seasonal weather forecasts, agriculture in India, and the broader literature on crop selection. Section 4 discusses empirical evidence of anticipatory adaptation through crop choices. Section 5 investigates issues resulting from agents observing more local information than the econometrician. Section 6 compares traditional rainfall-based estimates to those using the methodology described in this paper. Section 7 concludes.

1.2 Modeling Rainfall Expectations

1.2.1 How Economists Think About Rainfall

Ever in search of exogenous variation for identification, economists have had a long and flirtatious relationship with rainfall. Explicit description of rainfall as "exogenous" goes at least as far back as Koopmans (1949). Even earlier, Working (1927) suggests tracing a demand curve for agricultural commodities by using weather as a factor which shifts supply more than demand. Rainfall has been employed in many high-quality papers of enduring popularity.[2] Today rainfall continues to a be popular source of identifying variation. Table 1.1 presents a comprehensive review of all papers published in ten top economics journals over

[2]A selection of popular papers which instrument a form of income with a form of rainfall include Wolpin (1982), Paxson (1992),Jacoby and Skoufias (1997), Jacoby and Skoufias (1998), Kochar (1999), Jensen (2000), Miguel et al. (2004), Newhouse (2005), Jayachandran (2006), Yang and Choi (2007), and Hidalgo et al. (2010). For a few popular examples of papers using variation for direct or reduced-form variation, see Roll (1984), Cramer et al. (1997), Rose (1999), Miguel (2005), Conlin et al. (2007), Connolly (2008), and Maccini and Yang (2009). Even the number of overview articles is large and growing quickly. For examples, see Rosenzweig and Wolpin (2000), Auffhammer et al. (2013), and Dell et al. (2014)

three recent years which use rainfall-driven variation of annual or finer frequency for identification of main results.[3] Table 1.1 finds 19 such papers, which span a wide variety of topics and are generally well-cited for young papers.

Why is rainfall such a popular source of identifying variation? This is perhaps due to the perception of weather outcomes as random i.i.d. draws which cannot be altered by human behavior. Sovey and Green (2011) state the reason using rainfall as an instrument is intuitively appealing is that we think of rainfall as patternless." As economists, we often want to know how a shock to income or more aggregate production measures such as GDP impacts some future economic behavior or outcome. Rainfall appears to provide a convenient source of variation which impacts agricultural productivity but is exogenous with respect to agents' behavior and unobservables. Rainfall is most commonly treated as a shock to various measures of consumption or production. Despite concerns with changes in relative prices, even reduced form impacts of rainfall are often explicit that rainfall is a proxy for income variation. For example, Björkman-Nyqvist (2013) states "In an ideal setting, I would use rainfall as an instrument for household income in a first-stage regression and income as a determinant of investment in education in a second-stage regression. Unfortunately, district-specific income data over time are not available and I will therefore study the reduced form relationship between rainfall shocks and investment in boys' and girls' education." Rainfall is also used for other purposes, such as a shock to political behavior or variation in

[3] Papers which use rainfall only as a control or only in robustness checks are not included.

the desirability of attending an event.

Criticism over the use of rainfall for identification in specific cases is not new. Trombley (1997), Ciccone (2011), and Sovey and Green (2011) all worry about spurious relations resulting in incorrect estimates. As discussed above, Rosenzweig and Wolpin (2000) criticize rainfall's use in estimating income effects, citing evidence that rainfall also alters relative prices. Rainfall may have other direct effects as well. Residual soil moisture may impact future yields, or standing water may spread disease. These types of direct effects violate the exclusion restriction necessary in an instrumental variables setting, and Table 1.1 shows recent papers more commonly pursue reduced form estimation. While year-by-region fixed effects would also capture average expectations for that year and region, these are rarely if ever used. In most cases, they would eliminate all identifying variation and are hence infeasible. Broad geographic fixed effects or year fixed effects over many regions will also not capture important local variation in expectations. The contribution of this paper is to document adaptation due to expectations significantly alters the interpretation of reduced form estimates in a way which makes them much less informative about exogenous shocks to income. The reason is that rainfall is not an unpredictable shock, as evidenced by farmers adapting their crop choice in advance of weather realizations.

In economists' standard conceptual model, "climate" refers to the known distribution of rainfall and "weather" refers to the unpredictable draw from that distribution. Rainfall in a given time period, R_t, is drawn from an arbitrary distribution with finite mean and variance, such that $R_t = \overline{R} + \xi_t$ and $\xi_t \sim (0, \sigma^2)$. When ξ_t is viewed as a random i.i.d. draw, it seems like a reasonable candidate for satisfying the weak exogeneity and Granger non-causality conditions necessary for strong exogeneity. Meeting Granger non-causality is generally regarded as rather intuitive and will not be discussed in this paper.[4] Weak Exogeneity states the production function of rainfall does not need to be known in order to estimate unbaised parameters on ξ_t. As documented in Table 1.1, research papers commonly cite exogeneity or unpredictability to justify their identification strategy. Assumption 1 is the standard assumption providing sufficient conditions for weak exogeneity.

Assumption 1 (Unpredictability)**.** *$\overline{R}$ and the distribution of ξ_t may be known, but ξ_t is an unpredictable i.i.d draw from an underlying climatic distribution.*

Usually this assumption is implicit within the identification strategy, although a few papers explicitly model rainfall as unpredictable.[5] Rosenzweig and Wolpin (2000) models rainfall as follows. Weather is random and iid over time. A weather realization is drawn each period independently from a known distribution

[4]Successful attempts to manipulate weather would obviously violate strong exogeneity. At present this risk to identification seems small. Hail cannons, while popular in areas, have not been proven statistically effective. Evidence on the effectiveness of cloud seeding is mixed at best.

[5]One rare exception is Kochar (1999), who explicitly interacts rainfall with crop acreage to remove the expected portion of income shocks. Conversely, papers such as Jacoby and Skoufias (1997) interact ξ_t with farm characteristic to identify "unanticipated income changes". A very small number of paper including Roll (1984) and Rosenzweig and Udry (2013) examine the impact of predictable rather than unpredictable rainfall variation.

with finite moments. Similarly, Rose (2001) states "In period 1, the household does not know what the value of [rainfall] will be, but it knows its distribution. It knows the average over time (μ), and it knows the variability of the distribution." Even in examining daily rainfall, Connolly (2008) is clear that "the model tested here is not interested in the effect of the climate on time allocation but rather on the impact of an exogenous weather shock, which cannot be predicted." Brückner and Gradstein (2013) are clear that they "use detailed year-to-year rainfall data as a transitory, unanticipated, and exogenous shock." Adhvaryu et al. (2013) are also particularly explicit that rainfall shocks are ideal for a number of reasons:... (2) they are unpredictable in nature and therefore not likely to induce anticipatory smoothing of employment."

Occasionally concerns are voiced about serial correlation, which would imply predictability. Paxson (1992) explicitly tests for serial correlation in rainfall, and is indeed "unable to reject the hypothesis that rainfall follows a white-noise process." Newhouse (2005) includes levels of rainfall in subsequent years as a control for serial correlation. The extent of serial correlation in rainfall data may depend on location and time period.Testing for serial correlation in rainfall data is simple, although Table 1.1 documents that it is rarely done. When used, such tests or controls for serial correlation are often touted as tests of predictability.[6]

[6] "To construct a transitory rainfall variable, it should also be known how current rainfall deviates from its expected value, $\overline{R}_{dt}$. If rainfall were serially correlated across years, one would have to forecast the expected value of rainfall for each region in each year. However, rainfall does not appear to be serially correlated: I am unable to reject the hypothesis that rainfall follows a white-noise process. Thus, I can set $\overline{R}_{dt} = \overline{R}_d$, historical rainfall over time in district d." (Björkman-Nyqvist, 2013)

While a lack of serial correlation is a necessary for being unpredictable by agents with memory, it is not a sufficient condition for unpredictability.

There is little consensus on the proper functional form of rainfall to use as a shock. Admittedly, if ξ_t is truly exogenous with respect to behavior, then transformations of ξ_t should also be exogenous with respect to behavior. Given this, some authors explicitly select a functional form with desirable qualities, such as a strong correlation with income.[7] Others choose a functional form which has been used by past researchers. There is certainly some risk that authors may try several functional forms and return the one that yields the most preferred results, in line with concerns noted by Brodeur et al. (2013). As a source of transitory shocks, Table 1.1 shows various functions of total rainfall over a fixed period are most common, such as demeaned annual rainfall or the natural log of rainfall. Indicators for amounts of rainfall above or below an arbitrary cut-off are also prevalent, such as indicators for a given year's rainfall being drawn from the tail of the rainfall distribution or number of days when more than a set amount of rainfall occurs. Other functions such as first-differenced rainfall appear occasionally, although without further controls, these functions draw identification from one or both of ξ_t and ξ_{t-1}.[8] Outside of transitory shocks, some papers use

[7] Jensen (2000), Miguel et al. (2004), and Hidalgo et al. (2010) are all explicit that they tested a variety of functional forms for rainfall, and ultimately chose the one most strongly correlated with their measure of income.

[8] For a few examples, Rose (1999) uses "deviation of rainfall from its 21-year mean for each district" as ξ_t. Paxson (1992) uses "deviations from average values of regional rainfall in each of four seasons (plus deviations from averages squared)". Maccini and Yang (2009) use $\ln(R_t) - \ln(\overline{R})$, and like many papers they exclude year t in calculating the mean rainfall. Grimard and Hamilton (1999) and Jensen (2000) interact an indicator for observations greater than one

regional or even temporal variance in $\overline{R}$ as a source of variation in permanent income or agricultural fertility.[9] This paper focuses on identification from ξ_t due to its prevalence as a source of transitory shocks and the concerns presented below about how the failure of Assumption 1 significantly alters how economists think about rainfall-based identification strategies.

1.2.2 A Simple Model

The following two-period model isolates the concerns presented in this paper.[10] A utility-maximizing agent optimizes by engaging in various economic behaviors. Let B^i represent a single economic behavior over which the agent may optimize, such as the hours invested in a particular type of labor or the amount of land devoted to a particular crop. All behavioral choices made by the agent in a discrete time period, $t = \{1, 2\}$, are represented by the vector $\mathbf{B}_t = (B_t^1, B_t^2, \ldots)$. The aggregate utility function, $U(\mathbf{B}_1, \mathbf{B}_2)$, may be such that agents are required to obtain some minimum level of utility in each period. The time periods in mind here are agricultural seasons or years, although t could represent any duration of time.

---

standard deviation from mean rainfall with the absolute value of the deviation from mean rainfall. Miguel et al. (2004) examine economic growth, and choose their shock to be "the proportional change in rainfall from the previous year, $(R_{it} - R_{i,t-1})/R_{i,t-1}$."

[9]For examples, see Wolpin (1982), Hornbeck (2012), Burke and Emerick (2013), and others.

[10]This model could easily incorporate additional generalizations which are not necessary for this setting. I_1 could represent a vector of constraints, such as time and health constraints. Here ξ is treated as a scalar measure of a rainfall outcome which occurs between periods, but it could easily represent a more general vector of weather outcomes and external shocks. $\mathbf{B}_1$ could include borrowing or savings behaviors, altering future resources and current costs. The entire optimization problem could be extended over a larger or infinite number of periods.

Behavioral choices face an income constraint, such that the costs of behavior this period must be less than income earned in the prior period. Income earned in the first period is a function of first period behavior and a mean-zero rainfall outcome, ξ, such that $I_1 = F(\mathbf{B}_1, \xi)$. Initial income is normalized such that $I_0 = 1$. Following commonly cited examples discussed below, the cost of behavior is permitted to be a function of current and past behavior choices. The cost function for behavior is expressed as $c_1(\mathbf{B}_1)$ and $c_2(\mathbf{B}_1, \mathbf{B}_2)$ for periods one and two, respectively. In the second period, the agent chooses a utility-maximizing vector of behaviors, $\mathbf{B}_2$, as a function of past behaviors, $\mathbf{B}_1$, and the rainfall shock, ξ. Let the value function, $V(\mathbf{B}_1, \xi)$, represent the optimal obtainable utility given $\mathbf{B}_1$ and ξ, which are both known to the agent in the second period.

$$V(\mathbf{B}_1, \xi) \equiv \max_{\mathbf{B}_2} \{U(\mathbf{B}_1, \mathbf{B}_2) \text{ s.t. } c_2(\mathbf{B_1}, \mathbf{B}_2) \leq F(\mathbf{B}_1, \xi)\} \tag{1.1}$$

The agent does not know the amount of rainfall, ξ, when choosing first period behavior, $\mathbf{B}_1$. However, the agent does observe a vector of signals, $\mathbf{X}$, which may contain information about the value of ξ. In the first period, the agent chooses the vector of behaviors, $\mathbf{B}_1$ which is expected to maximize aggregate utility, $U(\mathbf{B}_1, \mathbf{B}_2)$.

$$\max_{\mathbf{B}_1} \{E_{\mathbf{X}}\left[V(\mathbf{B}_1, \xi)\right] \text{ s.t. } c_1(\mathbf{B}_1) \leq 1\} \tag{1.2}$$

It will be convenient to decompose the rainfall shock as the sum of expected and unexpected components, conditional on $\mathbf{X}$.

$$\xi = E[\xi|\mathbf{X}] + \tilde{\xi} \tag{1.3}$$

Where $E[\xi|\mathbf{X}] \neq 0$, this represents a violation of Assumption 1. $\tilde{\xi}$ represents the component of rainfall which is unexpected conditional on $\mathbf{X}$. Notice that in this model rainfall only directly impacts utility through altering income. This is chosen not because rainfall is unlikely to have non-income impacts, but to highlight the extent to which using ξ for identifying variation alters interpretation *even when rainfall only directly impacts income.* This highlighted concern is new to the literature.

Without further assumptions, optimal behavior in the first period is a function of other first period behavior and expectations over rainfall. A risk-neutral agent optimizes based on the expected value of ξ conditional on available information, such that optimal behavior is defined by $B_1^{i*}(\mathbf{B}_1^{-i}, E[\xi|\mathbf{X}])$ where $\mathbf{B}_t^{-i} = (\ldots, B_t^{i-2}, B_t^{i-1}, B_t^{i+1}, B_t^{i+2})$. A risk averse agent's optimal behavior may also consider further details about the distributions of $E[\xi|\mathbf{X}]$ and $\tilde{\xi}$, such as their variance. In the second period, optimal behavior is a function of all other behavior and second-period income, or $B_2^{i*}(I_1, \mathbf{B}_2^{-i}, \mathbf{B}_1)$.

1.2.3 What do rainfall estimates actually identify?

As discussed above, economists would like to use transitory rainfall events, ξ, to estimate how unexpected income shocks alter behavior, or $\frac{\partial B_2^i}{\partial I_1}$. Due to worries that ξ directly impacts outcomes other than income, many researchers now examine the reduced form impacts. For example, Maccini and Yang (2009) speculate about impacts to both income and the relative cost of food consumption. Yet even in this model where there are no other direct impacts, ξ does not provide an unexpected shock. So what is actually identifying when using rainfall for identifying variation? Consider the reduced form impact of a predicted increase in rainfall on income $\frac{dI_1}{dE[\xi|\mathbf{X}]}$.

$$\frac{dI_1}{dE[\xi|\mathbf{X}]} = \frac{\partial I_1}{\partial \xi} + \sum_i \left(\frac{\partial I_1}{\partial B_1^i} \frac{\partial B_1^i}{\partial E[\xi|\mathbf{X}]} \right) \tag{1.4}$$

The first portion of Equation 1.4, $\frac{\partial I_1}{\partial \xi}$, is the standard the impact researchers are interested in estimating. This term reflects the reduced form impact of one more unit of rain on income, because $\frac{d\xi}{dE[\xi|\mathbf{X}]} = 1$.

The second portion of Equation 1.4, $\sum_i \left(\frac{\partial I_1}{\partial B_1^i} \frac{\partial B_1^i}{\partial E[\xi|\mathbf{X}]} \right)$, refers to changes in income caused by first-period behavior re-optimizing over changes in expected future rainfall. This paper presents evidence of such behavior in farmers planting different crops to maximize future income when drought is expected. This portion of the income change is not a shock. Agents expected this income change, and indeed altered their first period behavior in order to obtain it. Hence this portion of the change in income is an endogenous function of behaviors chosen by the agent.

Now consider a reduced form estimation, such as the impact of rainfall variation on future migration outcomes, B_2^M. As shown in the first two expressions of Equation 1.5 below, the same exogenous and endogenous income effects will be captured in the reduced form estimate of rainfall on migration. In addition, adaptive behaviors may directly impact migration outcomes through non-income channels. For example, migration is a high cost and high return investment. If credit-constrained households expecting a positive income shocks are more likely to send migrants, then the learning about travel and migration opportunities which occurs may reduce the costs of future migration. The third expression of Equation 1.5, $\sum_i \left(\frac{B_2^M}{\partial B_1^i}\frac{\partial B_1^i}{\partial E[\xi|\mathbf{X}]}\right)$, shows this non-income impact of predictable rainfall variation captured by reduced-form estimates.

$$\frac{dB_2^M}{dE[\xi|\mathbf{X}]} = \frac{B_2^M}{\partial I_1}\frac{\partial I_1}{\partial \xi} + \sum_i \left(\frac{B_2^M}{\partial I_1}\frac{\partial I_1}{\partial B_1^i}\frac{\partial B_1^i}{\partial E[\xi|\mathbf{X}]}\right) + \sum_i \left(\frac{B_2^M}{\partial B_1^i}\frac{\partial B_1^i}{\partial E[\xi|\mathbf{X}]}\right) \tag{1.5}$$

There are an unknown number of behaviors through which agents might adjust in a way which alters income or other outcomes. As shown repeatedly in results below, one behavior not appearing to respond to rainfall variation does not imply other behaviors also do not respond. To be certain the desired estimate has been isolated, econometricians face the difficult task of proving they have either controlled for every behavior adaption, or that they have perfectly controlled for rainfall expectations. Proving that either has been done successfully seems infeasible.

Under what conditions might rainfall variation capture a purely unexpected change in income? Assumption 1 tells us that there are no informative signals of rainfall outcomes, or $E[\xi|\mathbf{X}] = 0$ for any $\mathbf{X}$. If Assumption 1 does not hold, as this paper suggests, the same estimate can be obtained if $\frac{\partial B_1^i}{\partial E[\xi|\mathbf{X}]} = 0 \ \forall i$. The income function being separable in rainfall is a sufficient alternative assumption to obtain this result, because this ensures optimal first-period behavior is not a function of rainfall expectations.

Assumption 2 (Separability)**.** *The income production function is such that* $\arg\max_{\mathbf{B}_1} F(\mathbf{B}_1, \xi)$ *is not a function of* ξ*.*

Whether income-maximizing behavior is independent of rainfall is a description of the agent's ability to adapt to changes in weather expectations. Assumption 2 is quite a high bar, because it requires not one but all ex ante adaptations, such as changes in crop acreage, to remain unchanged in the face of rainfall information. In many cases, it should be expected that rational economic agents whose utility depends on future weather outcomes do respond to weather expectations in economically meaningful ways, violating Assumption 2. Further, the absence of adaptation does not mean the income shock is unexpected. Predictability alone could alter estimates, as Jappelli and Pistaferri (2010) shows that consumption responses to expected income changes are different from unexpected shocks.

Assumption 1 and Assumption 2 are jointly testable, in that if behavior responds to rainfall variation, this implies rainfall is at least partially predictable and behavior adaptations are optimal. Assumption 1 and Assumption 2 would both be rejected, and ξ would not return the desired impact of transitory, unpredicted shocks to income, prices, or anything else.

Suppose $\tilde{\xi}$ could be measured and used in place of ξ. Because $\tilde{\xi}$ is an unexpected rainfall shock, it is orthogonal to first period behavior, or $\frac{dB_1^i}{d\tilde{\xi}} = 0 \; \forall i$, eliminating concerns about adaptive behavior. [11] Further, because $\tilde{\xi}$ is by definition uncorrelated with $E[\xi|X]$, by Equation 1.3 we know $\frac{d\xi}{d\tilde{\xi}} = 1$. So long as $\tilde{\xi}$ only impacts only impacts I_1 through ξ, this estimate isolates the desired estimate of the response of an economic behavior to an unexpected, transitory income shock, $\frac{dI_1}{d\tilde{\xi}} = \frac{\partial I_1}{\partial \xi}$.

Below, the difficulty of obtaining estimates of $\tilde{\xi}$ is discussed. As a reminder, this model has explicitly assumed neither ξ nor $\tilde{\xi}$ have non-income impacts on behavior or other outcomes of interest. If $\tilde{\xi}$ were obtainable, one would still need to worry about a variety of potential exclusion restriction violations, such as rainfall directly impacts health outcomes in rural areas through supporting the spread of malaria and other diseases. As discussed later in the paper, even admirably creative attempts to estimate $\tilde{\xi}$ are likely to encounter prohibitory data limitations.

1.2.4 Attempting to Estimate $\tilde{\xi}$

Imagine a panel of I agents each receive a set of signals, $\mathbf{X}_{i,t}$, during each of T periods. The values of the signals are known and observed for each individual, and $\mathbf{X}$ represents the full set of signals received by all agents over all time peri-

[11] Examples of such papers are rare given the ubiquity of Assumption 1. One example is Kochar (1999), who explicitly interacts rainfall with crop acreage to remove the expected portion of income shocks. Other papers, including Roll (1984) and Rosenzweig and Udry (2013) examine the impact of predictable rather than unpredictable rainfall variation.

ods. Assume agents correctly understand correlations between signals and future rainfall, and that these correlations are identical across individuals. Issues such as changes in correlations over time can be captured by appropriate choice of functional forms within $\mathbf{X}_{i,t}$. In this admittedly unrealistic case, a large enough sample would enable estimates of the correlations between observed variables and future rainfall, and hence enable estimates of predicted and unpredicted rainfall for each individual.

$$\xi_{i,t} = \beta \mathbf{X}_{i,t} + \tilde{\xi}_{i,t} \tag{1.6}$$

The fitted value $\hat{\beta}\mathbf{X}_{i,t}$ is the ideal prediction of deviation from mean rainfall given $\mathbf{X}_{i,t}$, and the residual converges to unpredicted rainfall, $\tilde{\xi}_{i,t}$.

Of course, econometricians do not face such an ideal environment. The main issue is that the econometrician does not observe $\mathbf{X}$. Instead, the econometrician may observe some subset of the signals $\mathbf{X}^{\mathrm{II}} \subseteq \mathbf{X}$, as well as additional signals, $\mathbf{X}^{\mathrm{III}}$, which are unobserved to the agents. How might econometricians use their set of signals, $\mathbf{W} = \{\mathbf{X}^{\mathrm{II}}, \mathbf{X}^{\mathrm{III}}\}$? One approach would be to estimate expectations using only information available to both the individual and the econometrician ($\mathbf{X}^{\mathrm{II}}$). There are several issues with this. First, definitively determining whether an econometrician's variable falls in $\mathbf{X}^{\mathrm{II}}$ or $\mathbf{X}^{\mathrm{III}}$ is difficult. Which variables are available to and used by individuals to form weather expectations remains a widely unanswered research question which this paper does not directly address. That issue aside, let $\mathbf{X}^{\mathrm{I}} = \mathbf{X} \backslash \mathbf{W}$ represent the set of signals available to the agent and not the econometrician. If there are variables in $\mathbf{X}^{\mathrm{I}}$ which can be used to form more accurate expectations, then individual are still able for form expectations over residual unpredicted rainfall, when predictions are estimated using only $\mathbf{X}^{\mathrm{II}}$.

This can be conceptualized in the framework of omitted variable bias. The econometrician estimates

$$\xi_{i,t} = \hat{\beta}^{\mathrm{II}}\mathbf{X}^{\mathrm{II}}_{i,t} + e_{i,t} \tag{1.7}$$

$\hat{\beta}^{\mathrm{II}}$ is biased from the corresponding coefficient in Equation 1.6. Instead of the ideal prediction, $\beta^{\mathrm{I}}\mathbf{X}^{\mathrm{I}}_{t-k,i} + \beta^{\mathrm{II}}\mathbf{X}^{\mathrm{II}}_{t-k,i}$, the predicted rainfall deviation is

$$\hat{\beta}^{\mathrm{II}}\mathbf{X}^{\mathrm{II}}_{i,t} = \left(\beta^{\mathrm{II}} + \beta^{\mathrm{I}}\frac{cov(\mathbf{X}^{\mathrm{I}}_{i,t}, \mathbf{X}^{\mathrm{II}}_{i,t})}{var(\mathbf{X}^{\mathrm{II}}_{i,t})}\right)\mathbf{X}^{\mathrm{II}}_{i,t} \tag{1.8}$$

Because accurate estimate of coefficients in Equation 1.6 is not the goal, the issue is not that $\hat{\beta}^{\mathrm{II}}$ is biased. Instead, the problem is that the estimate of unpredicted rainfall deviations, the residual $e_{i,t}$, still contains information which is predictable to the individual. In other words, rational expectations could still be formed over $e_{i,t}$ because both $\mathbf{X}^{\mathrm{I}}$ and $\mathbf{X}^{\mathrm{II}}$ still appear in the estimate of unpredicted rainfall.

$$e_{i,t} = \tilde{\xi}_{i,t} + \beta^{\mathrm{I}}\left(\mathbf{X}^{\mathrm{I}}_{i,t} - \frac{cov(\mathbf{X}^{\mathrm{I}}_{i,t}, \mathbf{X}^{\mathrm{II}}_{i,t})}{var(\mathbf{X}^{\mathrm{II}}_{i,t})}\mathbf{X}^{\mathrm{II}}_{i,t}\right) \tag{1.9}$$

So long as $\mathbf{X}^{\mathrm{I}}$ contains signals with predictive power ($\beta^{\mathrm{I}} \neq 0$), individuals who observe and respond to information in $\mathbf{X}_{i,t}$ will have behaviors correlated with $e_{i,t}$.

The econometrician has access to a additional set of signals, $\mathbf{X}^{\mathrm{III}}$, which may be correlated with $\mathbf{X}^{\mathrm{I}}$.[12] We can express $\mathbf{X}^{\mathrm{I}}$ as a function of $\mathbf{X}^{\mathrm{II}}$ and $\mathbf{X}^{\mathrm{III}}$

$$X^{\mathrm{I}}_{i,t} = \beta^{\mathrm{I,II}}\mathbf{X}^{\mathrm{II}}_{i,t} + \beta^{\mathrm{I,III}}\mathbf{X}^{\mathrm{III}}_{i,t} + \eta_{i,t} \tag{1.10}$$

[12]It is immaterial whether or not this correlation is causal. Regardless, the empirical method-

$\eta_{i,t}$ captures other factors which influence $\mathbf{X}^{\mathrm{I}}_{i,t}$. Because individuals observe $\mathbf{X}^{\mathrm{I}}$ directly, correlations between $\eta_{i,t}$ and rainfall are *not* exogenous. Hence there remains potential for individuals to form expectations using information which is simply not observable to the econometrician. A regression of $\mathbf{X}^{\mathrm{II}}$ and $\mathbf{X}^{\mathrm{III}}$ on $\xi_{i,t}$ yields

$$\xi_{i,t} =(\widehat{\beta^{\mathrm{II}} + \beta^{\mathrm{I}}\beta^{\mathrm{I,II}}})\mathbf{X}^{\mathrm{II}}_{i,t} + (\widehat{\beta^{\mathrm{III}} + \beta^{\mathrm{I}}\beta^{\mathrm{I,III}}})\mathbf{X}^{\mathrm{III}}_{i,t} + \beta^{\mathrm{I}}\eta_{i,t} + \epsilon_{i,t} \tag{1.11}$$

The residual from the regressing the econometrician's set of signals $\mathbf{W}$ on rainfall shocks $\xi_{i,t}$ in Equation 1.11 is $\tilde{\varepsilon}_{i,t} = \beta^{\mathrm{I}}\eta_{i,t} + \epsilon_{i.t}$, even if the econometrician cannot distinguish $\mathbf{X}^{\mathrm{II}}$ from $\mathbf{X}^{\mathrm{III}}$. The econometrician can attempt to minimize or eliminate $\eta_{i,t}$ by adding variables which are strongly correlated with $\mathbf{X}^{\mathrm{I}}$, while hopefully still retaining unpredicted rainfall variation, $\epsilon_{i,t}$. In the idealistic case where $\mathbf{W} = \mathbf{X}$, then $\tilde{\varepsilon}_{i,t} = \tilde{\xi}_{i,t}$. If $\mathbf{W} \supset \mathbf{X}$, the additional information can reduce the variance of $\tilde{\varepsilon}_{i,t}$, biasing estimates of $\tilde{\xi}_{i,t}$ towards zero. This is particularly troublesome in reduced-form settings, as it will inflate coefficients on $\tilde{\varepsilon}_{i,t}$. In an IV setting this particular issue is less troubling, as the first-stage coefficients are less often of direct interest. Additionally, it is not obvious how one would conclusively prove $\beta^{\mathrm{I}}\eta_{i,t} = 0$. Even if a number of behaviors appear orthogonal to $\tilde{\varepsilon}_{i,t}$, there are always additional unobserved behaviors. Results presented in this paper suggest a large amount of location-specific information is contained in $\mathbf{X}^{\mathrm{I}}$.

ology presented in the next section will use measures of ENSO as candidate $\mathbf{X}^{\mathrm{III}}$ signals. It is argued below in Section 3.1 that ENSO is strongly and causally correlated with potentially unobserved signals in $\mathbf{X}^{\mathrm{I}}$. Measures of ENSO are also explicit components of potential X^{II} variables such as government forecasts, although this does not imply either is a sufficient statistic for the other.

If researchers are instead interested in using the fitted value as an estimate of predicted rainfall, removing unpredicted variance from the residual implies adding unpredicted variance to the predicted value of rainfall, with corresponding estimation problems. Ideally $\mathbf{W}$ should be chosen in a way which minimizes $\eta_{i,t}$ relative to $\epsilon_{i,t}$ with minimal reduction in the variance of $\epsilon_{i,t}$. For analysis using $\tilde{\varepsilon}_{i,t}$, it is far more important that the residual only contain unpredicted variation than have the fitted value contain only predicted variation.

To reiterate a key distinction, this paper refers to $\tilde{\xi} \equiv \xi - E[\xi|\mathbf{X}^{\mathrm{I}}, \mathbf{X}^{\mathrm{II}}]$ as an unexpected rainfall shock which is useful for the theoretical model above but is often unobservable in practice. $\tilde{\varepsilon} \equiv \xi - E[\xi|\mathbf{X}^{\mathrm{II}}, \mathbf{X}^{\mathrm{III}}]$ refers to the unpredicted rainfall shock used as an estimate of the unobserved $\tilde{\xi}$ in the empirical examination below.

1.3 Background Information

1.3.1 Weather and Climate Systems

Although meteorologists can forecast this week's weather with general accuracy, the idea of longer-range forecasts is sometimes viewed as the meteorological equivalent to selling snake oil. Palmer (1993) acknowledges that "At first sight it might appear rather contradictory to suppose that the atmosphere is at all predictable beyond this deterministic limit." Yet over the past twenty years, meteorologists and physicists have made significant progress in describing the "coupled monsoon system" of oceanic and atmospheric factors which drive the El Niño Southern Oscillation (ENSO), as well as other global weather and climate systems.[13]

ENSO is of particularly importance in driving India's monsoon rainfall.[14] ENSO is a quasi-cyclical system of ocean surface temperatures and air surface pressures across the Southern Pacific Ocean which manifest themselves as the commonly known seasonal weather outcomes, El Niño and La Niña. In El Niño years countries in the Southwestern Pacific region (Indonesia, India, Australia, etc.) experience warmer temperatures and less rainfall, while areas in the Southeastern Pacific region (such as Peru) experience cooler temperatures and more rainfall. In La Niña years the global climatic teeter-totter reverses, with the Southwestern Pacific receiving cooler temperatures and more rainfall while the Southeastern Pacific receives warmer temperatures and less rainfall.

Academic meteorologists and physicists are well aware of the correlations between seasonal rainfall levels and the ENSO system. Yet it may still seem to be a huge leap of faith to believe farmers can form seasonal rainfall expectations more nuanced than a long-run mean. If the local weatherman armed with fancy

[13]See Wang, ed (2006) for a detailed description of recent advances in understanding this system.

[14]Rasmusson and Carpenter (1983) note correlation between below average monsoon rainfall and El Niño events. The percentage of variations in rainfall levels which ENSO models can account for varies over time, although it is always quite significant. On the upper end, Parthasarathy et al. (1988) develop statistical models which account between 70 to 83 percent of inter-annual rainfall variance for all of India, although some variables included ex post information.

satellite images and computer analysis can't perfectly predict the weather next weekend, how can illiterate farmers predict seasonal rainfall? The answer should begin by highlighting the distinction between deterministic prediction and simple correlations. Even if farmers cannot predict exactly how much rain will occur on any given day, they may still have reasons to believe that this year's summer will experience more or less rain than average.

There are also plenty of proxies for ENSO available to farmers. Many governments' meteorological agencies issue formal seasonal weather forecasts, including the United States, Peru, Australia, Brazil, Ethiopia and India.. Literature on informal proxies is sparse but suggestive. Moran et al. (2006) find that rural Amazonian farmers are reliant on personal experiences over official forecasts, and have accurate memory of ENSO events. They report that "Farmers monitor the behavior of the animals living on their farms, or wild animals living in the forests surrounding their farms; they learn to predict weather changes through monitoring cloud shapes, flowering or leaf dropping events in the local flora." Roncoli et al. (2002) report the use of similar forecasting techniques by farmers and herders in Burkina Faso, although these farmers perceive such traditional forecasts as becoming less reliable due to climate change. Orlove et al. (2000) investigate the mechanisms behind Andean farmers' centuries-old practice of using the visibility of the Pleiades star cluster to accurately forecast ENSO events. In El Nino years, an increase in subvisual high cirrus clouds reduces the number of visible stars, a recognizable signal farmers rely upon in determining when to plant their crops.

Empirically pinning down the particular mechanisms by which individuals form weather expectations is difficult. Yet one might hope that measures of ENSO could form a sufficient statistic for farmers' ENSO-related information, enabling researchers to control for weather expectations. Conveniently for the econometrician, El Niño and La Niña events are measured in practice by several highly correlated single-dimensional summary statistics. The Oceanic Niño Index (ONI) measures the difference between western and eastern ocean temperatures. The Southern Oscillation Index (SOI) measures the differences between western and eastern air pressures. Warmer ocean temperatures correspond with low air pressure, while cooler ocean temperatures come with high air pressure. There are also aggregate measures such as the Multivariate ENSO Index (MEI). Any of these measures seem plausible candidates for capturing the ENSO-related information available to rural farmers. While results will show this approach appears promising at aggregate levels, agents ultimately possess too much location-specific information for this approach to be successful.

1.3.2 Crop Selection, Weather Adaptation, and Agriculture in India

India has two major growing seasons, *kharif* and *rabi*, with specific crops grown in each season.[15] Across the districts of Madhya Pradesh, rice, sorghum, and maize are the most widely grown crops during the *kharif* season, while wheat is the main crop grown during the *rabi* season.[16] There is significant variance at the village level; the main crops grown in the villages examined in this setting also include cotton, pigeon pea, and castor bean. However, rice and sorghum are grown in almost all districts of Madhya Pradesh. This analysis focuses on the kharif season for several reasons. First, the meteorology literature suggests that monsoon seasons may be easier to generate seasonal rainfall expectations than dry seasons.[17] Second, the kharif season also exhibits greater variability in rainfall for potential identifying variation, as seen in Figure 2.1. It is also important to note that although both rice and sorghum perform best under similar aggregate water requirements, sorghum is relatively drought resistant while rice is relatively sensitive to drought.[18] Hence when farmers expect a dry season there should be less rice and more sorghum planted. Maize is neither as water-intensive as rice nor as drought-resistant as sorghum. For the crops prevalent among the ICRISAT

[15]These seasons can be roughly thought of as summer and winter season, respectively, but are more accurately monsoon season and dry season.

[16]Some varieties of sorghum can be grown during the rabi season as well. Some of the data sets examined in this paper separately measure rabi sorghum and kharif sorghum.

[17]Rodó et al. (1997), Stockdale et al. (1998).

[18]Brouwer and Heibloem (1986).

villages, both pigeon pea and castor bean are relatively drought resistant.

In 1966, high yield variety (HYV) crops were introduced in India, sparking India's "Green Revolution". Before the introduction of HYV crops, India was a major agricultural importer.[19] Thanks in large part to the increased yields of HYV crops, India turned the tables to become one of the worlds leading agricultural exporters. The requirement for HYV crops to realize these greatly increased yields was larger, more consistent supplies of water. HYV crops increased the marginal benefit of irrigation, and became most prevalent in areas which already had or subsequently introduced irrigation. One might expect farmers with access to irrigation to be less responsive to short run weather fluctuations.

There is a long and rich literature on farmers' crop choices. Agricultural economists, development economists, historical economists, agronomists, and even agricultural and biomechanical engineers have been writing on the subject for many decades. Many of these are not studies of human behavior or preferences, but rather attempts to suggest a profit-maximizing selection given a set of observable variables. The literature examining behavioral adaptation has largely focused on adaptation to long run climate change. There are many long-run estimates of new equilibria under climate scenarios, and few reliable methods of analyzing how these transitions occur over time. Prior research relied on long-run cross-sectional

[19] A major policy debate of the time was whether to continue researching India-specific high yield varieties, or to more quickly import varieties from other countries. Ultimately high yield wheat and rice were imported from Mexico in 1966. For an interesting discussion of this background, see Abler et al. (1994).

comparisons to estimate equilibria, and interest in other methodology is only beginning. The results presented here suggest that adaptation to annual rainfall expectations may be a major mechanism in driving long run adaptation to climate change.[20]

The list of research considering response to weather expectations is fairly short. In their text on the ICRISAT villages, Walker and Ryan (1990) note that "When the monsoon "plays-truant" or is initially erratic, the planned cropping strategy may no longer be optimal; farmers adapt to emerging information on rainfall events by changing crops or by fallowing land." Wallace and Vogel (1994) speculates that "a forecast of El Niño weather might induce farmers to sow more rice and less cotton than in a year without El Ni no." Kochar (1999) examines the ability of labor supply adjustments smooth unexpected income shocks. Concurrent research such as Rosenzweig and Udry (2013) find early-season agricultural investments respond to government forecasts. Their focus is on returns to early-season investment rather than general identification concerns.

[20]For examples of profit-maximizing selection, see Mohan and Arumugam (1994), Mjelde et al. (1996), or Cabrera et al. (2007). Examples of papers estimating long-run crop equilibria include Rosenzweig and Perry (1994), Mendelsohn and Dinar (1999), O'Brein et al. (2004), Tubiello et al. (2007), Seo and Mendelsohn (2008), and Kurukulasuriya and Mendelsohn (2008). Examples of papers seeking identification other than long-run cross sections include Taraz (2012), Burke and Emerick (2013), and others.

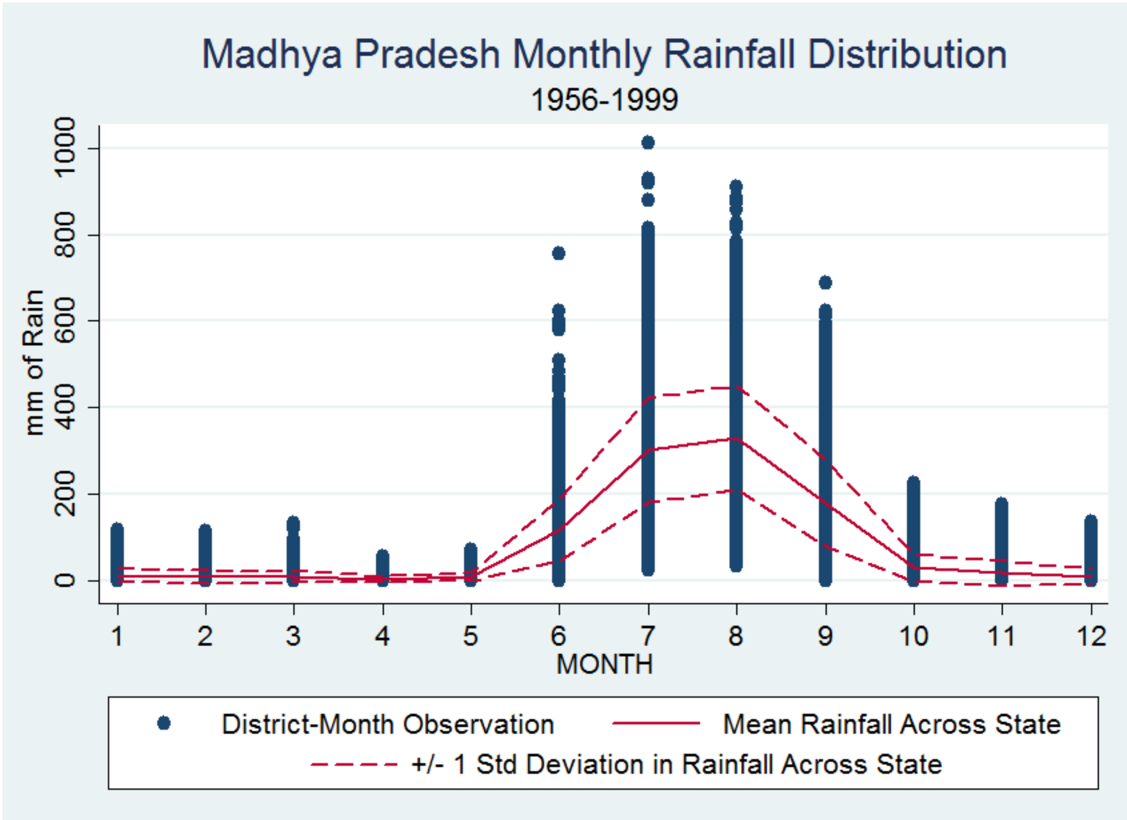

Figure 1.1: Madhya Pradesh Monthly Rainfall Distribution

1.4 Anticipatory Adaptation to Rainfall

1.4.1 Data

Given the common usage of rainfall as a shock to income through agricultural yields, rainfall is more commonly used for variation in countries with large populations of low-income farmers. Because Table 1.1shows India is the most common single-country setting in recent papers, this paper looks for evidence of adaptive behavior by Indian farmers. Results are confirmed across two district-level data sets. These data sets are commonly used by many researchers, including among papers using rainfall for identifying variation.

This paper combines many sources of agricultural and meteorological data. The International Crops Research Institute for the Semi-Arid Tropics Meso Level Database (ICRISAT Meso) contains annual district-level observations on the sown acreage of various crops, as well as yields and prices, over the years 1966-1999. ICRISAT Meso also includes useful measures of soil type, irrigation and land use. As a robustness check, district-level results are replicated using the World Bank India Agriculture and Climate Data Set (IAC). Similar in format to the ICRISAT Meso data, this data set contains annual district-level observations on the sown acreage of various crops, as well as yields, over the years 1957-1987. District-level analysis focuses on the state of Madhya Pradesh because rice and sorghum are major crops grown in almost every district of this state. For both ICRISAT Meso and IAC data sets, year refers to agricultural year, not calendar year. This means that the 1970-1971 rabi season is included in the 1970 data. This is potentially problematic because if a crop is also grown in the rabi season, changes in annually summed crop acreage could response ex post to kharif rainfall. Hence analysis is focused on crops which are typically grown only in one specific season or cases where data separately identifies kharif and rabi acreage. ICRISAT Meso data separates kharif sorghum from rabi sorghum, and indeed almost all sorghum is planted in the kharif season. Summary statistics for the ICRISAT Meso data and IAC data employed in this paper can be found in Table 1.2.

Historic weather data is obtained from India Water Portal's Meteorological Dataset (Met Data), in the form of monthly district-level observations.[21] Although observations are available for as early as 1901, the initial district-level analysis follows begins by following the standard practice of using mean rainfall over the observed period, 1956-1999. Later analysis also examines expectations over rolling means from the prior 20 years. Other than a gap in data availability for 2003, monthly district-level rainfall data is available through 2011.

District boundaries are defined by ICRISAT using 1966 parent boundaries, and IAC and Met Data appear to follow similar methodology. If for any reason the geography for rainfall data does not perfectly match the geography for the agricultural data, the impacts of rainfall and rainfall expectations on agricultural activity should be biased towards 0. India Water Portal excludes districts which are missing data on over 25% of district area. This would impact external validity if this exclusion is correlated with farmers' ability to form and respond to weather expectations.

Since 1886 the India Meteorological Department (IMD) has issued annual seasonal forecasts for the upcoming monsoon season. The India Meteorological Department, Pune provides historical forecasts dating back to 1932.[22] These fore-

[21]Met Data converts long-lat data to district-level observations. The underlying weather observations for 1901-2002 come from the Climate Research Unit (CRU) TS2.1 dataset, out of the Tyndall Centre for Climate Change Research, School of Environmental Sciences, University of East Anglia in Norwich, UK. The underlying data for 2004-2011 rainfall comes from the India Meteorological Department (IMD).

[22]A short and intriguing history behind these forecasts, as well as the raw data, is available at the website of IMD, Pune, http://www.imdpune.gov.in/research/ncc/longrange/longrange_index.html.

casts take a variety of forms, and were hence recoded to indicators for predictions of above normal, below normal, and normal rainfall for the corresponding geographic area. Where possible, above normal corresponds to a forecast of total monsoon rainfall being 110% or more of the long run average, while below normal corresponds to a forecast of 90% or less of long-run average. For the period 1950-2010 period, this makes 15.7% of forecasts for above average monsoon rainfall, and 15.6% of the forecasts for below average monsoon rainfall. Historical IMD forecasts might easily be perceived as an obvious indicator of the prevailing best expectations of the time, but are very geographically general and hence likely do not accurately reflect more geographically precise expectations.[23]

Finally, several common measures of ENSO are used. Results presented below use the Oceanic Niño Index (ONI), although adding or substituting the highly correlated Southern Oscillation Index (SOI) or Multivariate ENSO Index (MEI) makes little empirical difference. All ENSO measures are obtained directly from the NOAA Climate Prediction Center.[24] Summary statistics for IMD forecasts and monthly ONI measures can be found in Table 1.3.

[23]From 1932-1988 were issued for Northwest India and the Peninsula. From 1989-1998 forecasts were only issued for India as a whole. From 1999-2003 forecasts are issued for the country as a whole as well as Northwest India, Northeast India, and the Peninsula. From 2003-present Peninsula forecasts were replaced with separate forecasts for Central India and the Southern Peninsula.

[24]These and a variety of other climate indices useful for constructing expectations can be found at http://www.esrl.noaa.gov/psd/data/climateindices/list/.

1.4.2 Analysis

The main empirical analysis in this paper shows that individuals are able to form rainfall expectations and respond by altering the acreage sown with drought-resistant and water-intensive crops. Following the model described in Section 2.4, kharif rainfall is then decomposed into variation which is predictable and unpredictable given information available to the econometrician.[25] The response of crop acreage to these predicted and unpredicted estimates of rainfall is then examined. Because this hypothesis conflicts with the standard treatment of rainfall as an unpredictable exogenous shock, these results are replicated in a second separate and commonly used data set.

Predicted and unpredicted estimates of kharif rainfall are constructed using IMD forecasts, district-level deviations from mean rainfall over the past five years, and both monthly levels and squared monthly levels of ONI.[26] District fixed effects (α_i) and district-specific quadratic time trends ($\theta_\mathbf{i}\mathbf{Y_t}$) capture local climate trends of which local farmers are likely aware. $\mathbf{L}_{i,t}$ represents a vector of interactions between lagged rainfall and soil type as a control for residual soil moisture. Removing $\mathbf{L}_{i,t}$ from the information set does not significantly alter results, but is included because accurate Murphy-Topel standard errors require all second-stage

[25]Kharif rainfall is defined throughout this paper as the sum of June, July, August, and September rainfall. A robustness check shows excluding June rainfall from this calculation does not alter conclusions.

[26]May is used as the cut-off for ONI measures because kharif planting occurs subsequently in June, although as discussed above one could in theory add later ENSO measures. Lags of January through May of the observed year and January through December of the previous year are employed.

controls correlated with future rainfall be included in the first stage information set. The first stage regression can hence be written as

$$\xi_{i,t} = \beta^W \mathbf{W}_{i,t} + \alpha_i + \theta_{\mathbf{i}} \mathbf{Y_t} + \gamma \mathbf{L}_{i,t} + \tilde{\varepsilon}_{i,t} \tag{1.12}$$

Deviations from mean rainfall, $\xi_{i,t}$, are standardized for clear interpretation of second-stage results. For initial results, $\xi_{i,t}$ is defined following the standard approach in the literature of $\xi_{i,t} - \overline{\xi}_i$ where $\overline{\xi}_i$ is the mean in region i over the sample period, 1956-1999. Later, defining rainfall shocks as the deviation from a rolling mean will be discussed and implemented.

The particular coefficients from the first stage regression are not of interest in the context of this paper. More important is understanding how much variation in rainfall is explained. If the amount of predictable variation is negligible and there is no evidence of behavioral responses to rainfall, it becomes easier to believe that rainfall may indeed be unpredictable. To a researcher attempting to isolate unpredicted rainfall, if almost all of the variation is predictable then it is difficult to argue that an remaining unpredictable rainfall provides enough identifying variation to recovers coefficients of interest. While it is not clear what constitutes "enough" variation in either direction, Figure 1.2 provides a visualization of the available variation by plotting $\xi_{i,t}$ and $\tilde{\varepsilon_{i,t}}$ for the Raisen District of Madhya Pradesh under a variety of information sets, $\mathbf{W}$. Across all districts of Madhya Pradesh, R^2 is only .1336 when $\mathbf{W} = \{$IMD forecasts, $\xi_{i,t-1}, \ldots, \xi_{i,t-5}$, residual soil moisture, district fixed effects, district-specific time trends$\}$. The continued responsiveness of acreage decisions to rainfall shocks which are unpredicted given this information set is consistent with farmers observing some $\mathbf{X}^{\mathrm{I}}$ which includes additional information. Adding two years of monthly ONI levels and squares to $\mathbf{W}$ leaves no average response of rice, sorghum, or maize acreage and increases the R^2 value to 0.6720.

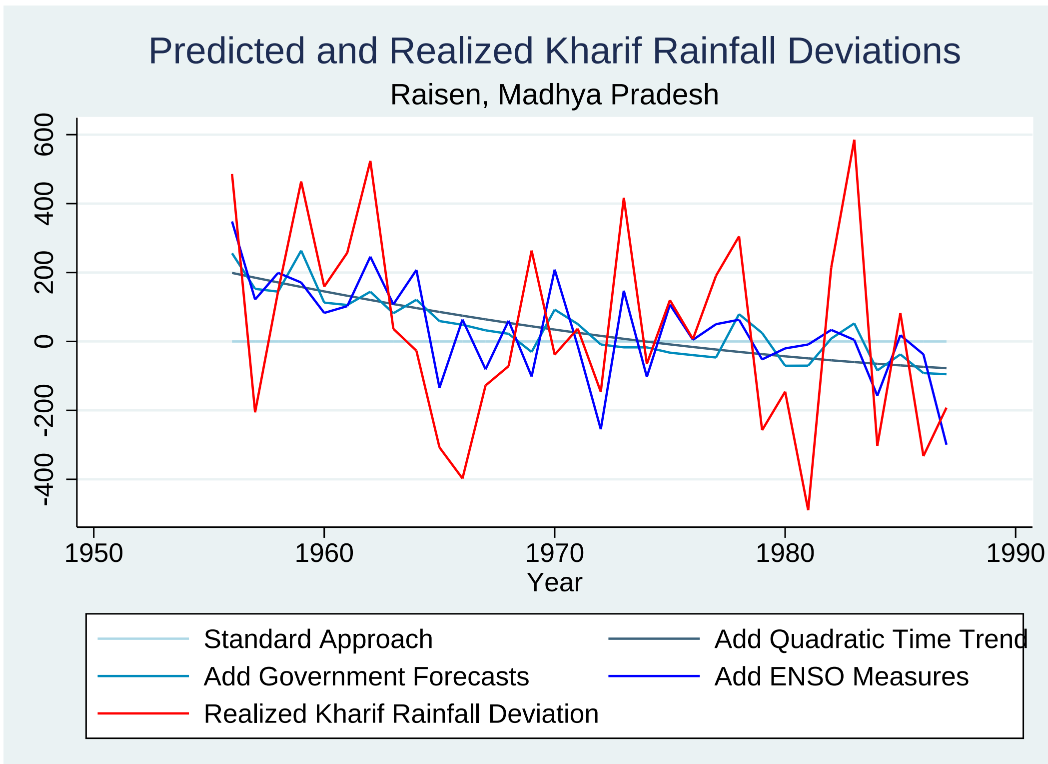

Figure 1.2: An Example of Expected Rainfall Data

To examine whether individuals are able to form and respond to rainfall expectations, hectares planted to a given crop in year t and district i are regressed on demeaned $(\xi_{i,t})$, predicted $(\xi_{i,t} - \tilde{\varepsilon}_{i,t})$, and unpredicted rainfall $(\tilde{\varepsilon}_{i,t})$.

$$\text{Hectares of Crop}_{i,t} = \beta^a(\xi_{i,t}) + \alpha_i + \theta_{\mathbf{i}}\mathbf{Y_t} + \gamma\mathbf{L}_{i,t} + e_{i,t} \tag{1.13}$$

$$\text{Hectares of Crop}_{i,t} = \mathbf{\Gamma}^a(\xi_{i,t} - \tilde{\varepsilon}_{i,t}) + \alpha_i + \theta_{\mathbf{i}}\mathbf{Y_t} + \gamma\mathbf{L}_{i,t} + e_{i,t} \tag{1.14}$$

$$\text{Hectares of Crop}_{i,t} = \mathbf{\Gamma}^b(\xi_{i,t} - \tilde{\varepsilon}_{i,t}) + \beta^b(\tilde{\varepsilon}_{i,t}) + \alpha_i + \theta_{\mathbf{i}}\mathbf{Y_t} + \gamma\mathbf{L}_{i,t} + e_{i,t} \tag{1.15}$$

Temporally fixed variables are captured by district fixed effects, α_i, while $\theta_{\mathbf{i}}\mathbf{Y_t}$ controls for district-specific quadratic time trends. Rainfall for the upcoming season is strongly correlated with last year's rainfall (violating the *i.i.d* portion of Assumption 1), so last year's rainfall deviations and residual soil moisture are important controls to ensure the coefficients on current rainfall shocks do not reflect income constraints from the prior year's rainfall.

Because demeaned rainfall shocks, $\xi_{i,t}$, are standardized before first and second stage estimation, β^a can be interpreted as the correlation between a standard deviation increase in future or expected rainfall on crop acreage. If rainfall is unpredictable, β^a should be both statistically and economically insignificant. If $\beta^a \neq 0$, this is consistent with anticipatory adaptation, particularly if the sign of β^a is consistent with agronomic practices.

Because predicted and unpredicted rainfall sum to total rainfall, a unit increase in either is associated with more than a one standard deviation increase in overall rainfall. If $\tilde{\varepsilon}_{i,t}$ truly gives no more information about economic behaviors than random noise, it should be the case that $\Gamma^a = \Gamma^b$ (henceforth referred to as Γ). The introduction of unpredicted variation into $\xi_{i,t} - \tilde{\varepsilon}_{i,t}$ will reduce Γ towards 0 as with classical measurement error. β^a represents the extreme example of this effect, including both predictable and unpredictable variation. If predictions perfectly estimate expectations, Γ would be interpreted as a causal estimate of the sensitivity of crop acreage to expectations over the pending season's rainfall. If all variation expected by farmers is removed from $\tilde{\varepsilon}_{i,t}$, it should be that $\beta^b = 0$. As discussed below, if $\beta^b = 0$ for one crop, this does not necessarily imply $\beta^b = 0$ for all crops or that farmers have no information over over $\tilde{\varepsilon}_{i,t}$.

Due to their importance and complexity in this setting, a few words on the proper standard errors when using generated regressors is worthwhile. The examinations of acreage response to both predicted and unpredicted rainfall involve both generated regressors and generated residuals. This paper employs the analytical standard error correction described by Murphy and Topel (1985) following the methods of Hole (2006). Confirming whether fitted regressors are significantly different from zero while fitted residuals are statistically insignificant does require care. Fitted regressors and other covariates require standard error corrections while generated regressors do not, so long as any second-stage covariates are also included in the first stage.[27] When OLS standard errors are smaller than Murphy-Topel standard errors, applying the Murphy-Topel correction to all variables except generated residuals is is also the most conservative approach for all hypotheses tested here. Particularly in linear models, certain conditions on the score functions of the first and second stages can result in the correction decreasing standard errors. To be as conservative as possible, the larger of corrected or uncorrected standard errors is applied to all variables except the generated residual. In

[27] See Pagan (1984)

all stages, the variance-covariance matrix is clustered at the district or household level as appropriate.

1.4.3 Results

Table 1.5 shows the response of the main crops in Madhya Pradesh to demeaned rainfall, as well as to the econometrician's estimates of predictable and unpredictable rainfall. In columns (1) and (4), β^a shows the response of crop acreage to the future rainfall is both statistically and economically significant for rice and sorghum. The interpretation is that rainfall one standard deviation below the mean is correlated with a 495 hectare decrease rice acreage and a 1,305 hectare increase in sorghum. As discussed above, these acreage changes are consistent with those a profit-maximizing farmer with information about future rainfall would make, since sorghum is relatively drought-resistant and rice is relatively water-intensive. The magnitudes are plausible and economically significant, representing a roughly 1-3% shift in total district-level acreage planted. However, the temporal order for these correlations is incorrect for a causal relationship – farmers must make decisions about how many hectares to sow *before* observing rainfall outcomes. An identical exercise using IAC data yields similar results (Table 1.6).

In column (7) of Table 1.5, the response of a third major crop, maize, show no acreage response to rainfall variation. If examined in isolation, a researcher might be tempted to conclude that the lack of a change in average maize acreage suggests no anticipatory adaptation of any sort occurs, and that rainfall is unpredictable. When examined in conjunction with rice and sorghum results, the coefficient on maize is consistent with maize lying between rice and sorghum in drought sensitivity. Indeed, the identical exercise using IAC data in Table 1.6 finds a small but significant response of maize acreage to rainfall variation. Because the econometrician will never observe all behaviors through which an individual may engage in anticipatory adaptation, this reinforces that a lack of adaptive behavior in one case implies a lack of adaptive behavior in all cases.

Columns (2), (5), and (8) of both Table 1.5 and Table 1.6 show that the response to the predictable portion of rainfall is always of larger magnitude than the response to the total rainfall shock. This is consistent with having removed attenuation bias caused by unpredictable variance. While the monotonicity of the change is compelling, a unit increase in predicted rainfall may also be associated with more than a one-standard deviation in total rainfall, so the two magnitudes are not directly comparable. It is also not surprising that a function of rainfall deviation has the same correlation structure as rainfall deviation itself. A better test is whether there is behavioral response to the unpredictable portion of rainfall. If the correlation is driven by something other than rainfall expectations, there does not appear to be an obvious reason why it would not also exist in unpredicted rainfall. Columns (3), (6), and (9) find no significant correlation between crop acreage and unpredictable rainfall variation, supporting the interpretation of response to rainfall expectations. Again, robustness checks using IAC data offers the same interpretation (Table 1.6).

Other results also also consistent follow the model's predictions. In Table 1.5 and Table 1.6, Γ^a is effectively identical to Γ^b, consistent with the idea that adding a variable containing unexpected variation should not influence other coefficients. In these tables, it also appears promising that the econometrician's information set may indeed be a sufficient statistic for the individual's information set. As discussed in the next section, this conclusion is premature because average responses at the district level may hide significant heterogeneity in local information.

Table 1.7 and Table 1.8 replicate Table 1.5 and Table 1.6, except they exclude ENSO measures from vector of information of rainfall information, $\mathbf{W}$. Columns (3), (6), and (9) now show that when sufficient information is not included in $\mathbf{W}$, individuals remain able to form and respond to expectations over $\tilde{\varepsilon}_{i,t}$. This is consistent with the story of ENSO measures residing in the $\mathbf{X}^{\text{III}}$ information set and serving as a proxy for unobserved $\mathbf{X}^{\text{I}}$. When ENSO measures were included in $\mathbf{W}$, the coefficients on $\tilde{\varepsilon}_{i,t}$ are economically and statistically insignificant.

One alternative explanation of the sensitivity of farmer crop acreage to rainfall outcomes is that farmers respond to medium-run expectations driven by meridional rainfall regimes, as described by Taraz (2012). These 30-40 year cycles of above or below average rainfall do influence rainfall deviations. Simply using deviation from a rolling mean over a shorter time period instead of a fixed mean over long periods removes the ability to form expectations based on such medium-run regimes or cycles, so that under such a model the remaining deviations should again be random shocks. Table 1.9 replicates the main analysis using deviations from 20-year rolling means rather than fixed means across the time period of crop acreage data. Coefficients decrease by only trivial amounts and the conclusions described above all remain robust. IAC data again provide similar results.This provides strong suggestive evidence that, at least in this setting, short run expectations are the driving force of any response to any such medium-run expectations.

A skeptical reader might be concerned that if crop acreage data is backed out from yields and not directly observed, then the district-level results will be biased in favor of finding response to predicted weather. There are a variety of reasons to believe this is not an issue. First, if this were entirely true, then it should not be found that unpredicted rainfall is uncorrelated with acreage planted. Second, acreage planted is still sensitive to rainfall shocks after adding controls for yields. Another concern might be that farmers are delaying planting and responding to early June rains. Again, if this were true then it should not be found that unpredicted rainfall is uncorrelated with acreage planted. Because Walker and Ryan (1990) discuss some farmers delaying or altering planting into late May, Table 1.10 addresses this concern by showing results when June rains are excluded from the analysis. Response to July-September rain remains consistent with the results presented above. Again, IAC data provides similar results.

The notion that farmers adjust their acreage decisions in anticipation of rainfall shocks contradicts common identification assumptions.[28] It should be noted that these results do not contain information about whether the occurring adaptation is within or between farmers. Regardless, the occurrence of any adaptive behavior implies that rainfall is not an unpredictable shock. Adaptation of crop selection to future rainfall implies the impact rainfall has on income is a function of adaptive behaviors made by observant farmers, and is hence neither unexpected nor exogenous.

1.5 Local Information Issues

The above results presented strong evidence of adaptation to expectations over future rainfall events. This means rainfall is not an unpredictable shock, implying the changes in income and other outcomes which result from rainfall are also not unpredicted. Because farmers are clearly adapting in ways which change their subsequent income outcomes, the changes in income resulting from rainfall are not an exogenous shock.

Given the insignificant response to rainfall unpredicted given the econometrician's data set, it is tempting to conclude that the econometrician's infor-

[28]For example, in examining take-up of high-yield crops, Suri (2011) is clear that "Central to identification is the fact that the hybrid seed choice is made before the farmer experiences most of the agricultural shocks to yields," or that transitory shocks to yields "do not affect the farmer's decision to use hybrid and, crucially, the farmer's switching behavior."

mation set may be a sufficient statistic for the information to which individuals respond. If that were the case, the econometrician could simply control for rainfall expectations, and use unexpected shocks for identification. The estimation and interpretation issues surrounding such an attempt are detailed in Section 2. It was noted in throughout this paper that an econometrician will never observe all potential adaptive behaviors, nor does the lack of a behavioral response imply the rainfall was unexpected. Because it seems almost impossible to every completely allay such concerns, they may appear overly cautious. Hence this section presents evidence of a third and very significant data concern faced by the econometrician.

The rainfall data observed by the econometrician, be it satellite data, gauge data, forecasts, or something else is almost always from a different or geographically broader area than the farmer. This issue is a familiar topic when working with rainfall index insurance, where basis risk refers to the uninsured risk due to difference between rainfall at the farmer's field and rainfall measured at the site used for determining insurance payouts. In using rainfall for identifying variation, econometricians face a similar problem.

Suppose farmers have location-specific knowledge ($\mathbf{X}^{\mathrm{I}}$) about how rainfall on their particular farm is correlated with the rainfall which is unpredicted at the district level given $\mathbf{W}$.These farmers would adjust their planting behavior based on information unobserved by the econometrician. The econometrician can ensure $\mathbf{W}$ contains enough information so that the average farmer does not expected rainfall above or below the predicted predicted rainfall outcome. But this average can continue to hide important heterogeneity. The econometrician can only decrease this concern by obtaining increasingly fine-grained data.[29]

In other words, local information may result in farmers adapting differently to what the econometrician perceives as a uniform rainfall expectation but local farmers correctly perceive as a heterogeneous expectation. It is possible to test whether households adapt at a local level over information which is appeared unexpected on average at the district level. If such behavior exists, it presents two problems for identification. First, it reiterates that the econometrician is unable to correctly estimate expectations as would be necessary to form an unexpected shock useful for identification. Second, it suggests suggests not only that adaptations vary over fine-grained geographic areas with different location and climate characteristics, but also that because such differences are known farmers may have sorted non-randomly over these areas.

1.5.1 Data

To examine this issue, "Generation II" ICRISAT Micro Level Data (ICRISAT Micro) is used for household-level analysis. In 1975, ICRISAT began a commonly used household-level surveys following 40 households in each of six villages from three districts in the semi-arid tropical parts of India. In 1980 four additional villages were added including two villages in the Raisen District of Madhya Pradesh.

[29]Note that Indian districts are on average slightly larger than U.S. counties.

Households are stratified on farm size, and detailed interviews record household demographics and the timing and value of various farming and economic activities.[30] In 2001 through 2008, a follow-up series of surveys, "Generation II," was targeted at household surveyed earlier and spin-off households, with additional households surveyed to fill gaps in farm size stratification due to attrition.[31] ICRISAT Micro data involves includes detailed time data which permits separate identification of kharif and rabi activities. The "Generation II" data is particularly useful because it provides detailed records of the amount of acreage farmers devote to each particular crop. Summary statistics for the ICRISAT Micro data can be found in Table 1.4.

1.5.2 Analysis

The cropping decisions of individual farmers can be examined in a similar fashion to the aggregate district-level acreage responses. The same set of information, $\mathbf{W}$, is used to predict rainfall outcomes.[32] Here, it is important to note that the most geographically precise information in $\mathbf{W}$ is at the district level. This means that at best, $\mathbf{W}$ can be used to estimate average district-level expectations.

[30]The ICRISAT Micro data interviews 10 households from each category of labor, small farm, medium farm, and large farm. The analysis presented below focuses on the farming households, following the standard story of rainfall impacting income through crop yields.

[31]Due to data limitations in linking households across all survey years, the households in the first and second rounds of the Generation II survey (2001-2004 and 2005-2008) are conservatively treated as new households in the assignment of household fixed effects. This is equivalent to the equally reasonable approach of interacting a dummy for changes in the survey instrument with the household fixed effects.

[32]Because ICRISAT Micro data includes villages from multiple states, the coefficients on $\mathbf{W}$ are allowed to vary at the state level.

Do these district-level expectations still appear to yield a sufficient statistic for anticipatory adaptation?

$$\begin{aligned}\text{Percentage of Acreage Devoted to Crop}_{i,t} = &\beta^c(\xi_{i,t}) + \alpha_i + \theta_{\mathbf{i}}\mathbf{Y_t} \\ &+ \gamma\mathbf{L}_{i,t} + e_{i,t} \qquad (1.16)\end{aligned}$$

$$\begin{aligned}\text{Percentage of Acreage Devoted to Crop}_{i,t} = &\mathbf{\Gamma}^c(\xi_{i,t} - \tilde{\varepsilon}_{i,t}) + \alpha_i + \theta_{\mathbf{i}}\mathbf{Y_t} \\ &+ \gamma\mathbf{L}_{i,t} + e_{i,t} \qquad (1.17)\end{aligned}$$

$$\begin{aligned}\text{Percentage of Acreage Devoted to Crop}_{i,t} = &\mathbf{\Gamma}^d(\xi_{i,t} - \tilde{\varepsilon}_{i,t}) + \beta^d(\tilde{\varepsilon}_{i,t}) + \alpha_i + \theta_{\mathbf{i}}\mathbf{Y_t} \\ &+ \gamma\mathbf{L}_{i,t} + e_{i,t} \qquad (1.18)\end{aligned}$$

Because different farmers with different crop preferences may own different amounts of land, the outcome variable of interest is percentage of sown crop acreage devoted to a particular crop rather than total hectares devoted to a particular crop. Coefficients should be interpreted as the additional percentage of sown acreage devoted to the given crop. Because fallow lands are not well-recorded in some years, the impact of demeaned rainfall shocks on the percentage of land devoted to a given crop would be biased if households are more likely to report fallow fields (mechanically increasing their total field area) in years with rainfall shocks. The reported acreage of fallow fields is indeed very strongly correlated with future rainfall shocks.To avoid this bias, the outcome variable is percentage of each year's total planted acreage devoted to a given crop. This identifies only the intensive margin of switching between crops, and not extensive margin of planting more or less area entirely. $\xi_{i,t}$ is defined as deviation from a rolling mean over the past 20 years.

As before, if rainfall is unpredictable, β^c should be both statistically and economically insignificant. If $\beta^c \neq 0$, this is consistent with anticipatory adaptation, particularly if the sign of β^c is consistent with agronomic practices. If all variation expected by farmers is removed from $\tilde{\varepsilon}_{i,t}$, it should be that $\beta^d = 0$. However, if farmers have additional information about rainfall which is unknown to the econometrician, the sign of β^d is unclear. On one hand, expectations of lower rainfall could result in a increase of drought-resistant crops and decrease in water-intensive crops. On the other hand, if farmers expect rainfall to be low but better in their area than other nearby areas, expectations of increased prices due to scarcity may encourage an increase rather than decrease in prices. This issue is less of a concern at broader geographic levels if increasing transportation costs reduce the benefits of arbitrage.

When the first and second stages do not involve one-to-one mappings of observations, standard error adjustments for generated regressors are potentially infeasible. The intuition is simplest in the logic of bootstrapping. An alternative to Murphy-Topel corrections would be a double-bootstrapping procedure which correctly accounts for sampling error by drawing a subsample of districts, estimating rainfall expectations, applying these expectations to the district or household-level data, and repeating this procedure a large number of times. However, in cases where the level of the panel variable in the second stage (households) is drawn from a small number of panel variables in the first stage (districts), block bootstrapping at the district level fails to generate regressors for large groups of second-stage observations on most re-samplings. While this means the standard errors on Γ^c and Γ^d may be slightly too small, the standard errors on β^c and β^d remain correct. Standard errors are clustered at the household level.

1.5.3 Results

Table 1.11 and Table 1.12 show household level results for the six most prominent crops in the ICRISAT Micro villages (cotton, pigeon pea, castor beans, soybeans, sorghum, and rice).[33] Again results show evidence that at the individual household level, farmers alter the acreage devoted to various crops in anticipation of future weather outcomes. When district-level measures of rainfall are one standard deviation lower than the mean, the acreage devoted to pigeon pea and castor bean, both drought-resistant crops, increase by 12% and 13% respectively. At the same time, rice yields decrease by over 11%. Changes in sorghum are not significant, perhaps because other drought resistant crops are more popular in these villages, but the sign of β^c remains consistent with sorghum being a drought-resistant crop. These magnitudes are larger than found at the district level as a whole, suggesting that farmers with less acreage may be disproportionately sensitive to seasonal weather expectations. This is consistent with the well-documented conclusion that

[33]Results from additional crop as well as decisions to rent out or fallow fields display similar signs of anticipatory adaptation and local information.

poor, small-scale farmers are quite risk averse.

Those crops which individuals adjust in anticipation of future rainfall also show significant response to predicted rainfall. Again, the magnitudes are not comparable because a one unit increase in predicted rainfall is not equal to a one unit increase in total rainfall. The main difference between these results and district level results is that at the individual farm level the acreage devoted the three most common crops exhibits significant response to variation which is unpredictable given the econometrician's information set, $\mathbf{W}$. This is consistent with farmers possessing location-specific information about how rainfall outcomes are correlated within the district. While the magnitude of β^d is difficult to interpret, standard errors on this key variable are correct and show a significant and agronomically sensible response to rainfall variation unpredicted by the econometrician. Both pigeon pea and castor bean are drought-resist crops, while cotton is comparatively water-intensive.

The important take-away from Table 1.11 and Table 1.12 is that individual farmers significantly alter their crop selections in response to location-specific information. While there is less evidence of response to local information in soybeans, sorghum, or rice, the response of any crop is troubling. Additional crops choices beyond these six most common crops also show evidence of anticipatory adaptation based on local information, including other pulses, sugarcane, sunflowers, onions, and even fallowing.[34] Looking at only a limited selection of potential behavior

[34]For conciseness, these results are available upon request.

adaptations, such as rice and sorghum acreage, could mistakenly cause researchers to mistakenly assume expectations have been completely captured. The researcher will never observe all possible adaptive behaviors, and hence can never be certain that unobserved adaptation has not occurred. These results further re-affirm that income and other outcomes which result from rainfall are not random, but are endogenous functions of decisions made by heterogeneous farmers. Further, they suggest that location-specific information unobserved by the econometrician is an important part of the endogenous crop selections made by farmers.

1.6 Measuring the Impacts of Adaptation

1.6.1 Methodology

Given that individuals engage in anticipatory adaptation, a practical concern is the extent to which adaptation alters rainfall-based identification strategies. Do the impacts of adaptation represent a trivial part of the correlation between rainfall and any outcomes of interest? If so, concerns about the interpretation of rainfall-based identification strategies are largely semantic. To what extent has adaptive behavior altered our understanding of these estimates?

To answer this question, this paper uses ICRISAT Micro data to compare estimates based on demeaned rainfall to estimates using unpredicted rainfall. This paper examines both reduced form estimates and IV approaches which instrument income with rainfall deviations. The particular outcomes examined here are chosen from popular topics in the literature which are available in the ICRISAT Micro data, in particular schooling, migration outcomes, and birth and death rates.[35] Equation 1.19 through Equation 1.23 show the explicit empirical approach. In each equation, β_i represents the coefficient of interest.

OLS

$$\text{Outcome}_{h,t} = \beta_1 \text{Income}_{h,t-1} + \theta_i^a \mathbf{Y_t} + \psi C_{h,t} + \theta^b \mathbf{D}_{h,t} + e_{h,t} \tag{1.19}$$

Reduced Form with de-meaned rainfall

$$\text{Outcome}_{h,t} = \beta_2 \xi_{i,t-1} + \theta_i^a \mathbf{Y_t} + \psi C_{h,t} + \theta^b \mathbf{D}_{h,t} + e_{h,t} \tag{1.20}$$

Reduced Form with unpredicted rainfall

$$\text{Outcome}_{h,t} = \beta_3 \tilde{\varepsilon}_{h,t-1} + \theta_i^a \mathbf{Y_t} + \psi C_{h,t} + \theta^b \mathbf{D}_{h,t} + e_{h,t} \tag{1.21}$$

$$\xi_{h,t} = \beta^W \mathbf{W} + \psi C_{h,t} + \tilde{\varepsilon}_{h,t}$$

IV with de-meaned rainfall

$$\text{Outcome}_{h,t} = \beta_4 \widehat{\text{Income}}_{h,t-1} + \theta_i^a \mathbf{Y_t} + \psi C_{h,t} + \theta^b \mathbf{D}_{h,t} + e_{h,t} \tag{1.22}$$

$$\text{Income}_{h,t} = \gamma f(\xi_{i,t}) + \theta_i^a \mathbf{Y_t} + \psi C_{h,t} + \theta^b \mathbf{D}_{h,t} + \nu_{h,t}$$

[35]Papers using rainfall variation to examine education include Jensen (2000), Maccini and Yang (2009), and Björkman-Nyqvist (2013). Papers using rainfall variation to examine migration-related outcomes include Munshi (2003), Yang and Choi (2007), and Beegle et al. (2011). Rainfall and mortality are examined by Rose (1999) and Burgess and Donaldson (2010). The impacts of rainfall shocks on both birth rates and infant mortality are examined by Bhalotra (2010)

IV with unpredicted rainfall

$$\begin{aligned} \text{Outcome}_{h,t} &= \beta_5 \widehat{\text{Income}}_{h,t-1} + \theta_i^a \mathbf{Y_t} + \psi C_{h,t} + \theta^b \mathbf{D}_{h,t} + e_{h,t} \\ \text{Income}_{h,t} &= \gamma f(\tilde{\varepsilon}_{i,t}) + \theta_i^a \mathbf{Y_t} + \psi C_{h,t} + \theta^b \mathbf{D}_{h,t} + \nu_{h,t} \\ \xi_{h,t} &= \beta^W \mathbf{W} + \psi C_{h,t} + \tilde{\varepsilon}_{h,t} \end{aligned} \tag{1.23}$$

β_1 is the OLS estimate of the correlation between income and the outcome. Endogeneity issues with such a regression are common knowledge. β_2 and β_3 are reduced form estimates of $\xi_{i,t}$ and $\tilde{\varepsilon}_{i,t}$, respectively. Recall that the unknown difference between the variance of $\tilde{\varepsilon}$ and $\tilde{\xi}$ should temper the temptation to compare the magnitudes of β_2 and β_3. Cases where the sign of β_2 andβ_3 differ remains suggestive of potential bias. The magnitudes of IV results from demeaned and unpredictable rainfall shocks, β_4 and β_5, can be directly compared.

Although alternative functions to $f(\cdot)$ are discussed for unpredicted rainfall shocks, in comparing IV estimates identical functions are used in order to isolate differences driven by expectations. As discussed above, there is no consistent choice of functional form, $f(\xi_{i,t})$, for a rainfall shock in literature. Some papers force positive and negative rainfall shocks to have identical impacts on income, while others force opposite effects.[36] The main results presented in this paper tests follow popular variants in using the levels and squares of rainfall shocks. Similar results are obtained using another popular approach of indicators for rainfall events

[36]For examples, Jensen (2000) and Hidalgo et al. (2010) use identical impacts of positive and negative shocks. Jayachandran (2006) tests for and rejects similar impacts in her data. Paxson (1992) permits such differences. The differences between positive and negative shocks on income in ICRISAT data are significant, so this paper permits them.

beyond cutoffs in standard deviation. Kharif season agricultural income itself is defined as the value of the household's kharif season output less the cost of kharif season inputs.[37] Explicitly, the income shock examined here is defined as deviation from mean household income, or Income $\text{Shock}_{h,t} \equiv \overline{I}_h - I_{h,t}$.

$C_{h,t}$ is a vector of the percent of utilized land devoted to each crop. While this is helpful in isolating individual's rainfall expectations, it would serve as a poor instrument for income and hence needs to be carried through each stage of the IV estimates. This crop mix vector is also included in non-IV regressions for comparability. In an ideal setting, no additional controls would be necessary. Yet small samples either may differentially risk picking up time trends in income. District-level time trends, $\theta_{\mathbf{i}}\mathbf{Y_t}$, should account for differences driven by spurious correlation with time trends rather than adaptation-based differences.[38] A second issue is the presence of serial correlation in rainfall shocks. For households with few observations to calculate mean income, estimates may also capture correlation between mean income and a rainfall shock in another year. This is problematic because it causes the instrumented income measure to be mechanically correlated with rainfall-induced behavior in other periods. For this reason, households with

[37]Such a measure of income may be biased if input-values include the cost of expensive long-term investments. If long-term investments are more likely following good harvest years and good harvest years are more likely when rainfall is near expected values, then this should bias income to be lower in years following rainfall near expected values and higher in years where rainfall departs significantly from expected values.

[38]Data limitations mean that rainfall data is still measured at the district level. Hence rainfall should be viewed as a potentially noisy proxy for actual rainfall received by the individual household. As a silver lining, this means district level trends are sufficient to control for biases from small-sample trends in rainfall.

only two or three income observations have controls for rainfall in other observed income periods, denoted as the vector $\mathbf{D}_{h,t}$.

As discussed in Section 5, isolating an unpredictable rainfall shocks which is uncorrelated with average adaptive behavior is far easier at more geographically general levels. It was shown that the set of $\mathbf{W}$ variables sufficient at the district level are likely unable to capture all local information.While one could add an increasing number variables to the information set $\mathbf{W}$, such an approach is poorly founded and over-fitting predicted rainfall risks leaving little remaining identifying variation for unpredicted rainfall shocks. At the same time, because there is not household level data for most districts, some elements of $\mathbf{W}$ correlated with behavior at the district level may not be correlated with behavior in the micro data available in this setting.

Instead, the reduced form of potential $\mathbf{W}$ variables on adaptive behavior such as crop choice can help inform which predictors contain utilized information. Hierarchical stepwise analysis of potential $\mathbf{W}$ variables on the proportion of sown acreage devoted to each of the ten most popular household-level crops as adaptive behaviors confirms intuition that large numbers of lags and powers of ENSO measures are unnecessary in the micro data. Eliminating excess proxies which have no significant correlation with adaptive behavior reduces concerns of overfitting. The household level analysis uses quadratic district-specific time trends, last period's rainshock, government forecasts, and January-May ONI levels (each interacted with the village's state) as elements of $\mathbf{W}$. If these are insufficient to eliminate correlation between rainfall and adaptive behavior, the adaptive behavior itself can be used a final proxy for rainfall information. At the individual farm level, the chosen crop mix itself can be used as a proxy for $\mathbf{X}^{\mathrm{I}}$ information, as used by Kochar (1999) as a proxy for predictable profits. But note that using only the adaptive behavior as a predictor of rainfall expectations creates far too weak a proxy for this setting.

1.6.2 Results

Given the focus on the use of rainfall for identifying variation in agricultural income, comparisons between results based on demeaned rainfall, $\xi_{i,t}$, and unpredicted rainfall, $\tilde{\varepsilon}_{i,t}$, should begin with an examination of first-stage results. Table 1.13 shows the first stage results for the impact of rainfall on income. Results both with and without controls for crop portfolio are shown, because this control is not necessary for correct interpretation of first stage results but rather maintains consistent functional form across OLS, reduced form, and IV estimates. Above-mean rainfall is found to be associated with higher levels of income, although at a decreasing rate. The view that more rain uniformly increases income appears to be incorrect, as rainfall above predicted levels is not associated with increased levels of income. Further, Table 1.14 separates above and below mean rainfall by whether predicted rainfall was also above or below the mean. Consistent with income increases being driven by adaptation, above mean rainfall is only associated with increased income when above mean rainfall was also predicted.

Returning to Table 1.13, the main income impact of unpredicted rainfall is rainfall below predicted levels causing decreases in income, potentially at a decreasing rate. Reported F-statistics are for the four rainfall variables and exclude other controls which do not serve as instruments in the IV settings. Deviations from mean rainfall appear to be a slightly stronger instrument. This raises an important caveat that because unpredicted rainfall is a weaker instrument, there may be more measurement error in those generated income measures. This would also result in less significant, smaller IV estimates. However, the standard errors found on IV estimates using unpredicted rainfall are often smaller, suggesting measurement error may not be a larger issue in the unpredicted estimates than the tradional demeaned estimates. Additionally, unpredicted rainfall explains slightly more of household's variation in agricultural income.

Table 1.15 presents the main results of this section. Each cell represents a separate regression. Each row represents a different outcome variable. The first column shows the endogenous OLS estimate of income elasticity, β_1 from Equation 1.19. The second and third column examine the reduced form impact of demeaned and unpredicted rainfall on the same outcome, β_2 from Equation 1.20 and β_3 from Equation 1.21. The magnitudes of β_2 and β_3 are not comparable because a one-unit change in unpredicted rainfall does not necessarily correspond to a one-unit change in unexpected rainfall. The fourth and fifth columns show the IV estimates of β_4 from Equation 1.22 and β_5 from Equation 1.23. The magnitudes of β_4 and β_5 are comparable, and column 6 shows the p-value for the null hypothesis that the two IV estimates are equal.

For four of the the nine outcomes presented here, β_4 and β_5 are statistically different at standard confidence levels. This ratio of significant differences should not be taken as averages across all possible outcomes. Instead, researchers should consider potential types of adaptation and whether the expressions presented in Equation 1.4 and Equation 1.5 may be non-zero in their particular setting. Examples of such considerations are discussed below.

Migration seems to be a promising place to look for impacts of anticipatory adaptation. Matching the literature on migration, demeaned rainfall instruments find positive income shocks improve migration outcomes. These estimates find households are more likely to have a migrant, have more off them, have them travel farther, and make more money. Yet estimates of how income shocks affect the number of migrants in a household and the net income from those migrants are significantly different at standard confidence levels when based on unpredicted rainfall variation. In each case both estimates are positive and significant, but the estimates on predictable income shocks are larger.

These results are consistent with the possibility that expectations of higher future income at the end of year $t-1$ caused credit-constrained households to increase migration in year $t-1$, while unexpected rainfall shocks did not. Both rainfall shocks provided increased agricultural income enabling increases in migration in year t. Under this story, demeaned rainfall shocks would also be correlated with non-agricultural income shocks from migration in year $t-1$. This is comparable to the endogenous income effect discussed in Section Equation 1.4. Further, migration in year $t-1$ may be complementary to future migration through channels such as information about migration opportunities. Lowered costs of migration due to past migration would be a direct, non-income impact of adaptive behavior as described in Equation 1.5. It would also mean variation based on de-meaned rainfall is correlated with non-income impacts, violating the exclusion restriction of IV estimates. Further research would be necessary to isolate the extent to which differences are driven by complementary behavior versus expectable income shocks enabling further investment.

For another example of how predictability alters estimation, consider the example of the impact of income shocks in period $t-1$ on births which occur between period $t-1$ and t. Although it is not surprising for household size to increase in the face of positive income shocks, the length of gestation suggests that most births in this period were conceived before the income shock occurs. Hence increases in birth rate during this period correlated with positive income shocks can likely be assigned to expectations of increased income rather the realization of the income shock. Indeed, the third row of Table 1.15 shows positive expected income shocks cause a rise in births, while unexpected income shocks do not.

Whether outcomes are significantly different depends on whether adaptive behavior alters the outcome of interest, either directly or through endogenous income effects. Many estimates do not find significant differences. This does not imply no adaptation is occurring. Rather, it is learned that neither the adaptive behaviors nor the endogenous factors associated with these adaptations are correlated with the outcome of interest. Hence even similar estimates are informative about the type of adaptive behavior occurring, and the factors which drive these adaptations.

1.7 Conclusion

Rainfall and other weather events have long been employed as sources of exogenous and unpredictable identifying variation. Rainfall is most popularly used for variation in agricultural income, although rainfall's direct impacts on relative prices, health, and other outcomes are known to make it an imperfect instrumental variable. This paper's contribution is to document that even in a reduced form setting, estimates of the effect of rainfall deviations require significant reinterpretation. This is due to rainfall outcomes being partially predictable, meaning that anticipatory adaptation makes the subsequent impacts of rainfall endogenous.

Evidence of anticipatory adaption was presented in the form of Indian farmers adjusting their crop selections efficiently in advance of seasonal rainfall outcomes. When seasonal rainfall was one standard deviation below the mean, district-wide acreage of drought-resistant sorghum increased by almost 3%, while acreage of water-intensive rice decreased by over 1%. The average farmer among ICRISAT villages increasing their sown acreage of drought-resistant and water-intensive crops by as much as 10% each, consistent with poorer, smaller-scale farmers being more risk averse. If rainfall were unpredictable at the time of planting, farmers would not be able to make such anticipatory decisions about how much acreage to sow with various crops before observing rainfall outcomes. The interpretation of this behavior as due to seasonal weather expectations was robust to a variety of tests and alternative explanations.

Once such anticipatory adaptation has occurred, the impacts of rainfall on income and other outcomes become an endogenous function of agents' anticipatory adaptations, and may then be correlated with a wide variety of heterogeneous capital constraints, information constraints, or even risk preferences. Results show that attempts to reduce these issues by controlling for individuals' expectations require more detailed location-specific data than is likely available to the econometrician. Reduced form estimates should be interpreted as reflecting not only the impact of the physical rainfall event through income and relative prices, but also all impacts of anticipatory adaptation.

The impacts of anticipatory adaptation on estimates are not trivial. Because anticipatory adaptation begins at least as early as the start of the agricultural season, there is a large window of time for heterogeneous constraints or preferences to influence outcomes. Removing as much anticipatory adaptation as possible may result in significantly altered point estimates, reducing some estimates of income elasticity to half their previous size. The extent to which estimates are significantly altered depends on whether the adaptation or endogenous constraints associated with it are correlated with the outcome variable of interest. Hence examining differences between standard estimates and estimates which remove anticipatory adaptation is informative about the types of adaptation which occur.

1.8 Acknowledgments

Chapter 1, in part is currently being prepared for submission for publication of the material. Miller, Benjamin M. The dissertation author was the primary investigator and author of this material.

1.9 Tables

Table 1.1: Survey of Papers Using Rainfall for Identifying Variation 2011-2013

Estimation Strategy	Number
IV	6
Reduce Form	10
Other	3
Most Common Outcome Variables	Number
Household Consumption/Investment	6
Political Behavior	5
Aggregate Consumption/Revenue/Growth	3
Most Commonly Instrumented Variables	Number
Income Shock	3
Attendance	2
Most Commonly Studied Regions	Number
One or more African countries	10
India	5
Most Commonly Used Form of Rainfall	Number
Total or De-meaned Rainfall	10
Indicator for Above\Below Cutoff	6
First Difference	1
Additional Information	Number
Cites exogeneity or unpredictability to justify identification	8
Acknowledges non-mean expectations are possible	1
Tests or controls for serial correlation in rainfall	3
Includes Year-by-region FE	0
Includes Year FE	11
Includes Region FE	12

All papers published in 2011-2013 in the following 10 journals containing the phrase "rain" or "rainfall" or "precipitation" were examined: American Economic Journal: Applied Economics, American Economic Review (excluding P&P), Econometrica, International Economic Review, Journal of Development Economics, Journal of International Economics, Journal of Labor Economics, Journal of Political Economy, Quarterly Journal of Economics, Review of Economics and Statistics. Papers which used rainfall-driven variation of annual or finer frequency for identification of main results were included. Papers using mean rainfall over longer periods, using rainfall as a control, or only using rainfall variation for robustness checks are not included. Comment papers are not included. Fixed effects numbers only reflect papers with multiple time periods and/or regions.

Table 1.1: Survey of Papers Using Rainfall for Identifying Variation 2011-2013, continued

Common JEL Codes and Citations	Number
D - Microeconomics	9
E - Macroeconomics and Monetary Economics	2
F - International Economics	3
H - Public Economics	3
J - Labor and Demographic Economics	2
O - Economic Development, Technological Change, and Growth	13
Q - Agricultural and Natural Resource Economics; Environmental and Ecological Economics	2
"h-index": maximum number of papers h each with h or more citations	13

All papers published in 2011-2013 in the following 10 journals containing the phrase "rain" or "rainfall" or "precipitation" were examined: American Economic Journal: Applied Economics, American Economic Review (excluding P&P), Econometrica, International Economic Review, Journal of Development Economics, Journal of International Economics, Journal of Labor Economics, Journal of Political Economy, Quarterly Journal of Economics, Review of Economics and Statistics. Papers which used rainfall-driven variation of annual or finer frequency for identification of main results were included. Papers using mean rainfall over longer periods, using rainfall as a control, or only using rainfall variation for robustness checks are not included. Comment papers are not included. Fixed effects numbers only reflect papers with multiple time periods and/or regions.

Table 1.2: District-Level Summary Statistics

	Panel A: ICRISAT Meso, 1966-1999			
	mean	sd	min	max
Hectares of Rice Planted during Agricultural Year	39997.93	57676.18	0	255500
Hectares of Sorghum Planted during Kharif Season	48192.98	48451.93	100	198200
Hectares of Sorghum Planted during Agricultural Year	48257.89	48494.9	100	198200
Hectares of Maize Planted during Agricultural Year	17724.32	23732.49	0	144100
Kharif Season Rainfall (mm)	907.9472	249.5626	351.892	1946.572
Soil Type: USTALFUSTOLLS	.0351351	.1375343	0	.75
Soil Type: VERTISOLS	.1	.1790321	0	.8
Soil Type: VERTICSOILS	.8527027	.2128356	.2	1
Soil Type: INCEPTISOLS	.0121622	.0512293	0	.25
Number of District-Year Observations	1258			

	Panel B: IAC, 1956-1987			
	mean	sd	min	max
Hectares of Rice Planted during Agricultural Year	37554.22	52866.3	100	233300
Hectares of Sorghum Planted during Agricultural Year	55455.07	49543.99	200	198800
Hectares of Maize Planted during Agricultural Year	14455.24	19289.24	0	130300
Kharif Season Rainfall (mm)	946.6578	243.7576	351.892	1689.737
Number of District-Year Observations	1184			

Table 1.3: Forecast and ONI Summary Statistics

	Panel C: ONI & IMD Forecasts, 1950-2012			
	mean	sd	min	max
IMD Forecasts Above-Average Rainfall on Penninsula (Indicator)	.1147541	.32137	0	1
IMD Forecasts Below-Average Rainfall on Penninsula (Indicator)	.1311475	.340363	0	1
IMD Forecasts Above-Average Rainfall in NW India(Indicator)	.1803279	.3876509	0	1
IMD Forecasts Below-Average Rainfall in NW India (Indicator)	.1147541	.32137	0	1
Oceanic Niño Index, January Level	.0311475	1.089578	-2	2.3
Oceanic Niño Index, February Level	.0098361	1.018938	-1.9	2.2
Oceanic Niño Index, March Level	-.004918	.87358	-1.6	1.9
Oceanic Niño Index, April Level	.0032787	.6809227	-1.2	1.5
Oceanic Niño Index, May Level	.0065574	.560318	-1	1.2
Oceanic Niño Index, June Level	.0180328	.5264688	-.9	1
Oceanic Niño Index, July Level	.0344262	.5618082	-1.2	1.2
Oceanic Niño Index, August Level	.0491803	.6238651	-1.3	1.5
Oceanic Niño Index, September Level	.0409836	.7225826	-1.2	1.8
Oceanic Niño Index, October Level	.0409836	.8313397	-1.4	2.1
Oceanic Niño Index, November Level	.047541	.9529964	-1.6	2.3
Oceanic Niño Index, December Level	.0393443	1.058817	-1.9	2.4
Annual Observations		61		

Table 1.4: ICRISAT Micro Summary Statistics

	ICRISAT Micro, Gen I 1975-1985, Gen II 2001-2008				
	mean	sd	min	max	count
Percent of sown acreage devoted to Cotton	16.92762	27.42556	0	100	2864
Percent of sown acreage devoted to Pigeon Pea	14.2202	19.96824	0	100	2864
Percent of sown acreage devoted to Castor	8.405448	19.21192	0	100	2864
Percent of sown acreage devoted to Soybean	2.333735	11.98798	0	100	2864
Percent of sown acreage devoted to Sorghum	14.63662	20.69164	0	100	2864
Percent of sown acreage devoted to Rice	8.306733	21.29819	0	100	2864
Student in household starts school (Indicator)	.0407324	.1977066	0	1	2676
Student in household stops school (Indicator)	.0429746	.2028377	0	1	2676
Birth in household (Indicator)	.1057549	.3075811	0	1	2676
Death in household (Indicator)	.0478555	.2135087	0	1	2215
Household reports migrant worker (Indicator)	.2179177	.4130813	0	1	826
Number of migrant workers reported	.3837772	.8830681	0	6	826
Mean distance migrated (km)	37.21155	113.5588	0	1505	824
Total days of migrant labor	267.0172	237.6022	0	1355	348
Net Income from Migrants (Rupees)	13747.65	14843.7	302.5	104310	1493
Small Farm Households (Indicator)	.3526536	.4778795	0	1	2864
Medium Farm Households (Indicator)	.3159916	.4649908	0	1	2864
Large Farm Households (Indicator)	.3313547	.4707825	0	1	2864
Kharif Season Income (Output - Input, Rupees)	11536.7	58516.84	-97746.7	2344830	2864
Kharif Season Rainfall (mm)	592.1535	172.99	352.035	938.7	2864

The number of observations vary due to non-response and differences in survey structure between and within Generation I and Generation II of the ICRISAT Micro data. Migration data comes purely from Generation II.

Table 1.5: Evidence of Anticipatory Adaptation and Response to Econometrician's Rainfall Predictions: Main Results

	(1) Rice	(2) Rice	(3) Rice	(4) Sorghum	(5) Sorghum	(6) Sorghum	(7) Maize	(8) Maize	(9) Maize
Realized Kharif Rainfall Deviation	495.0*			-1305.4***			-217.2		
	(210.0)			(318.1)			(154.7)		
Predicted Kharif Rainfall Deviation		738.8**	738.8**		-1873.1***	-1872.4***		-254.3	-253.7
		(272.4)	(272.2)		(514.2)	(529.0)		(159.3)	(130.8)
Unpredicted Kharif Rainfall Deviation			12.04			-182.2			-144.7
			(186.2)			(294.0)			(243.6)
L.Realized Kharif Rainfall Deviation	1262.1***	1313.6***	1312.8***	724.6	594.9	606.7	250.2	233.2	242.6
	(205.0)	(213.9)	(207.0)	(821.2)	(790.2)	(800.2)	(157.6)	(169.7)	(161.2)
District FE	X	X	X	X	X	X	X	X	X
District-Specific Quadratic Year Trend	X	X	X	X	X	X	X	X	X
Residual Soil Moisture	X	X	X	X	X	X	X	X	X
N	1258	1258	1258	1258	1258	1258	1258	1258	1258
r2	0.995	0.995	0.995	0.969	0.969	0.969	0.986	0.986	0.986

Dependent variable is the hectares of non-HYV rice or non-HYV sorghum. Expected Kharif Rainfall Deviation is the standardized expected deviation of kharif season rainfall from mean rainfall over 1956-1999. Realized Kharif Rainfall Deviation is the actual deviation that occurs. Kharif season is the wet or monsoon season, defined here as June through September. Murphy-Topel standard errors clustered at the district level.

$^{*}p < 0.05,^{**} p < 0.01,^{***} p < 0.001.$

Table 1.6: Evidence of Anticipatory Adaptation and Response to Econometrician's Rainfall Predictions: Secondary Data Set

	(1) Rice	(2) Rice	(3) Rice	(4) Sorghum	(5) Sorghum	(6) Sorghum	(7) Maize	(8) Maize	(9) Maize
Realized Kharif Rainfall Deviation	722.3**			-1170.3***			-71.52*		
	(253.3)			(258.0)			(29.33)		
Predicted Kharif Rainfall Deviation		1219.8**	1219.5**		-1417.0***	-1416.9***		-102.5	-102.5
		(432.8)	(391.6)		(381.9)	(396.1)		(60.82)	(61.52)
Unpredicted Kharif Rainfall Deviation			148.4			-49.22			-39.35
			(273.8)			(481.4)			(48.22)
L.Realized Kharif Rainfall Deviation	669.1***	828.1**	832.9***	-1819.0***	-2084.8***	-2086.3***	-80.63**	-105.1**	-106.3**
	(181.6)	(261.0)	(236.0)	(275.3)	(336.8)	(336.9)	(28.31)	(36.50)	(37.06)
District FE	X	X	X	X	X	X	X	X	X
District-Specific Quadratic Year Trend	X	X	X	X	X	X	X	X	X
N	1147	999	999	1147	999	999	1147	999	999
r2	0.995	0.995	0.995	0.974	0.975	0.975	0.996	0.996	0.996

Dependent variable is the hectares of non-HYV rice or non-HYV sorghum. Expected Kharif Rainfall Deviation is the standardized expected deviation of kharif season rainfall from mean rainfall over 1956-1999. Realized Kharif Rainfall Deviation is the actual deviation that occurs. Kharif season is the wet or monsoon season, defined here as June through September. All regressions include district fixed effects and district-specific quadratic time trends to control for cross-sectional differences and trends in crop choices. Note that IAC data comes from an overlapping but slightly earlier time period than ICRISAT data, does not contain soil × lagged rainfall controls, and sorghum grown in the kharif season cannot be separated from sorghum grown in the rabi season. Murphy-Topel standard errors clustered at the district level.

$^{*}p < 0.05,^{**}p < 0.01,^{***}p < 0.001$.

Table 1.7: Evidence of Anticipatory Adaptation and Response to Econometrician's Rainfall Predictions: ICRISAT, Weaker Predictions

	(1) Rice	(2) Rice	(3) Rice	(4) Sorghum	(5) Sorghum	(6) Sorghum	(7) Maize	(8) Maize	(9) Maize
Realized Kharif Rainfall Deviation	495.0*			-1305.4***			-217.2		
	(210.0)			(318.1)			(154.7)		
Predicted Kharif Rainfall Deviation		423.2	396.5		1810.1	1899.8		801.7	819.9
		(309.2)	(310.6)		(1306.3)	(1117.6)		(2753.6)	(3202.3)
Unpredicted Kharif Rainfall Deviation			507.2*			-1700.9***			-345.1
			(215.0)			(463.7)			(1099.3)
L.Realized Kharif Rainfall Deviation	1262.1***	1287.8***	1253.7***	724.6	885.8	1000.2	250.2	316.1	339.3
	(205.0)	(218.5)	(202.2)	(821.2)	(783.8)	(815.1)	(157.6)	(675.9)	(850.6)
District FE	X	X	X	X	X	X	X	X	X
District-Specific Quadratic Year Trend	X	X	X	X	X	X	X	X	X
Residual Soil Moisture	X	X	X	X	X	X	X	X	X
N	1258	1258	1258	1258	1258	1258	1258	1258	1258
r2	0.995	0.995	0.995	0.969	0.968	0.969	0.986	0.986	0.987

Dependent variable is the hectares of non-HYV rice or non-HYV sorghum. Expected Kharif Rainfall Deviation is the standardized expected deviation of kharif season rainfall from mean rainfall over 1956-1999. Realized Kharif Rainfall Deviation is the actual deviation that occurs. Kharif season is the wet or monsoon season, defined here as June through September. Murphy-Topel standard errors clustered at the district level.

$^{*}p < 0.05, ^{**}p < 0.01, ^{***}p < 0.001$.

Table 1.8: Evidence of Anticipatory Adaptation and Response to Econometrician's Rainfall Predictions: Secondary Data Set, Weaker Predictions

	(1) Rice	(2) Rice	(3) Rice	(4) Sorghum	(5) Sorghum	(6) Sorghum	(7) Maize	(8) Maize	(9) Maize
Realized Kharif Rainfall Deviation	722.3**			-1170.3***			-71.52*		
	(253.3)			(258.0)			(29.33)		
Predicted Kharif Rainfall Deviation		537.1	669.1*		937.1	760.7		-18.14	-31.75
		(277.6)	(313.4)		(866.9)	(842.1)		(86.42)	(86.38)
Unpredicted Kharif Rainfall Deviation			803.8**			-1074.5***			-82.86*
			(295.8)			(228.0)			(40.75)
L.Realized Kharif Rainfall Deviation	669.1***	796.3***	808.0***	-1819.0***	-1875.8***	-1891.4***	-80.63**	-99.45**	-100.6**
	(181.6)	(215.1)	(228.1)	(275.3)	(299.0)	(312.2)	(28.31)	(31.94)	(34.05)
District FE	X	X	X	X	X	X	X	X	X
District-Specific Quadratic Year Trend	X	X	X	X	X	X	X	X	X
N	1147	999	999	1147	999	999	1147	999	999
r2	0.995	0.994	0.995	0.974	0.974	0.975	0.996	0.996	0.996

Dependent variable is the hectares of non-HYV rice or non-HYV sorghum. Expected Kharif Rainfall Deviation is the standardized expected deviation of kharif season rainfall from mean rainfall over 1956-1999. Realized Kharif Rainfall Deviation is the actual deviation that occurs. Kharif season is the wet or monsoon season, defined here as June through September. All regressions include district fixed effects and district-specific quadratic time trends to control for cross-sectional differences and trends in crop choices. Note that IAC data comes from an overlapping but slightly earlier time period than ICRISAT data, does not contain soil × lagged rainfall controls, and sorghum grown in the kharif season cannot be separated from sorghum grown in the rabi season. Murphy-Topel standard errors clustered at the district level.

$^{*}p < 0.05,^{**}p < 0.01,^{***}p < 0.001$.

Table 1.9: Evidence of Anticipatory Adaptation and Response to Econometrician's Rainfall Predictions: Rolling Means Over Shorter Time Periods

	(1) Rice	(2) Rice	(3) Rice	(4) Sorghum	(5) Sorghum	(6) Sorghum	(7) Maize	(8) Maize	(9) Maize
Realized Kharif Rainfall Deviation	488.9*			-1118.9***			-183.7		
	(218.8)			(295.1)			(131.1)		
Predicted Kharif Rainfall Deviation		696.3**	696.1**		-1713.2***	-1713.4***		-238.4	-238.1
		(260.7)	(260.1)		(506.6)	(503.5)		(148.9)	(133.1)
Unpredicted Kharif Rainfall Deviation			81.35			50.93			-76.63
			(216.0)			(270.1)			(191.0)
L.Realized Kharif Rainfall Deviation	1402.4***	1447.3***	1441.7***	716.8	607.5	604.0	278.5	262.8	268.1
	(228.9)	(235.2)	(227.0)	(807.1)	(788.1)	(785.9)	(141.3)	(154.5)	(142.3)
District FE	X	X	X	X	X	X	X	X	X
District-Specific Quadratic Year Trend	X	X	X	X	X	X	X	X	X
Residual Soil Moisture	X	X	X	X	X	X	X	X	X
N	1258	1258	1258	1258	1258	1258	1258	1258	1258
r2	0.995	0.995	0.995	0.968	0.969	0.969	0.986	0.986	0.986

Dependent variable is the hectares of non-HYV rice or non-HYV sorghum. Expected Kharif Rainfall Deviation is the standardized expected deviation of kharif season rainfall from a 20-year rolling mean of rainfall. Realized Kharif Rainfall Deviation is the actual deviation that occurs. Kharif season is the wet or monsoon season, defined here as June through September. Murphy-Topel standard errors clustered at the district level.

$^{*}p < 0.05, ^{**} p < 0.01, ^{***} p < 0.001.$

Table 1.10: Evidence of Anticipatory Adaptation and Response to Econometrician's Rainfall Predictions: Robust to Delayed Planting Decisions

	(1) Rice	(2) Rice	(3) Rice	(4) Sorghum	(5) Sorghum	(6) Sorghum	(7) Maize	(8) Maize	(9) Maize
Realized Kharif Rainfall Deviation	356.6*			-1028.5***			-182.3		
	(170.6)			(277.6)			(122.3)		
Predicted Kharif Rainfall Deviation		604.6**	604.8**		-1565.2**	-1564.5**		-266.1	-265.9
		(234.0)	(233.9)		(504.2)	(516.3)		(168.0)	(147.2)
Unpredicted Kharif Rainfall Deviation			-37.04			-178.0			-49.56
			(168.0)			(247.6)			(163.5)
L.Realized Kharif Rainfall Deviation	1849.0	1767.2	1762.1	4928.1*	5140.4	5115.9*	933.1	969.3	962.4
	(1823.8)	(1918.9)	(1939.9)	(2140.3)	(2625.5)	(2606.8)	(678.5)	(935.3)	(707.1)
District FE	X	X	X	X	X	X	X	X	X
District-Specific Quadratic Year Trend	X	X	X	X	X	X	X	X	X
Residual Soil Moisture	X	X	X	X	X	X	X	X	X
N	1258	1258	1258	1258	1258	1258	1258	1258	1258
r2	0.995	0.995	0.995	0.968	0.968	0.968	0.986	0.986	0.986

Dependent variable is the hectares of non-HYV rice or non-HYV sorghum. Expected Kharif Rainfall Deviation is the standardized expected deviation of July-September rainfall from its mean over 1956-1999. Realized Kharif Rainfall Deviation is the actual deviation that occurs. Murphy-Topel standard errors applied to all but residual regressors. Murphy-Topel standard errors clustered at the district level.

$^{*}p < 0.05, ^{**}p < 0.01, ^{***}p < 0.001$.

Table 1.11: Household-Level Adaptation and Response to Econometrician's Rainfall Predictions: Main Results 1

	(1) Cotton	(2) Cotton	(3) Cotton	(4) PigeonPea	(5) PigeonPea	(6) PigeonPea	(7) Castor	(8) Castor	(9) Castor
Realized Kharif Rainfall Deviation	5.294			-12.39***			-13.12**		
	(5.285)			(3.580)			(4.840)		
Predicted Rainfall Deviation		3.680	3.996		-6.267*	-7.337*		-14.42**	-14.46**
		(5.972)	(5.952)		(3.002)	(3.063)		(5.310)	(5.326)
Unpredicted Rainfall Deviation			16.06*			-54.31***			-2.005*
			(6.938)			(15.41)			(0.951)
HH FE	X	X	X	X	X	X	X	X	X
Village-Specific Quadratic Year Trend	X	X	X	X	X	X	X	X	X
Village-Specific L.Realized Rainshock	X	X	X	X	X	X	X	X	X
N	1867	1867	1867	1867	1867	1867	1867	1867	1867
r2	0.346	0.345	0.346	0.113	0.107	0.126	0.389	0.389	0.390

Dependent variable is the percentage of household's farming acreage devoted to the given crop. Predicted Kharif Rainfall Deviation is the expected deviation of kharif season rainfall from the rolling mean of the prior 20 years. Realized Kharif Rainfall Deviation is the actual deviation that occurs. Kharif season is the wet or monsoon season, defined here as June through September. All regressions include household fixed effects and village-specific quadratic time trends to control for cross-sectional differences and trends in crop choices. Standard errors are clustered at the household level.

$^{*}p < 0.05,^{**} p < 0.01,^{***} p < 0.001.$

Table 1.12: Household-Level Adaptation and Response to Econometrician's Rainfall Predictions: Main Results 2

	(1) Soybean	(2) Soybean	(3) Soybean	(4) Sorghum	(5) Sorghum	(6) Sorghum	(7) Rice	(8) Rice	(9) Rice
Realized Kharif Rainfall Deviation	0.0924			-2.439			11.39**		
	(1.022)			(2.948)			(3.943)		
Predicted Rainfall Deviation		1.524	1.324		0.370	-0.0655		13.02**	12.98**
		(1.582)	(1.498)		(2.756)	(2.778)		(4.310)	(4.324)
Unpredicted Rainfall Deviation			-10.13			-22.12			-1.781
			(5.705)			(12.24)			(2.191)
HH FE	X	X	X	X	X	X	X	X	X
Village-Specific Quadratic Year Trend	X	X	X	X	X	X	X	X	X
Village-Specific L.Realized Rainshock	X	X	X	X	X	X	X	X	X
N	1867	1867	1867	1867	1867	1867	1867	1867	1867
r2	0.677	0.677	0.678	0.044	0.044	0.049	0.303	0.305	0.305

Dependent variable is the percentage of household's acreage actively farmed that year which was devoted to the given crop. Predicted Kharif Rainfall Deviation is the expected deviation of kharif season rainfall from the rolling mean of the prior 20 years. Realized Kharif Rainfall Deviation is the actual deviation that occurs. Kharif season is the wet or monsoon season, defined here as June through September. All regressions include household fixed effects and village-specific quadratic time trends to control for cross-sectional differences and trends in crop choices. Standard errors are clustered at the household level.

$^{*}p < 0.05, ^{**}p < 0.01, ^{***}p < 0.001$.

Table 1.13: First Stage: The Impact of Rainfall on Income

	De-meaned Rainfall, ξ		Unpredicted Rainfall, $\tilde{\varepsilon}$	
	(1)	(2)	(3)	(4)
	Income Shock	Income Shock	Income Shock	Income Shock
Rainfall Shock (+)	18.59***	27.53***	-8.284	-9.500
	(6.721)	(7.020)	(16.98)	(13.46)
Rainfall Shock^2 (+)	-13.62***	-23.59***	-18.17	-21.29*
	(5.205)	(6.013)	(13.43)	(11.99)
Rainfall Shock (-)	5.967	-1.363	-25.25*	-33.05**
	(4.304)	(4.054)	(15.33)	(12.98)
Rainfall Shock^2 (-)	6.611	17.72***	30.55	53.45***
	(5.208)	(5.641)	(22.59)	(20.28)
District Time Trend	X	X	X	X
Serial Correlation Controls	X	X	X	X
Crop Portfolio		X		X
Observations	2305	2305	2305	2305
R^2	0.00589	0.0176	0.0111	0.0233
F-Statistic	7.507	8.075	4.958	5.086

Dependent variable is deviation from mean household income. The first two columns examine deviation from mean rainfall, while the second two columns examine deviation from predicted rainfall. Deviations from mean rainfall are standardized so that a one unit change is associated with a one standard deviation change in total rainfall. Unpredicted rainfall is the residual from the first stage regression described in Equation 1.21. Hence a one-unit change in unpredicted rainfall may be associated with more than a one standard deviation change in total rainfall, so the magnitudes are not directly comparable. $(\cdot)$ indicates positive or negative rainfall deviations. Reported F-statistics are for the four rainfall variables, and do not include controls.

$^{*}p < 0.1, ^{**} p < 0.05, ^{***} p < 0.01$.

Table 1.14: First Stage: The Impact of Rainfall on Income, by Prediction

	(1) Income Shock	(2) Income Shock
De-meaned Rainfall (+)	3.875	4.231
Predicted Rainfall (-)	(12.51)	(12.05)
De-meaned Rainfall (+)	38.31***	70.24**
Predicted Rainfall (+)	(10.87)	(28.31)
De-meaned Rainfall^2 (+)	4.214	0.723
Predicted Rainfall (-)	(12.85)	(12.03)
De-meaned Rainfall^2 (+)	-33.77***	-57.70***
Predicted Rainfall (+)	(8.106)	(20.02)
De-meaned Rainfall (-)	13.29*	-8.169
Predicted Rainfall (-)	(7.315)	(6.886)
De-meaned Rainfall (-)	17.14**	10.24
Predicted Rainfall (+)	(8.460)	(8.153)
De-meaned Rainfall^2 (-)	2.917	31.48***
Predicted Rainfall (-)	(8.743)	(11.47)
De-meaned Rainfall^2 (-)	-12.58	-2.284
Predicted Rainfall (+)	(14.29)	(14.44)
District Time Trend	X	X
Serial Correlation Controls	X	X
Crop Portfolio		X
Observations	2305	2305
R^2	0.00860	0.0213

Dependent variable is deviation from mean household income. Both columns examine deviation from mean rainfall, and are standardized so that a one unit change is associated with a one standard deviation change in total rainfall. Each measure of deviation from mean rainfall has been interacted with an indicator for whether predicted rainfall was above or below 0.

$^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01.$

Table 1.15: Comparison of Reduced Form & Instrumental Variable Estimates

	(1) OLS	(2) Reduced Form Demeaned	(3) Reduced Form Unpredicted	(4) IV Demeaned	(5) IV Unpredicted	p(IV(ξ) = IV($\hat{\varepsilon}$))
Enroll Student	-0.00468	0.971	0.650	-0.161	-0.207	0.7867
	(0.0140)	(0.877)	(1.413)	(0.174)	(0.140)	
Withdraw Student	0.0108	0.984	-0.778	-0.0512	0.0732	0.4482
	(0.0114)	(0.902)	(1.485)	(0.173)	(0.117)	
Birth	-0.00596	0.377	2.849	1.074**	-0.0475	0.0119
	(0.00569)	(1.268)	(1.807)	(0.417)	(0.117)	
Death	0.0170	0.395	-2.318	0.289	0.260	0.9225
	(0.0170)	(1.030)	(1.526)	(0.256)	(0.272)	
Has Migrant	0.0190	28.65***	14.96***	0.831***	0.420**	0.0526
	(0.0138)	(3.334)	(5.779)	(0.271)	(0.210)	
Number of Migrants	0.000598**	0.380***	0.186**	0.0123***	0.00489	0.0923
	(0.000268)	(0.0808)	(0.0941)	(0.00423)	(0.00434)	
Mean Distance Migrated	0.0188	40.55***	-3.123	2.272***	1.615**	0.2953
	(0.0422)	(11.64)	(16.19)	(0.875)	(0.776)	
Days of Migrant Labor	0.371	60.35**	109.1**	2.802	0.560	0.2270
	(0.296)	(23.64)	(44.18)	(2.075)	(0.854)	
Net Migrant Income	6.911	16804.6***	-5021.9	922.0***	288.7**	0.0046
	(25.19)	(5354.9)	(5966.7)	(289.4)	(133.6)	
District Time Trend	X	X	X	X	X	
Serial Correlation	X	X	X	X	X	
Crop Portfolio	X	X	X	X	X	

Each cell represents β_i from Equation 1.19 through Equation 1.23. The dependent variable is listed on the far left of each row. The first column represents OLS estimates, the next two columns represent reduced form estimates, the next two represent IV estimates, and the final column shows the p-value for the null hypothesis that the two IV estimates are equal. Because a one-unit change in unpredicted rainfall is not necessarily equivalent to a one-unit change in demeaned rainfall, the magnitudes are not directly comparable.

$^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01.$

Chapter 2

The Salience of Complex Tax Changes: Evidence From the Child and Dependent Care Credit Expansion

2.1 Introduction

Preparing a U.S. personal income tax return can be complicated and time consuming. The IRS reports that 56 percent of taxpayers hired a paid tax professional to complete their federal personal income tax return in 2012.[1] Slemrod and Bakija (2008) estimate that taxpayers spend an average of 26 hours per year performing the recordkeeping and paperwork to complete their federal and state personal income tax returns. The complexity of the tax code makes it difficult for taxpayers to understand the tax implications of their economic choices.

The literature on tax salience, including papers by Duflo et al. (2006), Gallagher and Muehlegger (2011), Finkelstein (2009), and Chetty et al. (2009), concludes that when the financial incentives of a tax change are not highly salient, the tax change induces a smaller response than an otherwise equivalent price change. This paper adds some nuance to this literature by considering a complex tax change that consists of both a direct tax impact and indirect tax interactions. We propose a simple behavioral model in which taxpayers respond to the direct impact of a complex tax change and do not respond to the less salient interactions with other elements of the tax code. We then examine evidence of such behavior in taxpayers response to the 2003 expansion of the Child and Dependent Care Credit (CDCC).

The CDCC is an important child-care subsidy that likely influences the amount many families choose to spend on child care through both the quantity

[1]See the IRS Statistics of Income Bulletin Winter 2015, Selected Historical and Other Data Tables 1 and 22a.

and quality margins.[2] For taxpayers focusing only on the 2003 change to the credit itself (the direct impact), the expansion of the CDCC would have appeared as an unambiguous decrease in the after-tax price of child care. However, other tax changes, particularly the simultaneous expansion of the Child Tax Credit (CTC), interacted with the CDCC expansion to often reduce or even eliminate the child-care subsidy. Using individual-level survey data from before and after the CDCC expansion to employ a difference-in-differences estimation strategy, we present evidence showing that taxpayers increased their expenditure on child care in response to the expansion of the CDCC regardless of whether the actual after-tax price of child care increased or decreased.

Taxpayers in the model we present in Section II have limited attention and may choose to only focus on the direct impact of a change to a single tax provision rather than the actual financial implications of the change when the full tax code is considered in its entirety. Focusing on a part of the tax code rather than the whole is similar to what Liebman and Zeckhauser (2004) call spotlighting.[3] Whether individuals consider interactions between provisions of the tax code is distinct but complementary to the literature on whether individuals respond to average or marginal tax rates, such as De Bartolome (1995) and Ito (2014). Acquiring information about the change to the CDCC is low cost; figuring out how the

[2]See Blau and Robins (1988), Connelly (1992), Averett et al. (1997), Blau (2003), and Herbst (2010).

[3]Liebman and Zeckhauser (2004) define spotlighting as responding to the instantaneous payoff in the current period without considering the effects for the remainder of the accounting period. Here, we are using this term to describe taxpayers who respond to the direct implications of a single provision of the tax code without considering how their behavior affects total tax liability.

CDCC interacts with the rest of the tax code is far more costly. Taxpayers have access to all required information, but the effort needed to compute after-tax prices may lead rational taxpayers to adopt spotlighting behavior.

The rest of the paper proceeds as follows. Section II presents a model of spotlighting behavior with respect to the personal income tax. Section III provides a description of the Child and Dependent Care Credit, its 2003 expansion, and interaction with the Child Tax Credit. Section IV describes the data and methodology. Section V describes the results.

2.2 Model

Many deductions and credits have been introduced into the personal income tax code by lawmakers interested in encouraging certain activities. If the government wants to provide a subsidy for some activity it may be easier and more administratively efficient to introduce a targeted deduction or credit into the personal income tax system than to create an entirely new system to provide the subsidy.[4] But, as more targeted deductions and credits piggyback on the personal income tax, these tax provisions interact with each other and at times cause incentives to diverge from what was originally intended.

For taxpayers to make consumption decisions optimally, they must know what after-tax prices they face. Taxpayers who gather only the information re-

[4]Piggybacking a proposed subsidy or transfer payment onto the personal income tax system may not be efficient if the targeted beneficiaries of the proposed subsidy do not generally file tax returns. For example, the tax system would probably not be a good delivery mechanism for disability benefits.

quired to claim the relevant deductions and credits, but do not understand how they interact, may calculate a "naive" after-tax price that is far different than a "nuanced" after-tax price which considers the interactions.

We model the personal income tax as a function

$$\text{Tax} = f(y, X, \tau_1(y, X, Z), \ldots, \tau_n(y, X, Z)) \tag{2.1}$$

which depends on the taxpayers income, y, family size and other taxpayer characteristics, X, and n credits or deductions given by $\tau_i(y, X, Z)$, where Z denotes other taxpayer characteristics that influence the value of specific credits or deductions. The complexity of the function $f(\cdot)$ is primarily due to the fact that the credits and deductions interact with each other as well as with y and X. However, each of the individual credits and deductions are generally simple functions with few inputs.

Suppose that to encourage a specific action or to reduce the tax burden for a group of taxpayers, a particular tax credit is increased from $\tau_i(y, X, Z)$ to $\tau_i'(y, X, Z)$. The literature gives two explanations for why we observe a smaller aggregate response to a tax change than to an equivalent price change. First, some taxpayers are inattentive and may not realize that the particular tax provision has changed (a type of price misperception) and thus will not respond. Second, taxpayers who observe the change may believe that calculation and adjustment costs will be greater than the utility gain from the optimal response and thus choose to not respond to the tax change. We offer a third explanation which we call spotlighting behavior.

Taxpayers engaged in spotlighting behavior use an easy (low utility cost) way to approximate the effect of the tax change, holding all other factors constant:

$$\Delta \text{Tax}_S \approx -\tau_i'(y, X, Z) + \tau_i(y, X, Z) \tag{2.2}$$

where the s subscript denotes the use of the spotlighting approximation. An increase in the tax credit from $\tau_i(y, X, Z)$ to $\tau_i'(y, X, Z)$ often causes a proportional reduction in the tax liability which provides justification for spotlighting behavior. However, this is not always the case. The change in tax liability depends on a more nuanced understanding of how the tax credit interacts with the other arguments of the tax function. Given full information including end of year income, the change in tax liability from a change to tax provision i is given by:

$$\Delta \text{Tax} = f(y, X, \ldots, \tau_i'(y, X, Z), \ldots) - f(y, X, \ldots, \tau_i(y, X, Z), \ldots). \tag{2.3}$$

For example, suppose that $\tau_i(y, X, z)$ is a tax credit that provides partial reimbursement of expenditure on a specific good where z denotes expenditure on that good. Taxpayers using spotlighting would approximate the after-tax price of this good as:

$$\text{After-tax price}_S \approx p\left(1 - \frac{\partial \tau_i(y, X, z)}{\partial z}|_{z=z^*}\right) \tag{2.4}$$

where p is the pre-tax price of the good and z^* is the chosen level of expenditure. However, the actual after-tax price of the good is expressed as:

$$\text{After-tax price} = p\left(1 + \frac{\partial f(y, X, \ldots, \tau_i(y, X, Z), \ldots)}{\partial z}|_{z=z^*}\right). \tag{2.5}$$

Spotlighting may generally provide a good approximation of the after-tax price. It is likely that small deviations from the frictionless (no price misperception) optimum due to spotlighting cause only a small reduction in utility as in Chetty (2012). However, in situations with important interactions like the 2003 expansion of the CDCC, spotlighting can lead to a large misperception of the after-tax price. The low salience of interactions can lead to a large deviation from the frictionless optimum even when the direct financial implications are salient.

2.3 Child and Dependent Care Credit

2.3.1 Historical Background

The Child and Dependent Care Credit (CDCC) began in 1954 as an itemized deduction for work-related child-care expenses. Prior to this tax provision, the courts had ruled that child-care expenses were not deductible (Smith v. Commissioner, 1940). The itemized deduction was limited to households making less than $4,500 annually and was limited to $600 in total child-care expenses. An update to the deduction in 1964 increased these limits, but the value of the deduction was still quite small given the low marginal tax rates in this range of the income distribution. In practice, few households claimed the deduction as only those that itemized their deductions were eligible.

In 1971, the deduction's income ceiling tripled and the maximum allowable deduction increased to $4,800. However, this did little to increase the number of households that benefited, so in 1976, Congress replaced the child-care deduction with a credit. The credit value was set at 20 percent of qualified expenses, up to $2,000 per child, and the income cap was removed. As a credit, the benefits were no longer linked to itemizing, so in theory, households at any income level could receive the subsidy. But as a non-refundable credit, CDCC benefits remained limited to households with tax liability, excluding many low-income households.

In 1981, the 20 percent rate was changed to a schedule starting at 30 percent and then moving down to 20 percent in steps occurring at specific income levels. The limit was increased to $2,400 of qualified child-care expenses per child.[5] There were no changes to the CDCC from 1981 until 2003, which, because it is not inflation indexed, caused its value to taxpayers to decline substantially.

In 2001, Congress increased the qualifying expenses limit to $3,000 per child and increased the credit rate schedule for low-income families.[6] Though passed in 2001, the CDCC expansion was not scheduled to take effect until the beginning of 2003. As shown in Panel (a) of Figure 1, the CDCC credit rate schedule only increased for taxpayers with an adjusted gross income (AGI) below $43,000.[7]

[5]To qualify, the child care must enable parents in the household to work or look for work. The care provider cannot be a parent or an older sibling of the child. Taxpayers who participate in a dependent care assistance plan (childcare flexible spending account plan) through their employer are only eligible to claim child-care expenditure for the CDCC that is not paid out of the flex spending account, and this is limited to the CDCC max. A flex spending plan allows an employee to place up to $5,000 of pre-tax income into an account for child care expenses.

[6]The Economic Growth and Tax Relief Reconciliation Act of 2001 increase the maximum Child and Dependent Care Credit to 35 percent of child-care expenditure (from 30 percent) of up to $3,000 (from $2,400) for one child and of up to $6,000 (from $4,800) for two or more children. The phase-out of the credit rate was moved to begin at $15,000 of adjusted gross income (from $10,000).

[7]Married couples can only claim the CDCC if both spouses are working (or if the non-working spouse is a student or disabled) and the amount of child care expenses used in calculating the credit is limited to the amount of earned income of the lesser-earning spouse.

2.3.2 Interaction with the Child Tax Credit

The Child Tax Credit (CTC) is best described as a lump-sum transfer to taxpayers with children, while the Child and Dependent Care Credit (CDCC) is a partial reimbursement of child-care expenses. As mentioned above, the CDCC is a non-refundable credit, meaning that only taxpayers with tax liability benefit. In contrast, the CTC is refundable, meaning that taxpayers without remaining tax liability can still benefit. The refundable portion of the Child Tax Credit is called the Additional Child Tax Credit (ACTC). Taxpayers with no remaining tax liability who have not yet claimed the full value of the CTC can claim the remaining amount through the ACTC. However, prior to 2008, the ACTC was limited for low-income taxpayers.[8] For example, in 2003 the refund was limited to 10 percent of the taxpayers earned income in excess of $10,500. When this ACTC constraint binds, the taxpayer is not able to claim the full value of the CTC.

In 2002, the year before the CDCC expansion, the CTC provided a credit of $600 per child to taxpayers with children. At the time, the U.S. was experiencing a mild recession. With the primary motivation of stimulating the economy through advanced tax refunds, the Jobs Growth and Tax Relief Reconciliation Act of 2003 increased the CTC to $1,000 per child and provided advance tax refund checks of $400 per child (the amount of the increase in the CTC).

The timing of the CTC increase happened to coincide with the expansion

[8] Prior to 2001, only taxpayers with three or more children could receive the ACTC, and the ACTC was limited to their payroll tax liability.

of the CDCC, even though the CDCC expansion had been passed two years earlier. Taxpayers with children first appear to have received both a decrease in the marginal cost of child care through the CDCC and a lump sum transfer from the increase of the Child Tax Credit. However, taxpayers with insufficient tax liability did not fully benefit from the CDCC and CTC increases. The CDCC appears before the CTC on the tax form (see Figure 2). As a result, for some, the increase in the CDCC reduced the amount of tax that was left to be claimed for the CTC, which in turn shifted CTC benefits to the ACTC. But as soon as the income constraint on the ACTC became binding, any benefits from claiming additional child-care expenses through the CDCC were offset by an equivalent decrease in the CTC value and no change in ACTC value. In addition, the Economic Growth and Tax Relief Reconciliation Act of 2001 reduced tax rates and increased the standard deduction causing there to be even less tax liability for the non-refundable CDCC to soak up.

As soon as the income constraint on the ACTC becomes binding, any benefits from claiming additional child-care expenses through the CDCC were offset by an equivalent decrease the CTC value and no change in ACTC value. Because the final tax liability and refunds for these taxpayers were not affected by the amount of child-care expenditure claimed, the marginal subsidy on child care became zero. This is illustrated in Panel (b) of Figure 1 for a single-parent household with two children. This particular issue affects few taxpayers today, as the ACTCs income constraint has been significantly relaxed.[9]

The CDCC interaction with the CTC was not obvious to taxpayers. Using tax preparation software could help the taxpayer figure out the subsidy rate, but only if the taxpayer entered the information several times with different levels of child-care expenditure, and then compared the resulting tax liability or refund. Performing this type of hypothetical calculation is probably not common. While using tax preparation software was unlikely to result in taxpayers gaining a more nuanced understanding of the subsidy rate, it may have increased awareness of the change to the CDCC as several leading brands of tax preparation software ask specifically about child-care expenditure and give the value of the CDCC reported on the 1040 form. Because it focuses attention on the value reported on the 1040 form, tax preparation software may have increased the use of spotlighting by taxpayers.[10]

We are not aware of any evidence regarding the extent to which members of Congress understood that other changes in the tax code after 2001, including the CTC increase, would reduce the value of the CDCC expansion for low-income taxpayer. The extent of Congress awareness is not important to the identification strategy in this paper, and we do not wish to suggest it was the intention of Congress to leave low-income taxpayers with a reduced child-care subsidy rather

[9]The income constraint was partially relaxed in 2008. By 2009, the ACTC reached its present constraint of being limited to 15 percent of income in excess of $3,000.

[10]In 2003, 43 percent of personal income tax returns were filed electronically and most of these returns were prepared using tax preparation software (some were prepared by tax professionals). By 2008, 67 percent of returns were filed electronically (see the IRS Statistics of Income Bulletin Winter 2015, Table 1).

than the legislated increase.

2.3.3 Response of Child Care to Child-Care Subsidies

The literature shows that taxpayers respond to a reduction in the price of child care by purchasing more child care. Blau and Robins (1988) provide direct evidence in a model of labor supply, fertility, and child-care expenditure where the price variation comes from a child-care subsidy. Other papers including Connelly (1992), Averett et al. (1997), and Herbst (2010) examine the responsiveness indirectly through a change in the labor force participation of mothers with young children, under the assumption that these working mothers are consumers of child care. Blau (2003) surveys the literature on the elasticity of employment with respect to the price of child care and finds estimates ranging from 0.06 to -1.26.

How taxpayers respond to the 2003 expansion of the CDCC depends on their perception of how the after-tax cost of child care was affected. All else equal, the child-care expenditure decisions of taxpayers who are primarily ignorant of the 2003 CDCC expansion should remain unchanged. Taxpayers who primarily use the spotlighting method should increase their child-care expenditure in response to an increase in the “naive” measure of the value of the CDCC. Taxpayers who account for interactions between elements of the tax code should increase or decrease their child-care expenditure in response to a “nuanced” measure which considers interactions between the CDCC and other elements of the tax code. If there are a substantial number of both fully-informed taxpayers and those who are spotlighting, then we would expect to see a response to both the naive and the nuanced change in the value of the CDCC.

2.4 Data and Empirical Strategy

2.4.1 Data

We use data from the diary portion of the U.S. Bureau of Labor Statistics Consumer Expenditure Survey (CES). Each survey participant records all household expenditures for a one-week period in a provided diary. This diary is collected at the end of the week and an interview is conducted to obtain demographic and income information. The participant then records all household expenditure for a second one-week period. Note that because each household is surveyed only once, the data is a series of cross sections rather than a true panel. We select three years, 2000-2002, to represent the pre-CDCC expansion period and the following three years, 2003-2005, to represent the post-CDCC expansion period.

Only households with at least one child under age 13 are included in the analysis.[11] The tax interaction between the CDCC and the CTC was generally limited to taxpayers with $10,000 to $50,000 of family income, thus we only include families within this income range.[12] Married taxpayers with only one earner are not eligible for the CDCC and have much lower rates of using child care, so these

[11]This matches requirements to claim the CDCC, as the dependent qualifying child must be under age 13.

[12]We use the wage and salary income received by all household members in the past 12 months as the measure of family income. The consumer expenditure survey began imputing some missing income component values in 2004. To make the income measure comparable over the years of our study we remove imputed incomes which makes the income measure comparable across all years of the this study.

taxpayers are also excluded from the analysis. In this sample of low-income households with children, 26.2 percent were subject to the ACTC income constraint and thus were unable to claim their maximum CTC benefits through the ACTC. When also considering the CDCC, 46.0 percent of the sample were unable to claim the combined maximum value of the CDCC, CTC, and ACTC.

Summary statistics for our sample of households are given in Table 1. Our sample contains 2,682 households with young children, 268 of which paid for child care during the two-week survey. The child-care measure includes all expenditure for daycare, nursery, and preschool, including any tuition payments for preschool. The child-care measure does not include tuition payments for K-12 education, but would include other forms of formal child care. Babysitting is not included in the child-care measure as babysitting expenditure for non-work purposes cannot be used to claim the CDCC. A limitation of the CES two-week diary data is that some households that use child-care services pay for those services monthly, which will cause us to incorrectly categorize some households as not having any child-care expenditure. However, it should not do so in a way that is correlated with the CDCC expansion. Tests for differences in the means reported in Table 1 show that the pre and post periods are largely comparable, particularly for households with expenditure on child care. Inflation likely plays a role in the increase in spending over time as these figures are not inflation adjusted.

For each household, regardless of the year in which we observe them, we calculate a naive and nuanced measure of the value of the CDCC under both the pre-expansion (we use 2001) and the post-expansion (we use 2005) tax rules. Because we do not observe the chosen consumption level of each household under both tax rules, we cannot directly observe the change in total claimed benefits or marginal price. We can calculate two alternative measures, the change in maximum claimable benefits and the change in the cent per dollar discount on first-dollar marginal price. We refer to these measures respectively as the value of the CDCC and the child-care discount.

For maximum claimable benefits, the naive value of the CDCC is calculated as the statutory value of the credit if the taxpayer spent enough on childcare to reach the qualifying expenses limit for their household income.[13] This method does not consider any tax interactions and is how a taxpayer using spotlighting would approximate the value of the CDCC. The nuanced value of the CDCC is calculated as the difference in final tax liability by changing child-care spending from zero to the qualifying expense limit, holding all other factors constant. This method allows for interactions with other tax provisions.

For the discount on first-dollar marginal child-care price, the naive discount is calculated as the CDCC credit the taxpayer could claim if their spending changed from zero to one dollar of expenditure on child care.[14] In calculating the naive

[13]Spending levels of $3,000 for one young child and $6,000 for two or more young children are sufficient for claiming maximum benefits in both periods.

[14]Using the discount amount (1 marginal price) rather than the marginal price makes the interpretation of coefficients similar to the CDCC value approach.

discount, only the marginal credit rate shown in Panel (a) of Figure 1 is considered. Whether the taxpayer has any remaining tax liability is ignored. A taxpayer with less than $10,000 of income would have a naive discount of 30 cents for the first dollar of child-care expenditure before the tax change, and a naive discount of 35 cents after the tax change. This method does not consider any tax interactions and is how a taxpayer using spotlighting would approximate the first-dollar marginal cost of child care under the CDCC. The nuanced discount is calculated as the total change in tax liability or refund for the taxpayer if their child-care spending changed from zero to one dollar of expenditure. As shown in Panel (b) of Figure 1, the nuanced discount is zero if the ACTC income constraint binds, as any benefits from the CDCC would be offset by losses in the CTC.[15]

For both the pre- and post-expansion groups, the naive value of the CDCC is about $400 larger on average (a 50 percent increase) when calculated using the post-expansion tax rules as compared to the pre-expansion tax rules. The change from pre- to post-expansion tax rules in the nuanced value of the CDCC was significantly smaller than the naive value for both the pre- and post-expansion groups (t-values of 42.12 and 33.68, respectively). Similarly, Table 1 shows that the naive marginal price of child care decreased by a little less than five cents for the first dollar of child-care expenditure. Again, the nuanced change in the

[15]We do not adjust for the minimum value of the ACTC which may apply to families with three or more children. This means some households with binding minimums may be assigned non-zero changes in the first-dollar discounts when their true change is zero. Our results are robust to excluding all households with three or more children.

discount was significantly smaller than the naive change for both the pre- and post-expansion groups (t-values of 19.06 and 15.13, respectively).

Figure 3 shows the income distribution of households in our sample which saw an increase or a decrease in the nuanced value of the CDCC in Panel (a) and similarly in the nuanced child-care discount in Panel (b). As expected, households with an increase in the nuanced CDCC value had higher incomes than households which saw a decrease (t-value 36.54) and households with an increase in the nuanced child-care discount were also more likely to have higher incomes (t-value 18.17). Yet, there is extensive overlap in the income distributions in both Panels (a) and (b).

Figure 4 plots the changes in the naive and nuanced CDCC values by family income. Panel (a) shows that every household in our sample would have experienced an increase in the naive CDCC value between the pre- and post-expansion period with the largest increases concentrated among low-income households. The lower grouping of data points in Panel (a) is for households with one young child while those with more than one young child are in the higher grouping. Panel (b) shows the change in the nuanced value of the CDCC for each household in our sample and illustrates the heterogeneity of the change for households with similar levels of income. In our sample, 22.7 percent of households experienced a decrease in the nuanced value of the CDCC, 22.4 percent experienced no change, and 54 percent experienced an increase.

Figure 5 plots the naive and nuanced change in the first-dollar discount by family income. Panel (a) shows that the naive child-care discount increased for every family in our sample with income below $43,000. Panel (b) shows the change in the nuanced discount and illustrates that households with similar income can experience very different changes in the nuanced after-tax price of child care. In our sample, 7.5 percent of households experienced an increase in the first-dollar marginal price of child care, 38.6 percent of households experienced no change, and 54.0 percent of households experienced a decrease.

Differences in both the value of the CDCC and the price of child care are based only on change in the tax code and not on household differences over time. These figures describe a tax change that appeared to provide (if spotlighting) a large child-care subsidy to the low-income households in our sample. Yet for many low-income taxpayers, the nuanced value of the CDCC and nuanced price of child care remained unchanged or even moved in the opposite direction of the naive change.

2.4.2 Empirical Specification

By estimating the response of child-care spending to changes in the naive and nuanced value of the CDCC we are testing whether taxpayers are primarily ignorant of the CDCC change, are engaging in spotlighting, or are well-informed about the financial implications of the CDCC expansion. We estimate regression models of the following form where the $\Delta CDCC$ term is defined as either the change in the naive value, as indicated by the V superscript, or the nuanced value, as indicated by the U superscript:

$$\begin{aligned} E_{it} =& \beta_0 + \beta_1(\text{Post}_t \times \Delta CDCC_{it}^V) + \beta_2 \Delta CDCC_{it}^V + \beta_3(\text{Post}_t \times \Delta CDCC_{it}^U) \\ & + \beta_4 \Delta CDCC_{it}^U + \gamma \mathbf{X}_{it} + \theta_t + \varepsilon_{it} \end{aligned} \tag{2.6}$$

Households are indexed by i and time is indexed by t. The dependent variable is generally child-care expenditure or percentage of income spent on child care, though we use other spending measures in robustness checks.

The $\Delta CDCC$ variables are calculated for households in both the pre- and post-expansion periods holding all household characteristics constant. For those households that we observe in 2000-2002, this variable measures how the CDCC value would change if they faced the post-expansion tax rules. The variable *Post* is an indicator for the household being observed in 2003-2005. The coefficient on *Post* interacted with $\Delta CDCC$ is the difference-in-differences estimate of the causal effect of the change in the value of the CDCC on the measure of spending.

The identification comes from the assumption that households observed in 2003-2005 would have had the same spending on average as those observed in 2000-2002 had it not been for the tax change. To control for differences in the composition of the samples in the pre- and post-expansion periods we include a vector of observable characteristics, $\mathbf{X}$, including family income, race of the parent(s), educational attainment of the parent(s), and number of children. To account for inflation and trending we include a set of year fixed effects (given by θ). Reduced tax rates and the increased value of the CTC means taxpayers in the post period had lower tax liability on average than those in the pre-2003 period. Year fixed effects should also account for this income effect. We also include month fixed effects to control for seasonal variation such as differences in child-care spending during the summer versus the school year.

We also estimate specifications in which the change in the value of the Child and Dependent Care Credit, $\Delta CDCC$, is replaced with ΔD, the change in the discount on the first-dollar of child-care expenditure:

$$\begin{aligned} E_{it} =& \beta_0 + \beta_1(\text{Post}_t \times \Delta D^V_{it}) + \beta_2 \Delta D^V_{it} + \beta_3(\text{Post}_t \times \Delta D^U_{it}) \\ &+ \beta_4 \Delta D^U_{it} + \gamma \mathbf{X}_{it} + \theta_t + \varepsilon_{it} \end{aligned} \tag{2.7}$$

The naive change in the discount is indicated by the V superscript and the nuanced change is indicated by the U superscript. This alternative specification relies on the same identification assumptions, but allows us to estimate a response to a change in the after-tax price of child care rather than a change in the maximum credit value.

In both specifications, measurement error may impact both the naive and nuanced parameter estimates. The nuanced measure may contain more measurement error than the naive measure because the nuanced measure requires more information. In addition to attenuation bias from classical measurement error that may be present in both parameter estimates, correlation between the naive and nuanced measures could result in the coefficient estimate for the less noisy measure capturing some of the impact of the noisier measure. Such bias caused by the combination of measurement error and correlation between the naive and nuanced measures would make separating spotlighting behavior from fully-informed responses difficult. While naive and nuanced measures of the change in CDCC value have a sample correlation of -0.46, the measures of the change in child-care discount have a sample correlation of only -0.02.[16]

2.5 Results

2.5.1 Evidence of Spotlighting

We find evidence of a large and statistically significant effect of the change in the *naive* value of the CDCC on child-care expenditure and find little evidence of any effect from the change in the *nuanced* value of the CDCC. These results are reported in Table 2 and are consistent with spotlighting behavior. In the first three columns the dependent variable is the dollar value of child-care expenditure during the two-week survey period. In the last three columns the dependent variable is the percentage of income spent on child care. While the specifications in columns (1) and (4) include an indicator for family type, we also estimate the models separately for married and single households.

Because the CDCC expansion was passed in 2001 and was advertised in 2002, it is possible that the response began before the 2003 implementation. If

[16]For both approaches, similar results can be obtained when running separate regressions for naive and nuanced measures, suggesting that multicollinearity is not making the estimates unstable.

this is the case, our estimates of both the naive and the nuanced effect would be biased downward. It is also possible that the full effect of the CDCC expansion is realized with a lag as taxpayers realize that a change has taken place only when doing their taxes the next year. This would also cause a downward bias in our results. Therefore, Panel (b) of Table 2 reports results when the years 2002 and 2003 are excluded from the sample. This leaves us with a 2000-2001 pre-expansion period and a 2004-2005 post-expansion period from which to estimate the naive and nuanced effects. Estimates in Panel (b) of Table 2 are similar to those presented in Panel (a).

Estimates of the parameter of interest for the naive change in the value of the CDCC are large and often statistically significant for both the full sample and the sample excluding the years 2002 and 2003. Because the dependent variable in columns (1) through (3) is measured over a two-week period, an annual interpretation requires multiplying by 26. For example, the coefficient estimate of 0.039 implies that a one dollar increase in the naive value of the CDCC causes a \$1.01 (0.039 x 26) increase in annual child-care expenditure with a 95% confidence interval of (\$0.20, \$1.83). Multiplying by 26 may not be appropriate if households pay for child-care expenses monthly rather than every two-weeks. If all households are reporting monthly expenditures paid during that two week period, the coefficient estimate of 0.039 implies that a one dollar increase in the naive value of the CDCC causes a \$0.47 (0.039 x 12) increase in annual child-care expenditure with a 95%

confidence interval of (\$0.09, \$0.84).

One possible explanation for the large magnitude of the estimated response is that workers may choose from a limited number of options for hours of work.[17] When workers face such labor supply constraints, even a slight increase in naive child-care subsidy rates could persuade marginal families to make a large discrete change in both work hours and child-care expenditure. Average child-care expenditure may increase if lumpy adjustments exceed non-adjustments among households unable to make continuous consumption choices.

Estimates of the effect of the nuanced change in the value of the CDCC on child-care expenditure are not statistically different than zero. Importantly, in most specifications, we are able to reject the hypothesis that the naive and nuanced parameters are equal (p-value reported for each specification). We interpret the results as providing strong evidence of an effect of the change in the naive value of the CDCC on child-care expenditure and no evidence of an effect of the change in the nuanced value of the CDCC.

This result is illustrated in Figure 6 which plots the average child-care expenditure as a percentage of income by year for four groups of taxpayers (not mutually exclusive groups). In Panel (a), the dashed line plots average child-care expenditure for taxpayers with a larger than median change in the naive value of the CDCC and the solid line is for those taxpayers with a smaller than median

[17] See, for example, Altonji and Paxson (1988) or Dickens and Lundberg (1993). Golden (2001) notes female, non-white, and less educated workers (a group targeted by the CDCC expansion) are less likely to have flexible work schedules.

change in the naive value of the CDCC. Panel (b) is similar in that it groups taxpayers by the change in the nuanced value of the CDCC. The econometric model is not used in creating the figure as it simply reports the average child-care expenditure as a percentage of income for the different groups.

Figure 6 suggests that those with a large increase in the naive value of the CDCC increased their child-care expenditure, while those with a large increase in the nuanced value of the CDCC did not increase their spending on child care.[18] The increase in child-care expenditure for those with an above median naive CDCC change may seem to have begun even before the implementation of the CDCC expansion (indicated in the figure by the vertical line). This could simply be normal variation in the series or it could be a response in advance of the implementation given that the CDCC expansion was passed in 2001. The decline in child-care expenditure in 2005 for those with an above-median change in the naive value of the CDCC may indicate that spotlighting is a temporary behavior for some taxpayers. Consistent with the regression results from Table 2, there is no corresponding increase in child-care expenditure for those with an above-median change in the nuanced value of the CDCC. Importantly, there are no obvious differences in child-care expenditure for the different groups before 2003.

Table 3 reports the estimated effect of an increase in the first-dollar discount as specified in Equation (8). The magnitudes reported in Panel (a) of Table 3 are

[18] Indeed, households with changes in naive values above the 75th percentile spent 16.6% more in post-expansion period than households with changes in naive values below the 25th percentile (p-value = 0.057).

similar to the effect sizes reported in Panel (a) of Table 2, although the estimates are noisier. The coefficient estimate of 1.909 suggests that each additional cent per dollar decrease in the naive marginal price of child care causes expenditures to increase by $49.63 (1.909 x 26) per year, with a 95% confidence interval of (-$1.68, $100.95). Recall the average household with child-care expenditure spends $3,796 per year on child care ($146 x 26), so a cent per dollar decrease in the price of child care saves them almost $38. Hence Table 3 reports that a $1 "naive" increase in government expenditure causes a $1.31 (49.63/38) increase in annual child-care expenditure for this average household, with a 95% confidence interval of (-$0.04, $2.66). This is very similar to the $1.01 estimate reported in Table 2. If all households are reporting monthly expenditures paid during the two week period, this suggests an increase in child-care expenditure of $22.91 per year (1.909 x 12), with a 95% confidence interval of (-$0.78, $46.59). This would imply a $1 naive increase in government expenditure causes a $0.60 (22.91/38) increase in annual child-care expenditure, with a 95% confidence interval of (-$0.02, $1.23).

The responses to naive and nuanced measures of the marginal price of child-care are not statistically different for the full sample in Panel (a) of Table 3. Panel (b) excludes the year immediately before and the year immediately after the expansion. Estimates of the naive effect in Panel (b) are statistically significant and larger than responses to nuanced estimates, consistent with concerns about downward bias. In both panels, there remains no evidence that the nuanced change in the first-dollar discount has any effect on child-care expenditure.

Because available data does not pair expenditure on child care with a measure of the quantity or quality of child care, we do not know whether increased expenditures on the intensive margin reflects a larger quantity of child care or higher quality child care. It is possible to gain some insight about adjustment along the extensive margin by replacing the dependent variable with a dummy for non-zero child-care expenditure. Panel (a) of Table 4 shows there is little response by any group to the naive or nuanced maximum value of the CDCC. This is unsurprising, since the maximum value of the CDCC is unlikely to be the binding constraint for the first-dollar consumption of child care. The first-dollar discount more accurately reflects the binding constraint at the point of consumption for this group. Panel B of Table 4 shows the full sample of single parents in particular were 2.3% more likely to have non-zero child-care expenditure for each additional cent per dollar increase in the naive first-dollar discount on child care, with a 95% confidence interval of (0.5%, 4.1%). Similar estimates are obtained when excluding 2002 and 2003. For single parents, the response to the naive first-dollar discount is significantly larger than the response to the nuanced first-dollar discount. Again, we find no evidence that any group responded to the nuanced first-dollar discount.

2.5.2 Falsification Exercises

We perform two falsification exercises. The first is designed to see if the naive CDCC expansion had any effect on expenditure for other goods. The second is designed to see if we find similar results when we apply the same methods to a sample of households that were all in the pre-expansion period. We present results for both tests for only the change in value of the CDCC; using discounts to marginal price provides similar results.

If the change in the child-care subsidy affect expenditure on other unrelated goods it would raise concern about the causal interpretation. We examine expenditure on babysitting, nondurables, and all eleven generic aggregations as defined by the CES (food, alcoholic beverages, fuel, etc.) Table 5 reports the difference-in-differences estimates from estimating the same specification as reported in Table 2 columns (1) and (4) of Panel (a) where only the dependent variable is changed. There is little evidence that changes in the naive or the nuanced value of the CDCC affected expenditure. This suggests the causal effect on child-care expenditure is not simply capturing an income effect.[19]

Our second falsification exercise uses an additional sample of households from 1996 to 1999. In this exercise we assume that a hypothetical change in the value of the CDCC occurs at the beginning of 1999. The households observed in 1999-2001 are "treated" while those in 1996-1998 are the "control" group. Our measures of the naive and nuanced change in the value of the CDCC are still calculated by comparing the 2001 to the 2005 tax code, even though we are only using

[19]Several papers, including Johnson et al. (2006), Shapiro and Slemrod (2003), and Agarwal et al. (2007), have addressed how households respond to a sudden decrease in tax liability (like the sudden increase in the Child Tax Credit in 2003). They focus on what fraction of a tax rebate is spent rather than saved and find that households typically spend about 60 percent within the next year.

pre-expansion data. If a statistically significant response in child-care spending is found, such a false positive would raise concern about the causal interpretation of our main results. Table 6 reports no statistically significant response in child-care spending to this hypothetical treatment, which increases our confidence in the main results.

2.6 Conclusion

This paper examines how consumers respond to a change in a personal income tax provision when interactions with other elements of the tax code obfuscate the true impact of the changed provision. We use data from the Consumer Expenditure Survey to provide evidence that taxpayers engage in spotlighting behavior; they respond to the change in the particular tax provision in isolation without considering the interactions with other parts of the tax code. The evidence comes from our examination of the 2003 change to the Child and Dependent Care Credit (CDCC) which spotlighting taxpayers would have perceived as reducing the after-tax price of child care. However, interactions with other elements of the tax code, including the simultaneous change to the Child Tax Credit, reduced or even reversed this decrease in the after-tax price of child care for some taxpayers.

Using household data, we employ a difference-in-differences strategy which exploits the heterogeneity in the size of naive and nuanced measures of the change in value of the CDCC. We find strong evidence of a child-care expenditure response to the *naive* measure of the change in the value of the CDCC, which does not consider interactions with other elements of the tax code. We find little evidence of any response to the *nuanced* measure, which does account for interactions with other elements of the tax code. Similar results are found exploiting heterogeneity in the naive and nuanced marginal price of child care. Falsification exercises find little evidence the CDCC expansion (either naive or nuanced) affected expenditure on other goods. We also find no evidence of a response to a hypothetical CDCC expansion using pre-expansion data. We interpret these results as evidence that taxpayers were engaged in spotlighting behavior.

This paper supplements the existing tax salience literature by showing taxpayers can misperceive after-tax prices due to important but low-salience interactions, even when the direct financial implications are salient. Tax preparation software may reinforce spotlighting in some instances by focusing attention on each deduction or credit in isolation rather than on how different economic behavior affects final tax liability. This issue applies to any tax interactions that taxpayers may ignore, including other non-refundable tax credits, deductions and credits with phase-outs, and credits with income eligibility requirements.

Because available data does not pair expenditure on child care with a measure of the quantity or quality of child care, we do not always know whether increased expenditures reflect a larger quantity of child care or higher quality child care. There is some evidence that decreases in the naive measures of the marginal price of child induced single parents to begin consuming child care, and this may have increased female labor force participation rates.

Despite any influence the naive measure of the CDCC had on taxpayers child-care expenditures and labor decisions, the government did not bear the cost associated with the naive value of the CDCC. Instead, the government bore the cost of the nuanced value which in our sample of low-income families with children was just 47 percent of what the government would have born had they paid the full cost associated with the naive value. Our results indicate that taxpayers significantly increased their expenditure on child care in response to the 2003 expansion of the CDCC regardless of whether their after-tax price of child care decreased.

2.7 Acknowledgments

Chapter 2, in full, has been submitted for publication of the material as it may appear in the National Tax Journal, 2015, Miller, Benjamin M.; Mumford, Kevin. The dissertation/thesis author was a primary investigator and author of this paper.

2.8 Figures and Tables

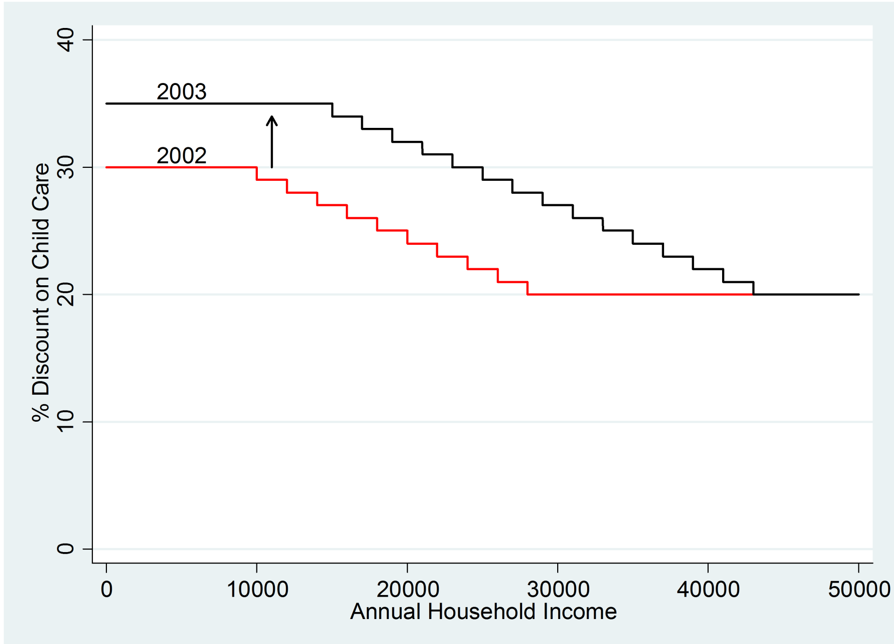

(a) Naive Credit Value

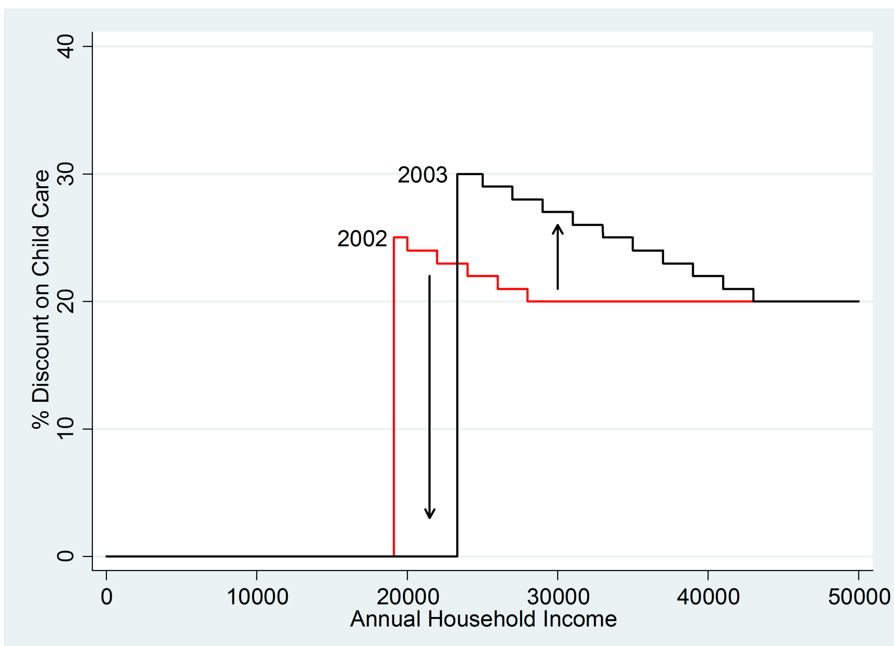

(b) Nuanced Credit Value

Figure 2.1: Child and Dependent Care Credit Rate Increase

Notes: These figure illustrates the percent of the first dollar of child-care expenditure that is refunded through the CDCC to a single-parent household with two children. Panel (a) presents the naive value of the CDCC which does not consider interactions with other elements of the tax code. Panel (b) presents the nuanced value of the CDCC where interactions with all other elements of the tax code are considered. Because the CDCC is a non-refundable tax credit, many low-income taxpayers do not benefit from this credit.

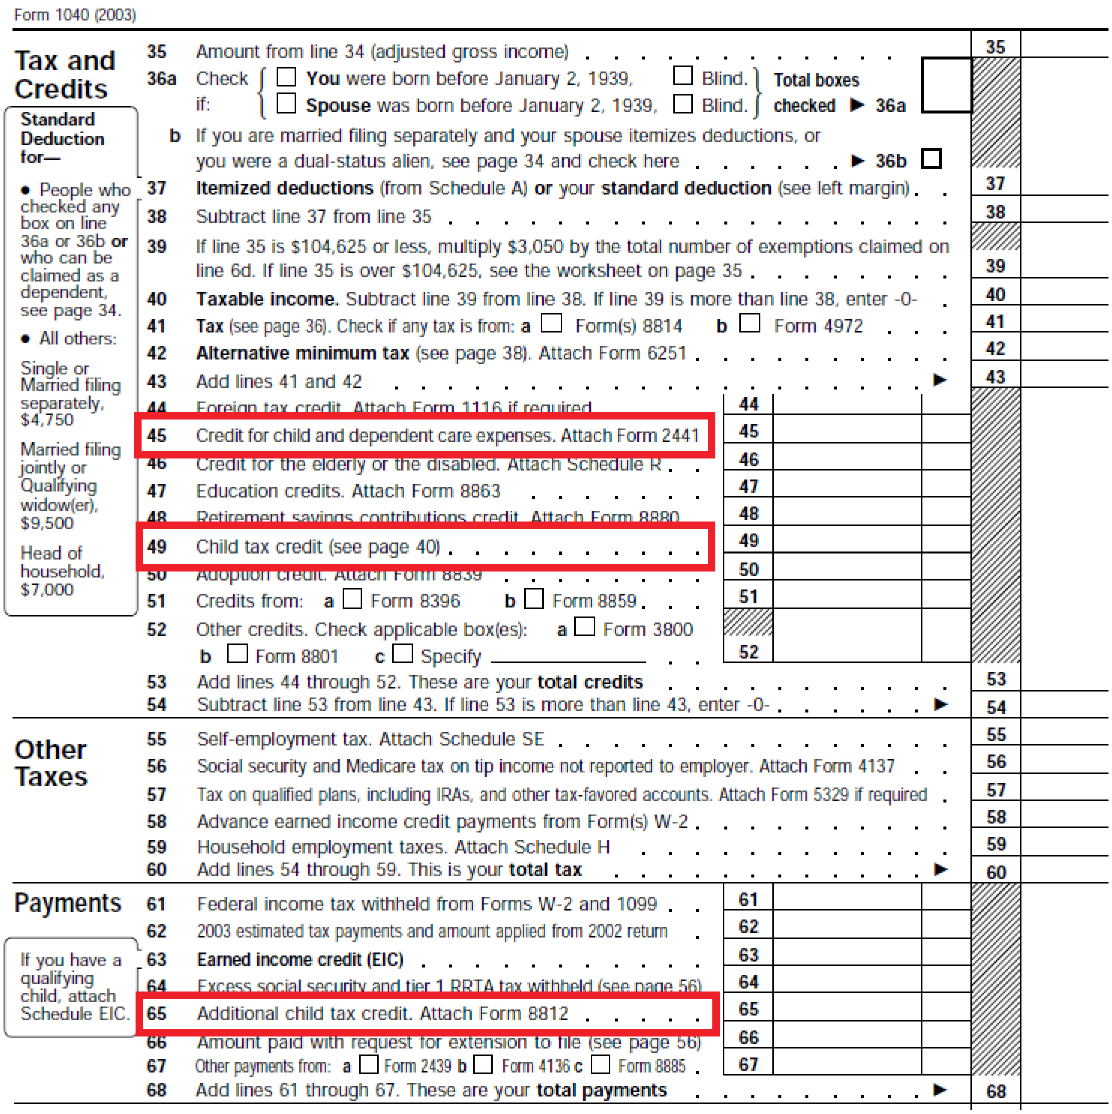

Form 1040 (2003)

Tax and Credits

Standard Deduction for—
- People who checked any box on line 36a or 36b **or** who can be claimed as a dependent, see page 34.
- All others:

Single or Married filing separately, $4,750

Married filing jointly or Qualifying widow(er), $9,500

Head of household, $7,000

35	Amount from line 34 (adjusted gross income)	35
36a	Check if: ☐ **You** were born before January 2, 1939, ☐ Blind. ☐ **Spouse** was born before January 2, 1939, ☐ Blind. **Total boxes checked** ▶ **36a**	
b	If you are married filing separately and your spouse itemizes deductions, or you were a dual-status alien, see page 34 and check here ▶ **36b** ☐	
37	**Itemized deductions** (from Schedule A) **or** your **standard deduction** (see left margin)	37
38	Subtract line 37 from line 35	38
39	If line 35 is $104,625 or less, multiply $3,050 by the total number of exemptions claimed on line 6d. If line 35 is over $104,625, see the worksheet on page 35	39
40	**Taxable income.** Subtract line 39 from line 38. If line 39 is more than line 38, enter -0-	40
41	**Tax** (see page 36). Check if any tax is from: **a** ☐ Form(s) 8814 **b** ☐ Form 4972	41
42	**Alternative minimum tax** (see page 38). Attach Form 6251	42
43	Add lines 41 and 42 ▶	43
44	Foreign tax credit. Attach Form 1116 if required	44
45	Credit for child and dependent care expenses. Attach Form 2441	45
46	Credit for the elderly or the disabled. Attach Schedule R	46
47	Education credits. Attach Form 8863	47
48	Retirement savings contributions credit. Attach Form 8880	48
49	Child tax credit (see page 40)	49
50	Adoption credit. Attach Form 8839	50
51	Credits from: **a** ☐ Form 8396 **b** ☐ Form 8859	51
52	Other credits. Check applicable box(es): **a** ☐ Form 3800 **b** ☐ Form 8801 **c** ☐ Specify	52
53	Add lines 44 through 52. These are your **total credits**	53
54	Subtract line 53 from line 43. If line 53 is more than line 43, enter -0- ▶	54

Other Taxes

55	Self-employment tax. Attach Schedule SE	55
56	Social security and Medicare tax on tip income not reported to employer. Attach Form 4137	56
57	Tax on qualified plans, including IRAs, and other tax-favored accounts. Attach Form 5329 if required	57
58	Advance earned income credit payments from Form(s) W-2	58
59	Household employment taxes. Attach Schedule H	59
60	Add lines 54 through 59. This is your **total tax** ▶	60

Payments

If you have a qualifying child, attach Schedule EIC.

61	Federal income tax withheld from Forms W-2 and 1099	61
62	2003 estimated tax payments and amount applied from 2002 return	62
63	**Earned income credit (EIC)**	63
64	Excess social security and tier 1 RRTA tax withheld (see page 56)	64
65	Additional child tax credit. Attach Form 8812	65
66	Amount paid with request for extension to file (see page 56)	66
67	Other payments from: **a** ☐ Form 2439 **b** ☐ Form 4136 **c** ☐ Form 8885	67
68	Add lines 61 through 67. These are your **total payments** ▶	68

Figure 2.2: IRS 1040 Form (2003)

Notes: This is the top portion of page 2 of the 1040 form for year 2003. Line 45 is the Child and Dependent Care Credit (CDCC), line 49 is the Child Tax Credit (CTC), and Line 65 is the Additional Child Tax Credit (ACTC). The stacking order of the credits on the 1040 form has remained the same since 2003.

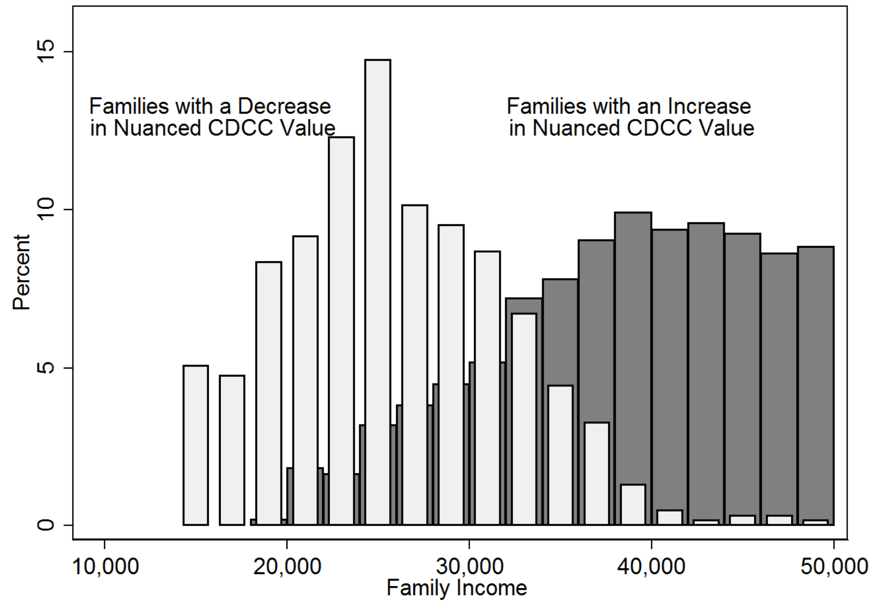

(a) Income Distribution by Change in Nuanced CDCC Value

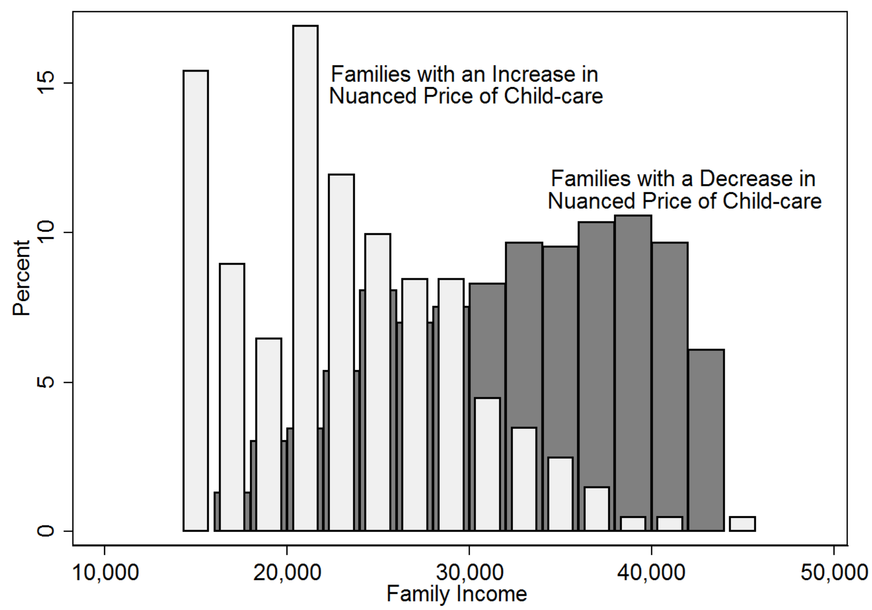

(b) Income Distribution by Change in Nuanced Price of Child-care

Figure 2.3: Income Distribution of Households by Group

Notes: Includes all CES households from 2000 to 2005 with at least one child under age 13 and income between \$10,000 and \$50,000. The income distribution for the two groups (those with a decrease in the nuanced price of child-care and those with an increase in the nuanced price of childcare) were graphed separately and then combined into this figure. Households with no change in the nuanced price of child-care do not appear on this figure.

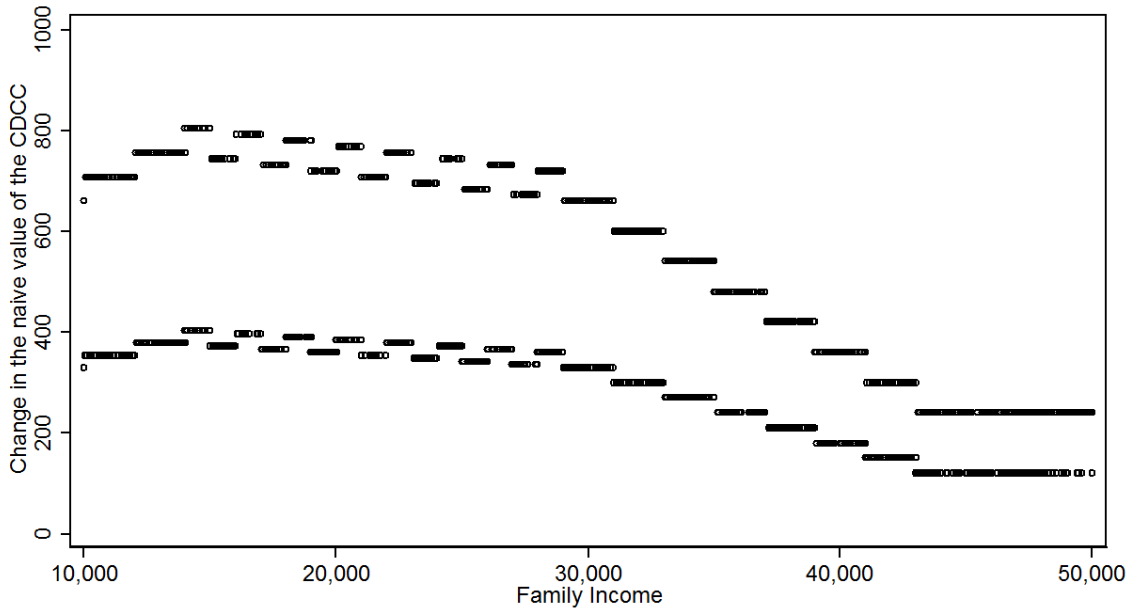

(a) Naive CDCC Change

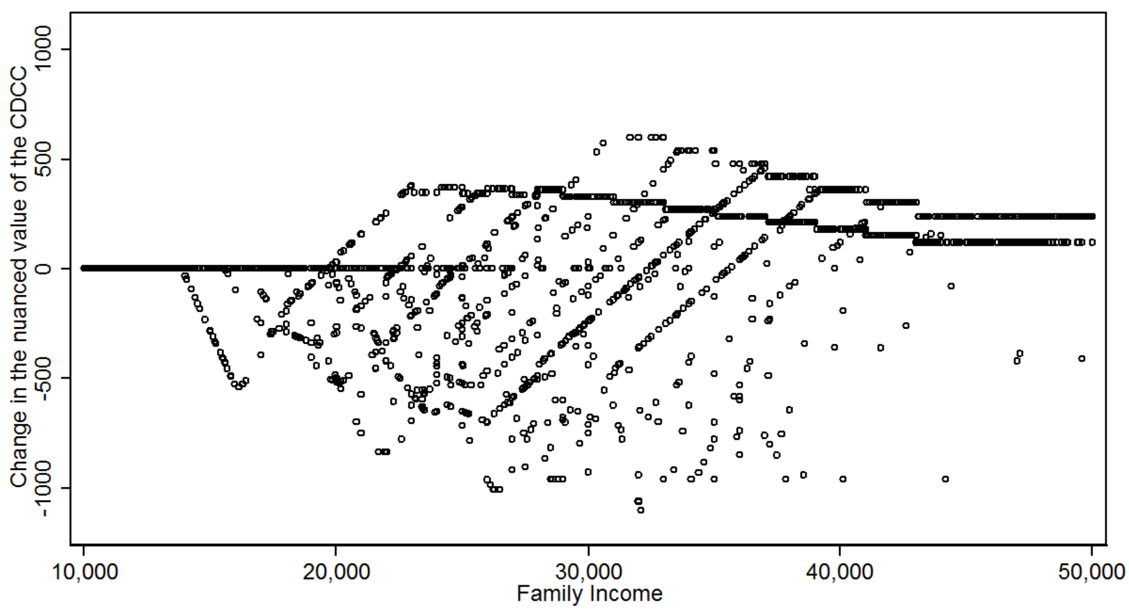

(b) Nuanced CDCC Change

Figure 2.4: Change in the Naive and Nuanced Value of the CDCC and Income

Notes: Panel (a) and Panel (b) depict each family in the data as a circle with family income on the x-axis. The y-axis in Panel (a) is the change in the maximum value of the child and dependent care credit between 2000 and 2005 if it were a fully refundable credit. The y-axis in Panel (b) is the change in the nuanced value of the child and dependent care credit (a non-refundable credit) between 2000 and 2005.

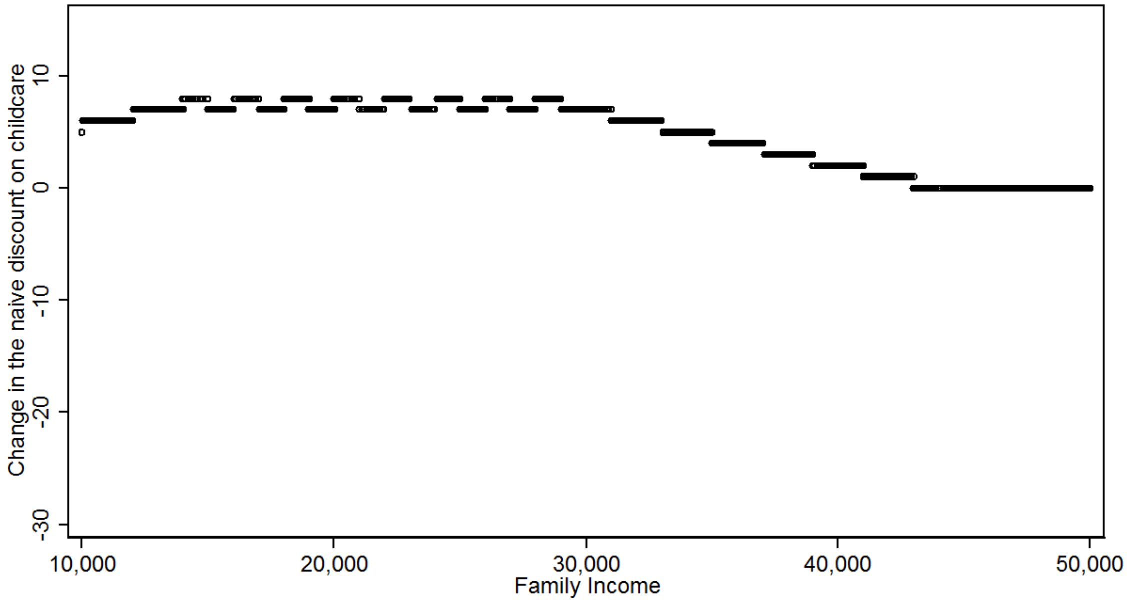

(a) Naive Change in Discount on Child-Care

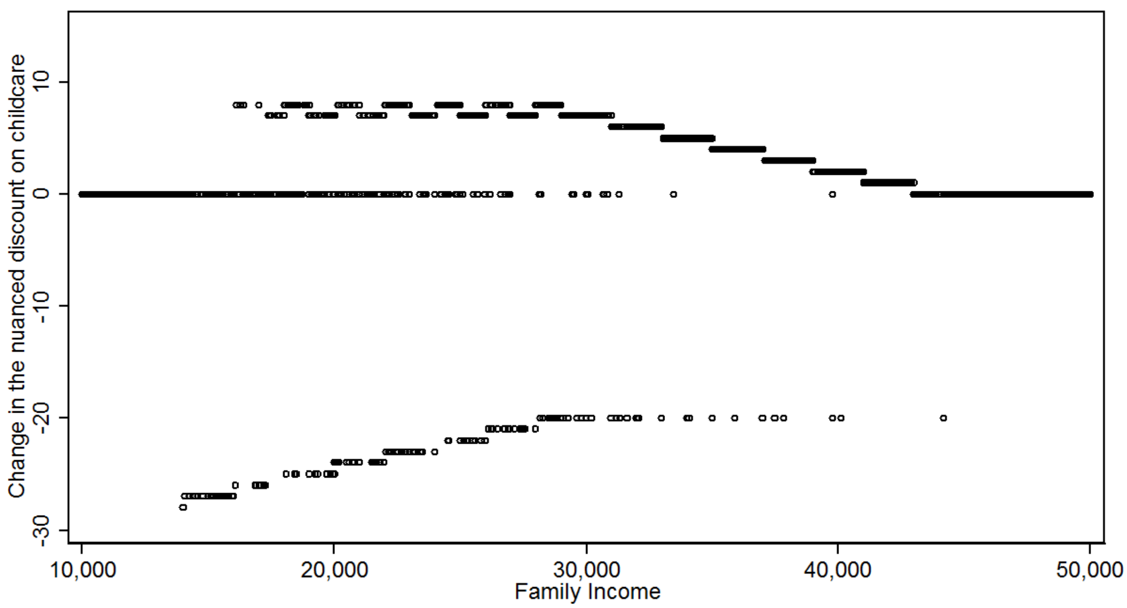

(b) Nuanced Change in Discount on Child-Care

Figure 2.5: Change in the Naive and Nuanced Value of the CDCC and Income

Notes: Panel (a) and Panel (b) depict each family in the data as a circle with family income on the x-axis. The y-axis in Panel (a) is the change between 2000 and 2005 in the value of the Child and Dependent Care Credit in cents if the credit were fully refundable assuming one dollar of child-care expenditure. The y-axis in Panel (b) is the change between 2000 and 2005 in the nuanced value of the CDCC in cents assuming one dollar of child-care expenditure.

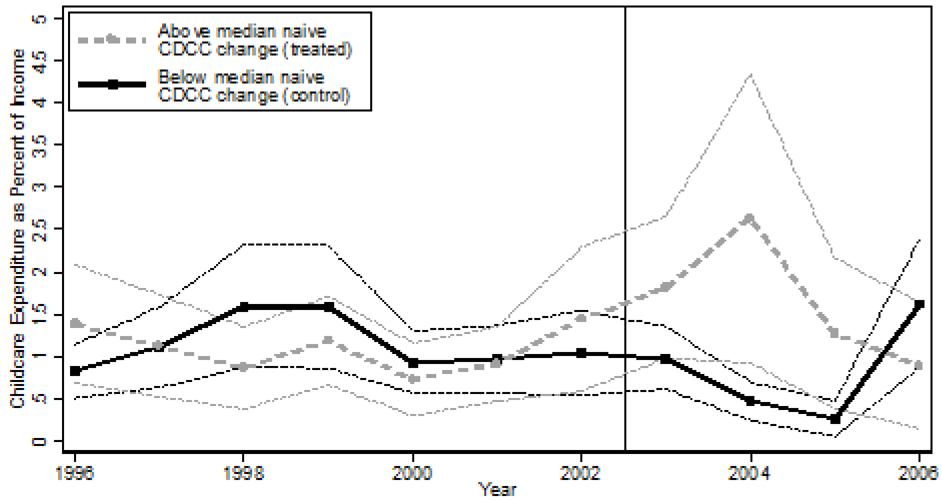

(a) By Naive CDCC Change

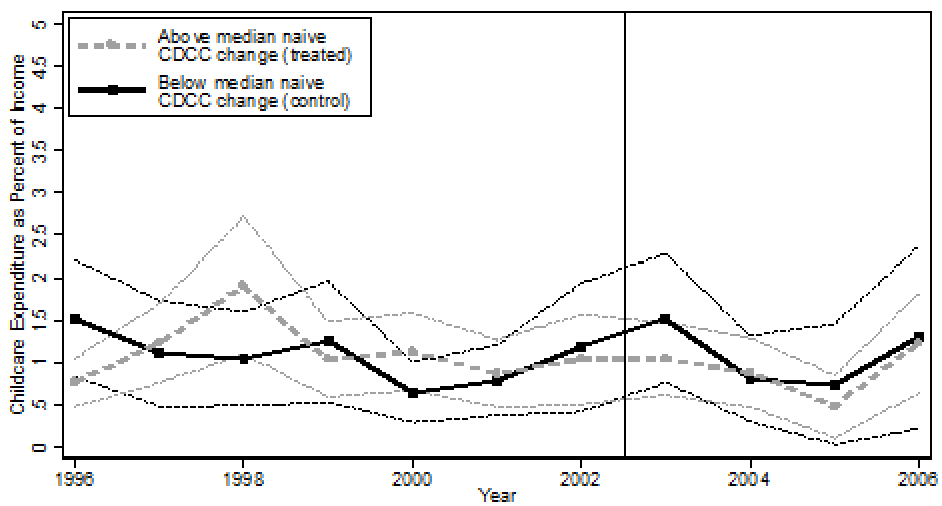

(b) By Nuanced CDCC Change

Figure 2.6: Average Child Care Expenditure by Year

Notes: Panel (a) and Panel (b) plot the average childcare expenditure as a percentage of income by year for CES households with income between $10,000 & $50,000 for two groups. The treated group is defined as those individuals with an above median change in naive or nuanced value of the CDCC. The control group is defined as those individuals with a below median change in naive or nuanced value of the CDCC. 90 percent confidence intervals are indicated in gray for the treated group and black for the control group. Households with a very large or small (top or bottom 5%) change (defined separately for each panel) are excluded.

Table 2.1: Summary Statistics

Full Sample					
	2000-2002		2003-2005		H_0:
Variables	Mean	Std. Dev.	Mean	Std. Dev.	Pre = Post
Expenditure on child care	15.89	92.47	13.50	53.55	0.4046
Expenditure on babysitting	10.59	46.58	8.90	41.95	0.3262
Expenditure on nondurables	445.50	279.73	473.42	304.55	0.0067
Expenditure on all categories	1,139.14	1,014.49	1,256.50	1,453.99	0.0231
Naive CDCC value pre	791.00	287.46	779.35	288.76	0.3314
Naive CDCC value post	1,206.92	470.81	1,185.16	473.03	0.2674
Naive CDCC value change	415.92	206.80	405.81	208.84	0.2414
Nuanced CDCC value pre	512.88	355.94	515.21	351.21	0.8732
Nuanced CDCC value post	551.00	463.70	575.41	456.18	0.2005
Nuanced CDCC value change	38.12	289.50	60.20	270.51	0.0525
Naive discount pre	22.22	3.04	22.12	3.00	0.4344
Naive discount post	27.12	5.22	26.91	5.22	0.3365
Naive discount change	4.90	2.91	4.79	2.96	0.3660
Nuanced discount pre	16.28	8.88	16.50	8.65	0.5434
Nuanced discount post	17.18	11.82	17.63	11.50	0.3534
Nuanced discount change	0.90	7.80	1.13	7.49	0.4721
Income	30,131.68	11,352.05	30,514.57	11,322.12	0.4207
Married (indicator variable)	0.59	0.49	0.50	0.50	0.0000
Number of Children	1.70	0.92	1.65	0.83	0.1428
Number of Observations	1573		1109		

Notes: The data comes from the Consumer Expenditure Survey and only includes households with at least one child under age 13 and self-reported family income between $10,000 and $50,000. Married couples with only one working spouse are excluded from the data. Expenditure values are from a two-week diary from years 2000-2005. Spending on nondurable goods is defined as in Johnson, Parker, and Souleles (2006) as spending on goods and services which can only be used once and last no more than 3 years at most. The final column reports the p-values from the null hypothesis that the mean is the same in both the pre and post periods. Sample weights used in calculations.

Table 2.2: Summary Statistics

Households with Expenditure on Child Care

	2000-2002		2003-2005		H_0:
Variables	Mean	Std. Dev.	Mean	Std. Dev.	Pre = Post
Expenditure on child care	146.38	244.89	145.79	108.53	.9788
Expenditure on babysitting	13.07	38.47	12.7	39.42	0.9413
Expenditure on nondurables	494.73	273.78	526.54	336.40	0.4236
Expenditure on all categories	1,485.19	1,086.08	1589.13	1,154.43	0.4850
Naive CDCC value pre	782.97	281.00	782.48	294.12	0.9900
Naive CDCC value post	1,178.13	470.15	1,186.78	475.59	0.8927
Naive CDCC value change	395.16	214.98	404.30	211.13	.7504
Nuanced CDCC value pre	586.50	335.68	536.77	352.85	.2924
Nuanced CDCC value post	663.34	453.58	625.14	447.48	0.5312
Nuanced CDCC value change	76.84	275.80	88.37	257.86	0.7313
Naive discount pre	21.67	2.70	22.01	2.83	0.3787
Naive discount post	26.00	5.18	26.83	5.28	0.2442
Naive discount change	4.33	3.13	4.82	3.16	0.2431
Nuanced discount pre	17.74	7.64	16.78	8.45	0.4018
Nuanced discount post	19.37	10.50	16.78	11.50	0.8795
Nuanced discount change	1.64	6.97	2.36	5.96	0.3944
Income	32,708.51	11,303.98	30,805.96	11,400.94	0.2166
Married (indicator variable)	0.55	0.50	0.53	0.50	0.8374
Number of Children	1.70	0.87	1.63	0.78	0.5259
Number of Observations	167		101		

Notes: The data comes from the Consumer Expenditure Survey and only includes households with at least one child under age 13 and self-reported family income between \$10,000 and \$50,000. Married couples with only one working spouse are excluded from the data. Expenditure values are from a two-week diary from years 2000-2005. Spending on nondurable goods is defined as in Johnson, Parker, and Souleles (2006) as spending on goods and services which can only be used once and last no more than 3 years at most. The final column reports the p-values from the null hypothesis that the mean is the same in both the pre and post periods. Sample weights used in calculations.

Table 2.3: Effect of CDCC Value on Child-Care Expenditure

	Full Sample					
	Dollars of Two-Week Expenditure			Percent of Annual Income		
	all	marries	single	all	marries	single
	(1)	(2)	(3)	(4)	(5)	(6)
Naive diff-in-diff	0.039**	0.030	0.062***	0.005**	0.003	0.008**
	(0.016)	(0.027)	(0.024)	(0.002)	(0.003)	(0.004)
Naive ΔCDCC	-0.023	0.019	-0.064	-0.003	0.003	-0.008
	(0.023)	(0.023)	(0.039)	(0.004)	(0.002)	(0.006)
Nuanced diff-in-diff	0.010	0.002	0.019	0.002	0.000	0.003
	(0.009)	(0.012)	(0.016)	(0.001)	(0.001)	(0.002)
Nuanced ΔCDCC	-0.004	0.003	-0.009	-0.001	0.000	-0.002
	(0.007)	(0.005)	(0.014)	(0.001)	0.001)	(0.002)
Income ($1,000s)	0.020	0.483*	-0.204	-0.063	0.025	-0.116
	(0.304)	(0.256)	(0.482)	(0.051)	(0.027)	(0.083)
H_0: Naive = Nuanced	0.021	0.158	0.016	0.041	0.258	0.027
Observations	2,682	1,483	1,199	2,682	1,483	1,199
R-squared	0.026	0.047	0.030	0.019	0.048	0.026

Notes: All specifications include month and year fixed effects as well as indicators for the race of the parent(s), education of the parents(s), family type, number of young children, and the presence of a child age 13 or more (potential babysitting by sibling). The data only includes households with young (under age 13) children and an annual income between $10,000 and $50,000. Both parents must earn income in two-parent households to be included in the data. Child-care expenditure as a percentage of annual income is calculated as 26 times the reported two-week child-care expenditure divided by annual income and multiplied by 100. P-values from Wald tests for the equality of Naive and Nuanced diff-in-diff estimates are reported. Robust standard errors in parentheses: $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$

Table 2.4: Effect of CDCC Value on Child-Care Expenditure

Excluding Years 2002 and 2003

	Dollars of Two-Week Expenditure			Percent of Annual Income		
	all	marries	single	all	marries	single
	(1)	(2)	(3)	(4)	(5)	(6)
Naive diff-in-diff	0.050**	0.076	0.042**	0.006*	0.008	0.004
	(0.023)	(0.047)	(0.021)	(0.003)	(0.006)	(0.003)
Naive ΔCDCC	-0.005	0.024	-0.050	0.000	0.004	-0.004
	(0.020)	(0.032)	(0.040)	(0.002)	(0.003)	(0.003)
Nuanced diff-in-diff	0.005	0.020	-0.005	0.001	0.003	0.000
	(0.010)	(0.020)	(0.014)	(0.001)	(0.002)	(0.001)
Nuanced ΔCDCC	0.005	0.005	0.009	0.000	0.000	0.001
	(0.006)	(0.007)	(0.010)	(0.000)	(0.001)	(0.001)
Income ($1,000s)	0.310	0.522	0.211	-0.007	0.025	-0.027
	(0.190)	(0.383)	(0.301)	(0.023)	(0.043)	(0.038)
H_0: Naive = Nuanced	0.012	0.078	0.016	0.043	0.152	0.050
Observations	1,700	949	751	1,700	949	751
R-squared	0.046	0.065	0.069	0.045	0.077	0.053

Notes: All specifications include month and year fixed effects as well as indicators for the race of the parent(s), education of the parents(s), family type, number of young children, and the presence of a child age 13 or more (potential babysitting by sibling). The data only includes households with young (under age 13) children and an annual income between $10,000 and $50,000. Both parents must earn income in two-parent households to be included in the data. Child-care expenditure as a percentage of annual income is calculated as 26 times the reported two-week child-care expenditure divided by annual income and multiplied by 100. P-values from Wald tests for the equality of Naive and Nuanced diff-in-diff estimates are reported. Robust standard errors in parentheses: $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$

Table 2.5: Effect of Child-Care Discount on Child-Care Expenditure

	Full Sample					
	Dollars of Two-Week Expenditure			Percent of Annual Income		
	all	marries	single	all	marries	single
	(1)	(2)	(3)	(4)	(5)	(6)
Naive diff-in-diff	1.909*	2.011	3.786*	0.127	0.187	0.177
	(1.007)	(1.237)	(1.982)	(0.123)	(0.130)	(0.229)
Naive Δ Discount	-0.435	-0.237	-1.655	-0.004	0.017	-0.013
	(0.945)	(0.797)	(2.357)	(0.141)	(0.069)	(0.356)
Nuanced diff-in-diff	0.879	-0.142	1.780	0.145	-0.023	0.311
	(0.830)	(0.308)	(1.711)	(0.147)	(0.036)	(0.305)
Nuanced Δ Discount	-0.752	0.055	-1.597	-0.136	0.011	-0.292
	(0.870)	(0.140)	(1.795)	(0.156)	(0.014)	(0.321)
Income ($1,000s)	0.226	0.254	0.274	-0.033	-0.009	-0.041
	(0.161)	(0.208)	(0.238)	(0.026)	(0.024)	(0.042)
H_0: Naive = Nuanced	0.512	0.105	0.528	0.940	0.152	0.791
Observations	2,682	1,483	1,199	2,682	1,483	1,199
R-squared	0.028	0.047	0.038	0.023	0.045	0.036

Notes: All specifications include month and year fixed effects as well as indicators for the race of the parent(s), education of the parents(s), family type, number of young children, and the presence of a child age 13 or more (potential babysitting by sibling). The data only includes households with young (under age 13) children and an annual income between $10,000 and $50,000. Both parents must earn income in two-parent households to be included in the data. Child-care expenditure as a percentage of annual income is calculated as 26 times the reported two-week child-care expenditure divided by annual income and multiplied by 100. P-values from Wald tests for the difference between Naive and Nuanced diff-in-diff estimates are reported. Robust standard errors in parentheses: $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$

Table 2.6: Effect of Child-Care Discount on Child-Care Expenditure

Excluding Years 2002 and 2003

	Dollars of Two-Week Expenditure			Percent of Annual Income		
	all	marries	single	all	marries	single
	(1)	(2)	(3)	(4)	(5)	(6)
Naive diff-in-diff	3.661***	4.249**	5.890***	0.299***	0.383*	0.418***
	(1.118)	(1.890)	(2.049)	(0.113)	(0.222)	(0.152)
Naive Δ Discount	-0.804	0.186	-3.651*	-0.080	0.050	-0.335*
	(0.914)	(1.002)	(1.964)	(0.108)	(0.091)	(0.193)
Nuanced diff-in-diff	0.052	-0.820	0.217	-0.001	-0.099	0.021
	(0.236)	(0.609)	(0.286)	(0.025)	(0.071)	(0.029)
Nuanced Δ Discount	0.133	0.214	0.129	0.024**	0.030	0.026
	(0.123)	(0.190)	(0.182)	(0.012)	(0.019)	(0.020)
Income ($1,000s)	0.348*	0.380	0.414	-0.023	-0.002	-0.029
	(0.208)	(0.271)	(0.287)	(0.031)	(0.031)	(0.046)
H_0: Naive = Nuanced	0.001	0.024	0.005	0.010	0.076	0.007
Observations	1,700	949	751	1,700	949	751
R-squared	0.047	0.063	0.080	0.041	0.070	0.056

Notes: All specifications include month and year fixed effects as well as indicators for the race of the parent(s), education of the parents(s), family type, number of young children, and the presence of a child age 13 or more (potential babysitting by sibling). The data only includes households with young (under age 13) children and an annual income between $10,000 and $50,000. Both parents must earn income in two-parent households to be included in the data. Child-care expenditure as a percentage of annual income is calculated as 26 times the reported two-week child-care expenditure divided by annual income and multiplied by 100. P-values from Wald tests for the difference between Naive and Nuanced diff-in-diff estimates are reported. Robust standard errors in parentheses: $^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$

Table 2.7: Effect on Extenisve Margin

Maximum Value of CDCC

	Full Sample			Excluding Years 2002 and 2003		
	all	marries	single	all	marries	single
	(1)	(2)	(3)	(4)	(5)	(6)
Naive diff-in-diff	0.006	0.002	0.011	0.012	0.013	0.015
	(0.007)	(0.009)	(0.011)	(0.009)	(0.013)	(0.012)
Naive Δ Discount	0.007	0.010	-0.001	-0.001	0.005	-0.015
	(0.009)	(0.012)	(0.016)	(0.011)	(0.016)	(0.019)
Nuanced diff-in-diff	0.000	-0.002	0.001	0.002	0.005	-0.004
	(0.005)	(0.006)	(0.007)	(0.006)	(0.008)	(0.008)
Nuanced Δ Discount	-0.003	0.006	0.002	-0.001	0.008	
	(0.003)	(0.004)	(0.005)	(0.004)	(0.005)	(0.005)
Income ($1,000s)	0.207*	0.284*	0.119	0.124	0.158	0.063
	(0.109)	(0.163)	(0.157)	(0.137)	(0.226)	(0.201)
H_0: Naive = Nuanced	0.699	0.514	0.114	0.380	0.716	0.065
Observations	2,682	1,483	1,199	1,700	949	751
R-squared	0.045	0.051	0.061	0.050	0.055	0.072

Notes: All specifications include month and year fixed effects as well as indicators for the race of the parent(s), education of the parents(s), family type, number of young children, and the presence of a child age 13 or more (potential babysitting by sibling). The data only includes households with young (under age 13) children and an annual income between $10,000 and $50,000. Both parents must earn income in two-parent households to be included in the data. Dependent variable is an indicator for non-zero childcare expenditure, and is multiplied by 100. P-values from Wald tests for the difference between Naive and Nuanced diff-in-diff estimates are reported. Robust standard errors in parentheses: $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$

Table 2.8: Effect on Extensive Margin

First-Dollar Marginal Price of Childcare

	Full Sample			Excluding Years 2002 and 2003		
	all	marries	single	all	marries	single
	(1)	(2)	(3)	(4)	(5)	(6)
Naive diff-in-diff	0.679	0.472	1.880**	1.091**	0.964	2.381**
	(0.452)	(0.590)	(0.834)	(0.545)	(0.771)	(0.988)
Naive Δ Discount	-0.156	0.423	-1.361*	-0.206	0.692	-1.828*
	(0.455)	(0.552)	(0.814)	(0.560)	(0.693)	(0.966)
Nuanced diff-in-diff	0.083	-0.016	0.059	0.005	-0.450	0.130
	(0.145)	(0.285)	(0.159)	(0.177)	(0.453)	(0.176)
Nuanced Δ Discount	0.040	-0.115	0.196*	0.138	0.111	0.176
	(0.098)	(0.158)	(0.119)	(0.088)	(0.144)	(0.112)
Income ($1,000s)	0.127	0.259*	0.029	0.145	0.294	0.062
	(0.101)	(0.135)	(0.140)	(0.132)	(0.184)	(0.185)
H_0: Naive = Nuanced	0.398	0.996	0.014	0.093	0.373	0.019
Observations	2,682	1,483	1,199	1,700	949	751
R-squared	0.046	0.051	0.066	0.053	0.059	0.079

Notes: All specifications include month and year fixed effects as well as indicators for the race of the parent(s), education of the parents(s), family type, number of young children, and the presence of a child age 13 or more (potential babysitting by sibling). The data only includes households with young (under age 13) children and an annual income between $10,000 and $50,000. Both parents must earn income in two-parent households to be included in the data. Dependent variable is an indicator for non-zero childcare expenditure, and is multiplied by 100. P-values from Wald tests for the difference between Naive and Nuanced diff-in-diff estimates are reported. Robust standard errors in parentheses: $^{*}p < 0.1,^{**}p < 0.05,^{***}p < 0.01$

Table 2.9: Effect on Other Expenditure

Expenditure Category	Dollars of Two-Week Expenditure			Percent of Annual Income		
	Naive diff-in-diff	Nuanced diff-in-diff	H_0: Naive= Nuanced	Naive diff-in-diff	Nuanced diff-in-diff	H_0: Naive= Nuanced
Babysitting	-0.003	-0.005	0.793	-0.000	-0.000	0.851
	(0.011)	(0.005)		(0.001)	(0.001)	
Nondurables	-0.053	-0.046	0.902	0.000	-0.005	0.462
	(0.065)	(0.048)		(0.008)	(0.005)	
Food	-0.044	-0.031	0.663	-0.005	-0.004*	0.753
	(0.031)	(0.023)		(0.004)	(0.002)	
Alcohol	-0.005	0.001	0.242	-0.000	0.000	0.553
	(0.005)	(0.003)		(0.001)	(0.000)	
Fuel	-0.007	-0.008	0.988	0.002	-0.001	0.501
	(0.036)	(0.025)		(0.004)	(0.002)	
Household Supplies	-0.010	-0.002	0.139	-0.001	-0.000	0.253
	(0.006)	(0.005)		(0.001)	(0.001)	
Household Furnishings	-0.007	0.035	0.331	-0.001	0.003	0.208
	(0.041)	(0.037)		(0.004)	(0.003)	
Apparel	-0.032	-0.023	0.742	-0.005	-0.002	0.479
	(0.030)	(0.023)		(0.004)	(0.003)	
Gasoline and Motor Oil	0.002	-0.003	0.635	0.002*	-0.000	0.024
	(0.011)	(0.008)		(0.001)	(0.001)	
Medical Supplies	0.005	0.004	0.784	0.001	0.000	0.910
	(0.006)	(0.003)		(0.001)	(0.000)	
Entertainment	0.037	0.000	0.170	0.006	0.001	0.121
	(0.030)	(0.018)		(0.004)	(0.002)	
Personal Care	-0.005	-0.001	0.568	-0.001	-0.000	0.649
	(0.006)	(0.004)		(0.001)	(0.000)	
Miscellaneous	-0.034	0.040	0.051	-0.002	0.004	0.058
	(0.022)	(0.029)		(0.002)	(0.003)	

Notes: All specifications are identical to the main results and include month and year fixed effects as well as indicators for the race of the parent(s), education of the parents(s), family type, number of young children, and the presence of a child age 13 or more (potential babysitting by sibling). The data only includes households with young (under age 13) children and an annual income between $10,000 and 50$,000 over the full sample of years 2000-2005. Both parents must earn income in two-parent households to be included in the data. Child-care expenditure as a percentage of annual income is calculated as 26 times the reported two-week child-care expenditure divided by annual income and multiplied by 100. P-values from Wald tests for the difference between Naive and Nuanced diff-in-diff estimates are reported. Robust standard errors in parentheses: $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$

Table 2.10: Falsification Exercise (1996-2001 data)

	Dollars of Two-Week Expenditure			Percent of Annual Income		
	all	marries	single	all	marries	single
	(1)	(2)	(3)	(4)	(5)	(6)
Naive diff-in-diff	-0.006	-0.002	0.011	0.000	0.001	0.002
	(0.010)	(0.012)	(0.020)	(0.001)	(0.001)	(0.002)
Naive Δ Discount	-0.006	-0.010	-0.027	-0.002	-0.001	-0.005
	(0.017)	(0.018)	(0.036)	(0.002)	(0.002)	(0.004)
Nuanced diff-in-diff	-0.003	0.002	-0.010	0.000	0.000	-0.000
	(0.007)	(0.008)	(0.013)	(0.001)	(0.001)	(0.001)
Nuanced Δ Discount	0.009	0.005	0.014	0.000	0.000	0.001
	(0.006)	(0.006)	(0.011)	(0.001)	(0.001)	(0.001)
Income ($1,000s)	0.219	0.070	0.301	-0.043*	-0.032	-0.065
	(0.201)	(0.225)	(0.357)	(0.026)	(0.025)	(0.049)
H_0: Naive = Nuanced	0.744	0.651	0.332	0.784	0.801	0.217
Observations	2,989	1,914	1,076	2,989	1,914	1,076
R-squared	0.033	0.041	0.063	0.026	0.034	0.047

Notes: All specifications include month and year fixed effects as well as indicators for the race of the parent(s), education of the parents(s), family type, number of young children, and the presence of a child age 13 or more (babysitter). The data only includes households with young (under age 13) children and an annual income between $10,000 and $50,000. Both parents must earn income in two-parent households to be included in the data. Child-care expenditure as a percentage of annual income is calculated as 26 times the reported two-week child-care expenditure divided by annual income and multiplied by 100. P-values from Wald tests for the difference between Naive and Nuanced diff-in-diff estimates are reported. Robust standard errors in parentheses: $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$

Chapter 3

The Not-so Marginal Value of Weather Warning Systems

3.1 Introduction

From 2004-2013, weather-related hazards caused over 600 fatalities and 3,300 injuries per year in the United States. Tornadoes exceed hurricanes as the leading cause of storm-related deaths and injuries in the U.S., accounting for 1,091 fatalities and 12,407 injuries over this window.[1] Healy and Malhorta (2009) report that from 1985-2004, the United States federal government spent $195 million per year on disaster preparedness. It is often argued that disaster preparedness is woefully underfunded in the sense that the number of deaths and injuries prevented per dollar exceed the returns on many other government investments purportedly targeted at saving lives. To date this claim has been largely speculative due to the lack of robust statistical evidence regarding the benefits of weather warning systems. The main contribution of this paper is to provide the first causal estimates of the impacts of a weather warning system.

A large number of case studies and household surveys exists regarding how individuals respond to various warnings. While this literature very useful for understanding how individuals responded to the situation being examined, the counterfactual is not clear. How would people have fared without a particular warning system? If warnings from a NOAA weather radio were not available, would individuals have received and responded to warnings from a siren, television, or phone? Perhaps different warning systems convey different information which impacts how

[1] U.S. Natural Hazard Statistics, http://www.nws.noaa.gov/om/hazstats.shtml.

effectively individuals protect themselves.

This paper addresses the difficulty in obtaining a reliable counter-factual by exploiting variation in the initial broadcast dates of the radio transmitters for the National Oceanic and Atmospheric Administration's Weather Radio All Hazards (NWR). Because the timing and location of these installations is non-random, two different identification strategies are used to estimate the causal impact of NWR broadcasts on tornado deaths and injuries. The first is a cross-sectional comparison of deaths and injuries across all tornadoes. Variation in the percentage of the tornado's path covered by NWR broadcasts is conditional on a large number of controls including population, housing stock, and the order of transmitter installation. The second method is a county-level fixed-effects analysis which examines how fatality and injury rates change after broadcasts begin.

Both methods estimate that NWR broadcasts causally reduce fatalities and injuries by as much as 40% or more. Robustness checks confirm intuition that NWR broadcasts do not reduce property damage. Note that these results reflect the aggregate reduction in deaths and injuries due to the NWR broadcast and not the risk reduction for a single individual. Receiving the NWR broadcast requires the individual to own a NWR receiver.

Focusing the first impact evaluation of a weather warning system on NWR is appropriate for several reasons. In addition to the availability of quality data and a clear source of identifying variation, NWR has long been a flagship warning system of the National Weather Service. This research also points the way for further revealed preference impact evaluations in the weather community. Such studies are essential for enabling policy makers to choose optimal level of investment both within and between warning systems. This paper and others like it can shed objective light on hotly debated questions such as whether new types of warning systems, such as warnings delivered directly to cellular phones, have made older systems obsolete.

The rest of the paper proceeds as follows. Section II discusses the roll-out of the NWR program and existing knowledge on the effectiveness of weather warning systems. Section III discusses the data and methodology used in this study. Section IV discusses results and robustness checks. Section V concludes.

3.2 Background

3.2.1 The History of NWR

Federal government involvement in weather forecasting in the United States can be traced back to 1870 when a new national weather service was created within the U.S. Army. Sergeant John P. Finley published the first investigations into tornado forecasting, twice per day categorizing weather conditions across large areas of the United States as favorable or unfavorable for tornadoes. 28% of the favorable conditions in Finley (1884) produced confirmed tornadoes somewhere within their broad geographic area. Due to concerns that incited panic would outweigh the benefits of these forecasts, the use of the work tornado was banned from official forecasts from 1887 until 1938 (Coleman et al., 2011). Because of this policy, the first successful forecast of a potential tornadic event at a more precise location did not occur until a 1948 forecast by Major Ernest J. Fawbush and Captain Robert C. Miller at Tinker Air Force Base in Oklahoma was made possible by a "fortuitous series of events" (Grice et al., 1999). Following this much-lauded forecast's success, the Weather Bureau authorized public tornado alerts in 1950 and began issuing tornado forecasts in 1952.

The means of disseminating these warnings to the public has evolved over time and is well documented. "During the 1950s and 1960s, tornado warnings were disseminated to the public primarily by commercial television and radio stations. The TV and radio stations received these warnings from the [U.S. Weather Bureau] by telephone or teletype." (Coleman et al., 2011). Outdoor warning sirens, originally designed as World War II air-raid sirens, were re-purposed for weather warnings. Operated by a local emergency manager, outdoor sirens have also been used for tornado warnings since about 1970.

The National Oceanic and Atmospheric Administration's Weather Radio All Hazards (NWR) is a network of radio stations which broadcast weather information from the nearest National Weather Service (NWS) office. The first NWR transmitters were installed in New York City (Jan. 1, 1953) and Chicago (Apr. 1, 1953) to broadcast aviation weather. After aviation broadcasts moved to L/MF radio stations, the stations were became available for marine service. Over 1966 and

1967, nine additional coastal stations were added to support the maritime community. Partially in response to the 1974 "Super Outbreak" of 148 tornadoes in the 24 hours between 1:00pm EST, April 3rd and 1:00pm EST, April 4th, a January 1975 White House policy statement designated NOAA Weather Radio as the sole government-operated radio system to provide direct warnings into private homes for natural disasters. Figure 3.1 shows how the distribution of NWR transmitters across the country has changed over time, starting immediately before this policy change in 1975. The main concern in using initial broadcast dates as a source of identifying variation is that the areas which receive transmitters first may be more or less prone to fatalities and injuries from observably similar tornadoes.[2] The methodology section addresses this concern directly.

A special receiver is required to pick up the NWR signal, as the signal is broadcast in the very high frequency range of the radio spectrum. These receivers, which are produced by several private companies, are widely available in retail stores, and online. At present receivers cost about $20 or more depending on features. NWS does not manufacture, sell, or endorse any particular receivers. NWS does recommend several common features, including a tone alarm which may be activated even when the audio is turned off. Several important changes have occurred over the years of NWR use. In 1998, NWR began incorporating Specific

[2] In addition to more populated areas receiving transmitters first, in some cases collective action by concerned citizens was the impetus for transmitter installation. Following the Super Outbreak in 1974, citizens of Huntsville, AL raised money to buy their own transmitter which was then donate it to NWS. See http:www.srh.noaa.govhun?n=stationhistory.

Area Message Encoding (SAME), which enables users to select the county(s) or sub-county(s) for which they desired to receive warnings rather than receiving all warnings issued by the transmitter. A 2002 FCC rule change declared that "All existing and new models of EAS equipment manufactured after August 1, 2003 must be capable of selectively displaying and logging messages with state and local event codes." (FCC 02-64). This change enabled users to opt out of receiving certain types of warnings. For example, an individual who lives on a high floor of building many miles from the coast might choose to opt out of coastal flood warnings. Both the introduction of SAME codes and the FCC rule change gave individuals more control over which warnings they receive. If individuals opt out of important warnings, this could decrease the effectiveness of the warning system. But if individuals opt of less relevant warnings, they may pay more attention to those warnings they do receive. The FCC expressed confidence that users "will exercise good judgement in making these choices."

Another common concern is that new warning systems will cause older systems to become outdated. In 2006, Congress passed the "Warning, Alert, and Response Network (WARN) Act" establishing a system of emergency alerts sent to all phones using commercial cell phone towers. NWS began participation in this service in late June of 2012. There is not yet enough data to determine if this program has had altered the impact of NWR transmitters on tornado fatalities or injuries. The recently created Integrated Public Alert and Warning System bring

NWR, phone warnings, and other public alert systems under a single interface.

3.2.2 Current Literature

A large number of case studies and surveys examine how individuals respond to warnings in various situations. Balluz et al. (2000) found that roughly 45% of those who responded to a random telephone survey following March 1, 1997 tornadoes in Arkansas reported seeking shelter after learning of the tornado warning. Liu et al. (1996) survey in two Alabama areas after tornado warnings to learn the type of warning respondents heard first. Dow and Cutter (1998) survey residents about how they responded to repeated hurricane evacuation orders over the course of a single season. Case studies such as these are immensely helpful in determining what warnings individuals hear and how those warnings are perceived. But they are not helpful for determining the causal impact of warning systems, because they often suffer from serious selection bias problems and the counter-factual is unclear. Hence this paper answers a very different type of question: How would people have fared without a particular warning system?

While there has been little writing on the causal impact of warning systems directly, some papers have attempted to estimate the causal benefits of warnings themselves. Doswell et al. (1999) show that shortly after the beginning of public tornado forecasting there was a reduction in deaths relative to inflation-adjusted damage from major tornadoes. Improvements in forecast ability also matter. Simmons and Sutter (2008) find that longer lead times on tornado warnings reduce injuries, and Simmons and Sutter (2005) find that the installation of Doppler radar in the 1990's reduced fatalities and injuries by 45% and 40%, respectively. Sutter and Simmons (2014) find no measurable impact of tornado watches (issued when conditions are favorable for tornadoes) beyond the benefits already granted by warnings.

Various types of mitigating infrastructure have also received much attention. Czajkowski and Simmons (2014) find that building codes reduce hail damage. Merrell et al. (2002) perform a cost-benefit analysis of tornado shelters. Analysis of the benefits of mitigation infrastructure extends outside the United States. Smyth et al. (2004) discuss a variety of potential improvements in for a representative Turkish apartment building's earthquake resilience, each of which are found to be cost-efficient in the long run due to prevented fatalities.

While much research suggests the returns to investments in disaster mitigation may be large, "Governments do not routinely collect or monitor spending on disaster prevention" (Bank and Nations, 2010). Healy and Malhorta (2009) find that $1 of preparedness spending is worth $15 of future damage mitigation, yet U.S. voters reward political parties for disaster relief spending but not disaster preparedness spending. This is perhaps due to disaster preparedness spending by governments having low salience relative to disaster relief spending. There is some evidence that individuals are willing to pay for more salient direct receipt of disaster mitigation. Simmons et al. (2002) find both greater structural integrity and storm blinds increase the sale value of homes in a Gulf Coast city. And large numbers of individuals purchase NWR receivers.

3.3 Data & Methodology

3.3.1 Data

This paper uses data from the Storm Prediction Center's national tornado archive to examines over 58,000 tornadoes recorded between 1950-2014.[3] Each tornado observation has deaths and injuries reported by state, along with the date, time, and properties of the tornado such as start and end location, path length, and tornado width. Fujita scales are reported, with Enhanced Fujita scale reported after January 2007.[4] Property damage is reported beginning in 1996.[5] Affected counties are recorded by FIPS code.

Much gratitude is due to NOAA for directly providing the date on which each transmitter began broadcasting, as well as the handful of dates that any transmitters were permanently deactivated. The database which was used to gather this information was not developed until the 1990's. Installation dates prior to that time were gathered by NOAA from old records so some measurement error is possible. Annual county-level population data comes from the U.S. Decennial Census and the U.S. Census Bureau's Intercensal Estimates. Data on state-by-decade

[3] Only data from U.S. states are used in this analysis. Records from Puerto Rico and the Virgin Islands are excluded.

[4] The Fujita Scale is a measure of tornado intensity based on the damage indicators to buildings, trees, and infrastructure.

[5] Categorical damages ($\leq$ \$50, \$50-\$500, \$500-\$5,000, etc.) is available before 1996 but is missing in some cases so is not used.

housing stocks come from the Historical Census of Housing. Summary statistics can be found in Table 3.1 and Table 3.2.

3.3.2 Methodology

This paper estimates the causal impact of NWR transmitters on tornado injuries and fatalities by comparing outcomes between tornadoes which cross through areas with various levels of coverage by NWR transmitters. Differences in coverage are due to the tornado location and whether the tornado occurs before or after transmitter begin broadcasting. Because both the location and timing of transmitter installation is non-random, many important controls are included and discussed below. The identification assumption is that conditional on these controls, transmitter broadcast areas are only correlated with tornado injuries and fatalities through their transmission of warnings.

It is well-known that log-linear regression models produce biased estimates when the dependent variable has a binding lower bound of zero (Silva and Tenreyro (2006)). Here the outcomes of interest are the number of deaths and injuries attributed to each tornado. Both Poisson and negative binomial regressions analysis are commonly used to model count data (non-negative integers without an explicit upper limit). Each involves different assumptions about the conditional variance of the error term.[6] Because both the fatality and injury statistics have a large

[6]The Poisson distribution assumes the mean is equal to the variance, while different negative binomial distributions make a variety of alternative assumptions about the functional form of the variance, with Poisson being a special case of a negative binomial distribution. The properties discussed here are well known. See Winkelmann (2008) or Cameron and Trivedi (2013).

number of 0's, the data suffers from overdispersion – the variance is greater than the mean. The Poisson regression still produce consistent estimates regardless of any overdispersion as long as there exists some β such that, given a set of variables X, the conditional mean of fatalities and injuries can be correctly expressed as

$$E[Y|X] = e^{\beta X} \tag{3.1}$$

Not all negative binomial distributions have this property. For some negative binomial distributions, unbaised estimation of β requires correct specification of an additional variance parameter. One method of resolving this issue in certain types of negative binomial distributions is to use the two-step estimation process discussed in Wooldridge (2010). Poisson regression analysis is used here because it is simpler and negative binomial regressions have been found to provide no improvement over Poisson for identification in finite samples with overdispersion Blackburn (2014). Coefficients are estimated using the pseudo-maximum likelihood technique of Silva and Tenreyro (2010).

The dependent variable, Y_i, is a count of fatalities, injuries, or in some cases property damage from tornado i. The expected outcome conditional on all controls, X_i, can be expressed as

$$\begin{aligned} ln(E[Y_i|X_i]) =& \beta_1(\text{Coverage})_i + \gamma_1(\text{Properties})_i + \gamma_2(\text{Location})_i \\ &+ \gamma_3(\text{Timing})_i \end{aligned} \tag{3.2}$$

Coverage is a [0,1] treatment intensity variable representing the percentage of counties on a tornado's path which receive broadcasts from at least one NWR transmitter.[7] Properties$_i$ is a vector of tornado properties which accounts for tornadoes which have longer paths being mechanically more likely to enter a county with a broadcasting transmitters. It contains a quadratic measure of path length and the number of counties hit.

Location$_i$ is a vector of characteristics of the location of the tornado to account for transmitters being more likely to be built in certain areas such as densely populated areas or areas with certain types of housing infrastructure. This vector contains state fixed effects, several controls for population, and decade-by-state measures of the percentage and number of residences of different housing types.[8] To account for further unobserved reasons transmitters may be installed in certain areas first, it contains the date at which the first transmitter on the tornado's path was installed. Location$_i$ also contains the area of each of the three largest impacted counties is controlled for to account for correlation between a county's physical area and its probability of having a transmitter. Finally, Location$_i$ contains controls for the number and size of tornadoes a county has historically received, to account for areas more prone to tornadoes being more likely to have NWR transmitters.[9]

[7]Literally, $\frac{\text{number of treated counties on path}}{\text{number of counties on path}}$. In the present draft, transmitter broadcast areas are measured at the county level. Measurement error exists to the extent that a NWR transmitter's broadcast area does not fully cover cover a county.

[8]Specifically, population controls include a cubic polynomial and log of total population among the impacted counties, as well as county-level measures of density for each of the three most densely populated impacted counties. This last measure is interacted with indicators for the number of counties impacted control for the fact that most tornadoes only impact a single county. Limiting to three categorical indicators avoids collinearity issues, as less than 0.25% of tornadoes strike four or more counties. Housing types include detached homes, attached homes, two-to-four unit-per-building homes, five or more unit-per-building homes, and mobile homes.

[9]Specifically, controls for the total number of tornadoes and the number received in the last 5 years for each of the three most-impacted counties on each tornadoes' path are included.

Finally, Timing_i is a vector of variables related to the timing of the tornado to control for changes over time in the probability that tornadoes will be recorded. Year fixed effects control for coverage being higher in later years when fatalities or injuries may have decreased for other reasons.[10] While variables such as the month and time of day a tornado occurs are correlated with fatalities and injuries, there is no obvious reason that these are correlated with coverage and hence they are not included.

Even though both location and timing information are available for each tornado and are used in Equation 3.2, the analysis is at heart cross-sectional with each tornado as an independent observation rather than a county-level panel. I also constructs a county-level panel to control for all temporally constant differences between counties. This approach is limited by the fact that fatalities and injuries are measured at the state level rather than the county level. In order to correctly allocate outcomes, this second methodology examines only tornadoes which impacted a single county, and only counties where fixed effects are not collinear with treatment status across observed tornadoes. This accounts for 89.3% of the tornadoes used in the cross-sectional analysis. This analysis also requires at least one non-zero outcome within each county, so that the final sample is 72.1% of the cross-sectional sample for injuries, and 25.8% of the cross-sectional sample for fatalities. Hence these results are driven by the subset of counties which receive

[10]Year fixed effects begin in 1975, because prior to that point all areas are untreated. See section 2.1.

more dangerous tornadoes more frequently. The cost of resolving any concern about fixed differences between counties is some external validity.

Under this panel FE analysis, the expected outcome of a tornado in county j at time t, conditional on all controls, $X_{j,t}$, can be expressed as

$$\begin{aligned} ln(E[Y_{j,t}|X_{j,t}]) = &\beta_2(\text{Coverage})_{j,t} + \theta_1(\text{Properties})_{j,t} + \alpha_j + \theta_2(\text{Location})_{j,t} \\ &+ \theta_3(\text{Timing})_t \end{aligned} \tag{3.3}$$

Because these are single-county tornadoes, Coverage is a binary treatment indicator representing the whether the county was receiving broadcasts from at least one NWR transmitter at the time of the tornado. Properties$_{j,t}$ again contains a quadratic measure of path length and indicators for Fujita or Enhanced Fujita scale values. Because county FE, α_j, are included, Location$_{j,t}$ includes only time-varying area characteristics. These include the same controls for population and housing types found in Equation 3.2. Timing$_t$ again includes year fixed effects.

3.4 Results

It is easiest to first interpret the results from Equation 3.3. The coefficient on Coverage is approximately equal to the percent reduction in death and injuries caused by having a NWR transmitter broadcasting over the impacted county. Table 3.3 reports that having a NWR transmitters broadcasting over a county causally reduces injuries by approximately 40.2%, with a 95% confidence interval of (3.0%, 77.4%). Similarly, having an NWR transmitters broadcasting over a county causally reduces fatalities by 77.5%, with a 95% confidence interval of (12.4%, 142.6%). While these ranges are quite broad, it is clear that NWR transmitters cause a statistically and economically significant reduction in the injuries and fatalities caused by tornadoes.

For the cross-sectional analysis of Equation 3.2, interpretation is similar. However, 8.6% of tornadoes in the cross-sectional sample impact multiple counties, so Coverage measures intensity of treatment rather than a binary treatment indicator. The coefficient on Coverage is then approximately equal to the percent reduction in death and injuries caused by having a NWR transmitter(s) broadcasting over all impacted counties relative to none. Table 3.3 reports that having a NWR transmitter(s) broadcasting over all impacted counties causally reduces injuries by 38.6%, with a 95% confidence interval of (6.7%, 70.5%). Similarly, having an NWR transmitters broadcasting over a county causally reduces fatalities by 46.9%, with a 95% confidence interval of (5.2%, 88.6%). While these ranges are again quite broad, it remains clear that NWR transmitters cause a statistically and economically significant reduction in the injuries and fatalities caused by tornadoes. Applying the cross-sectional analysis to the panel sample yields very similar estimates to the panel estimates, suggesting differences are driven by external validity concerns rather than unresolved selection bias.

Because a tornado warning are often given with relatively short notice, it is likely difficult for NWR users to move or protect physical capital such as homes after hearing a warning. It should be expected that while NWR may reduce injuries and fatalities, it is unlikely to reduce property damage. As a robustness check, property damage is examined as the dependent variable. Analysis focuses on data beginning in 1996 because property damages are reported in more detail at this point. Table 3.3 reports that having a NWR transmitter(s) broadcasting over all impacted counties has no statistically or economically significant impact on property damage under either specification.

3.5 Conclusion

This paper provides the first estimates of the causal impact of a natural disaster warning system on fatalities and injuries. Two different identification strategies found the presence of National Oceanic and Atmospheric Administration's Weather Radio All Hazards (NWR) transmitters causally reduces deaths and injuries by as much as 40% or more. While the 95% confidence intervals are quite broad, it is clear that NWR transmitters cause a statistically and economically significant reduction in fatalities and injuries. These are results are consistent with the widespread belief that the return on investment for disaster warning and mitigation systems is quite large relative to other public health and safety investments.

3.6 Acknowledgments

Chapter 3, in part is currently being prepared for submission for publication of the material. Miller, Benjamin M. The dissertation author was the primary investigator and author of this material.

3.7 Figures and Tables

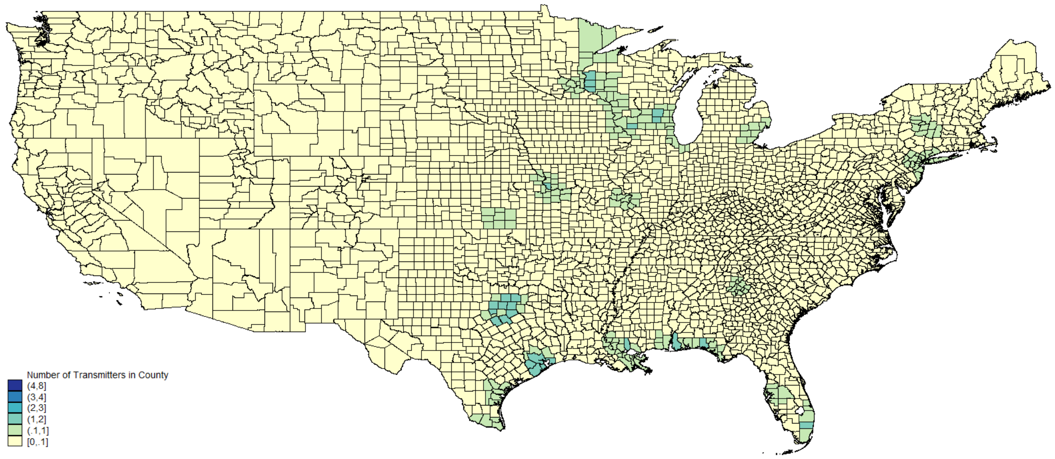

(a) Jan. 1, 1975

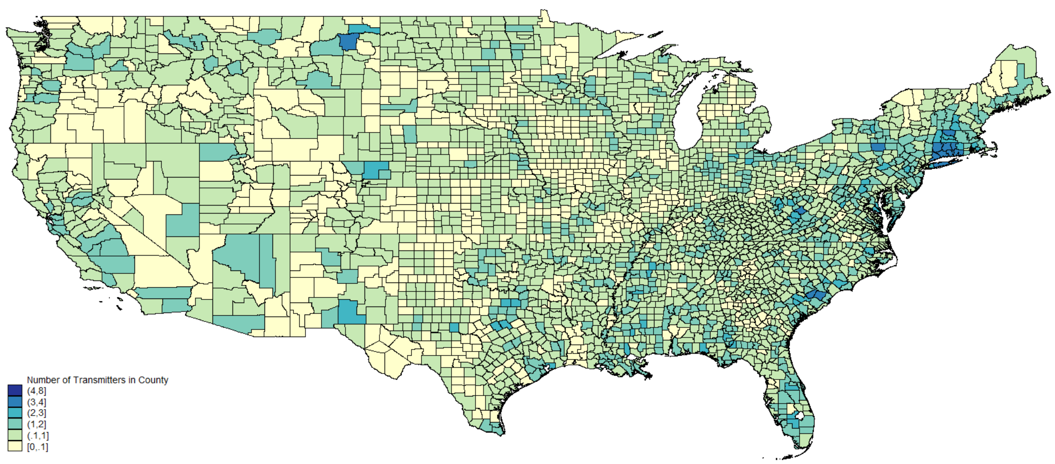

(b) Jan. 1, 1990

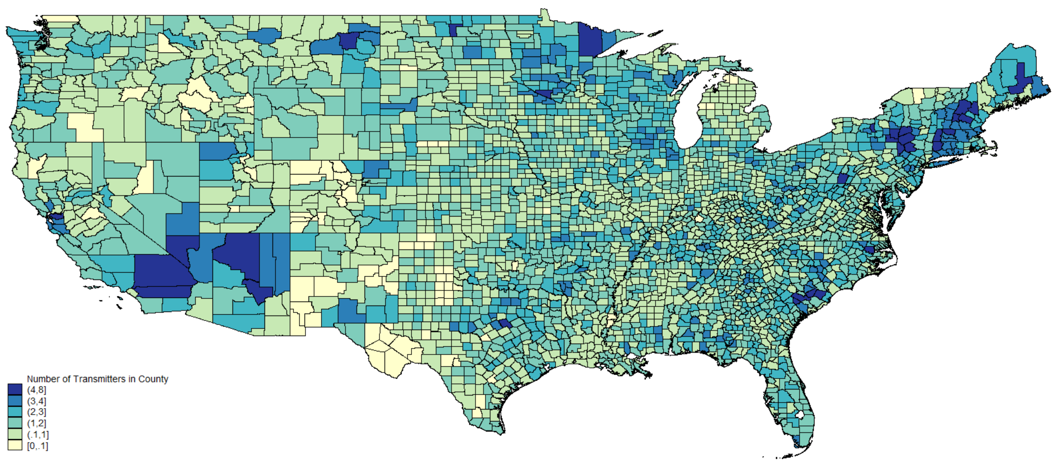

(c) Jan. 1, 2005

Figure 3.1: Counties Receiving NWR Broadcasts, by Date

Initial dates of Transmitter broadcasts provided directly by NOAA.

Table 3.1: Summary Statistics, Full Sample (Cross Sectional Sample)

	Coverage>0			Coverage=0		
	mean	sd	count	mean	sd	count
Injuries	1.120948	16.51491	35197	2.253604	22.73667	23651
Fatalities	.0668239	1.233654	35197	.1450679	1.948917	23651
Property Loss	2.776952	46.64212	11261	.4283062	1.643588	418
Length	3.087201	6.824699	35197	4.086531	10.62586	23651
Number of Counties Hit	1.10856	.4099218	35197	1.123843	.4879203	23651
Total Population of all Counties Hit	145314.1	433335.4	35197	79731.27	273864.4	23651
Detached House Count	1759899	1375400	35197	1263014	1062010	23651
Detached House Percent	.6604029	.0695362	35197	.7476643	.0899188	23651
Attached House Count	111848.7	151275.9	35197	60603.4	102912.9	23651
Attached House Percent	.0346784	.0271764	35197	.0309022	.0291586	23651
Two to Four Units Per Building Count	215433.7	213146.6	35197	176543.1	182093.2	23651
Two to Four Units Per Building Percent	.074879	.0276843	35197	.1008895	.0464029	23651
Five or More Units Per Building Count	473474.7	543494.1	35197	210764.8	344002	23651
Five or More Units Per Building Percent	.1406211	.0509969	35197	.0865989	.0550293	23651
Mobile Home Count	226188.8	215244.8	35197	65529.05	115990.3	23651
Mobile Home Percent	.0855338	.0421663	35197	.0329253	.0297342	23651
Land Area of Largest County Hit	1027.963	1204.346	35197	1058.408	1224.175	23650
Land Area of Second Largest County Hit	627.3252	412.6389	3010	629.6503	344.764	2022
Land Area of Third Largest County Hit	545	369.5758	557	575.6165	271.7003	558
Number of Pre-Coverage EF3+ Tornadoes in County with most such Tornadoes	1.229885	.5851879	522	1.331807	.6971672	4587
Number of Pre-Coverage EF3+ Tornadoes in County with second most such Tornadoes	1.066667	.2537081	30	1.138249	.3590918	217
Number of Pre-Coverage EF3+ Tornadoes in County with third most such Tornadoes	1	.	1	1.05	.2207214	40
Number of recorded Tornadoes in impacted County with most recorded Tornadoes	43.74427	36.41301	35197	40.04744	29.6819	23651
Number of recorded Tornadoes in impacted County with second most recorded Tornadoes	28.12392	17.95259	3010	28.28586	18.59084	2022
Number of recorded Tornadoes in impacted County with third most recorded Tornadoes	23.64991	14.58063	557	26.88172	17.67004	558
Number of Recent Tornadoes in impacted County with most Recent Tornadoes	5.283748	5.322252	30002	3.607905	3.225353	17432
Number of Recent Tornadoes in impacted County with second most Recent Tornadoes	3.128275	2.654836	2214	2.277682	1.883136	1156
Number of Recent Tornadoes in impacted County with third most Recent Tornadoes	2.546218	2.055745	357	2.013793	1.606588	290

This table contains data from the full sample of tornadoes in counties which eventually receive transmitters. Data from the NWS' Storm Prediction Center's national tornado archive, NOAA records, the Decennial Census and intercensal population estimates, and the Historical Census of Housing.

Table 3.2: Summary Statistics, Single-County Tornadoes (Panel Sample)

	Coverage>0			Coverage=0		
	mean	sd	count	mean	sd	count
Injuries	.4863386	4.926286	31073	1.161963	14.01464	21579
Fatalities	.021369	.3299059	31073	.0709023	1.31849	21579
PropertyLoss	.9093769	11.33192	9294	.400449	1.615476	392
Length	1.914879	3.296652	31073	2.1341	4.404986	21579
Number of Counties Hit	1	0	31073	1	0	21579
Total Population of all Counties Hit	140345.5	435484.1	31073	74088.55	267118	21579
Detached House Count	1774794	1395133	31073	1280915	1085771	21579
Detached House Percent	.6604267	.0687774	31073	.7468861	.0896655	21579
Attached House Count	111679.9	148308.5	31073	61561.04	103256.1	21579
Attached House Percent	.0346856	.0267405	31073	.0310228	.0291305	21579
Two to Four Units Per Building Count	213104.5	203696.6	31073	175825.3	181171.3	21579
Two to Four Units Per Building Percent	.0742635	.0267959	31073	.0996429	.0455856	21579
Five or More Units Per Building Count	477627	538406	31073	216233.1	351497.5	21579
Five or More Units Per Building Percent	.1415159	.0501818	31073	.0875406	.0553538	21579
Mobile Home Count	229184.2	220860.8	31073	68173.22	119716.9	21579
Mobile Home Percent	.085186	.0418049	31073	.0338209	.0302697	21579
Land Area of Largest County Hit	1043.422	1226.766	31073	1078.105	1266.797	21579
Number of Pre-Coverage EF3+ Tornadoes in County	1.215584	.5931609	385	1.324427	.6879235	3930
Number of recorded Tornadoes in impacted County	45.12564	37.13972	31073	39.97678	29.97021	21579
Number of Recent Tornadoes in impacted County	5.361686	5.450039	26617	3.607606	3.241706	15803

This table contains data from the sample of single-county tornadoes in counties which eventually receive transmitters. Data from the NWS' Storm Prediction Center's national tornado archive, NOAA records, the Decennial Census and intercensal population estimates, and the Historical Census of Housing.

Table 3.3: The Causal Impact of NWR Transmitters

	(1) log(Injuries)	(2) log(Injuries)	(3) log(Fatalities)	(4) log(Fatalities)	(5) log(Property Loss)	(6) log(Property Loss)
Coverage	-0.386*	-0.402*	-0.469*	-0.775*	0.907	0.0580
	(0.163)	(0.190)	(0.213)	(0.332)	(0.581)	(0.456)
Path Length	X	X	X	X	X	X
Population	X	X	X	X	X	X
County Tornado History	X	X	X	X	X	X
Housing Stock	X	X	X	X	X	X
Year FE	X	X	X	X	X	X
Installation Order	X		X		X	
State FE	X		X		X	
Number of Counties Hit	X		X		X	
County Area	X		X		X	
County FE		X		X		X
N	58492	43578	57738	16524	11648	9214

Coverage is defined as the percentage of impacted counties which receive transmitter broadcasts at the time of the tornado. Controls as described in Section 3.3.2.

$^{*}p < 0.05, ^{**}p < 0.01, ^{***}p < 0.001.$

Bibliography

Abler, David G., George S. Tollow, and G. K. Kripalani, *Technical Change & Income Distribution in Indian Agriculture*, Westview Press, 1994.

Adhvaryu, Achyuta, A. V. Chari, and Siddharth Sharma, "Firing Costs and Flexibility: Evidence from Firms' Employment Responses to Shocks in India," *The Review of Economics and Statistics*, 2013.

Agarwal, Sumit, Chunlin Liu, and Nicholas S Souleles, "The Reaction of Consumer Spending and Debt to Tax Rebates–Evidence from Consumer Credit Data," *Journal of Political Economy*, 2007, *115* (6).

Altonji, Joseph G and Christina H Paxson, "Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-offs," *Journal of Labor Economics*, 1988, *6* (2), 254–276.

Auffhammer, Maximilian, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel, "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change," *Review of Environmental Economics and Policy*, 2013, *7* (2).

Averett, Susan L, H Elizabeth Peters, and Donald M Waldman, "Tax Credits, Labor Supply, and Child Care," *Review of Economics and Statistics*, 1997, *79* (1), 125–135.

Balluz, Lina, Laura Schieve, Talmage Holmes, Stephanie Kiezak, and Josephine Malilay, "Predictors for peoples response to a tornado warning: Arkansas, 1 March 1997," *Disasters*, 2000, *24* (1), 71–77.

Bank, The World and The United Nations, "Natural Hazards, Unnatural Disasters: The Economics of Effective Prevention," 2010.

Bartolome, Charles AM De, "Which Tax Rate do People Use: Average or Marginal?," *Journal of Public Economics*, 1995, *56* (1), 79–96.

Beegle, Kathleen, Joachim De Weerdt, and Stefan Dercon, "Migration and Economic Mobility in Tanzania: Evidence from a Tracking Survey," *The Review of Economics and Statistics*, 2011.

Bhalotra, Sonia, "Fatal fluctuations? Cyclicality in infant mortality in India," *Journal of Development Economics*, 2010.

Björkman-Nyqvist, Martina, "Income shocks and gender gaps in education: Evidence from Uganda," *Journal of Development Economics*, 2013.

Blackburn, McKinley L., "The Relative Performance of Poisson and Negative Binomial Regression Estimators," *Oxford Bulletin of Economics and Statistics*, 2014.

Blau, David, "Child Care Subsidy Programs," in Robert A. Moffitt, ed., *Means-Tested Transfer Programs in the United States*, University of Chicago Press, 2003, pp. 291–364.

Blau, David M and Philip K Robins, "Child-care costs and family labor supply," *The Review of Economics and Statistics*, 1988, *70* (3), 374–381.

Brodeur, Abel, Mathias Lé, Marc Sangnier, and Yanos Zylberberg, "Star Wars: The Empirics Strike Back," 2013. http://anon-ftp.iza.org/dp7268.pdf.

Brouwer, C. and M. Heibloem, *Irrigation Water Management: Irrigation Water Needs* Food and Agricultural Organization 1986. Training Manual No. 3.

Brückner, Markus and Mark Gradstein, "Effects of transitory shocks to aggregate output on consumption in poor countries," *Journal of International Economics*, 2013.

Burgess, Robin and Dave Donaldson, "Can Openness Mitigate the Effects of Weather Shocks? Evidence from India's Famine Era," *American Economic Review: Papers & Proceedings*, 2010.

Burke, Marshall and Kyle Emerick, "Adaptation to Climate Change: Evidence from US Agriculture," 2013. http://www.ocf.berkeley.edu/~kemerick/burke_emerick_2013.pdf.

Cabrera, Victor E., David Letson, and Guillermo Podest, "The value of climate information when farm programs matter," *Agricultural Systems*, 2007.

Cameron, A. Colin and Pravin K. Trivedi, *Regression Analysis of Count Data*, second ed., Cambridge University Press, 2013.

Chetty, Raj, "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 2012, *80* (3), 969–1018.

_ , Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Ciccone, Antonio, "Economic Shocks and Civil Conflict: A Comment," *American Economic Journal: Applied Economics*, 2011, *3* (4).

Coleman, Timothy A., Kevin R. Knupp, James Spann, J. B. Elliott, and Brian E. Peters, "The History (and Future) of Tornado Warning Dissemination in the United States," *Bulletin of the American Meteorological Society*, 2011, *92* (5).

Conlin, Michael, Ted O'Donoghue, and Timothy J. Vogelsang, "Projection Bias in Catalog Orders," *American Economic Review*, 2007, *97* (4).

Connelly, Rachel, "The Effect of Child Care Costs on Married Women's Labor Force Participation," *The Review of Economics and Statistics*, 1992, *74* (1), 83–90.

Connolly, Marie, "Here Comes the Rain Again: Weather and the Intertemporal Substitution of Leisure," *Journal of Labor Economics*, 2008, *26* (1).

Cramer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler, "Labor Supply of New York City Cabdrivers: One Day at a Time," *Quarterly Journal of Economics*, 1997, *112* (2).

Czajkowski, Jeffrey and Kevin M Simmons, "Convective Storm Vulnerability: Quantifying the Role of Effective and Well-Enforced Building Codes in Minimizing Missouri Hail Property Damage," *Land Economics*, 2014, *90* (3), 482–508.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, "What Do We Learn from the Weather? The New Climate-Economy Literature," *Journal of Economic Literature (forthcoming)*, 2014.

Dickens, William T and Shelly J Lundberg, "Hours Restrictions and Labor Supply," *International Economic Review*, 1993, *34* (1), 169–192.

Doswell, Charles A., Alan R. Moller, and Harold E. Brooks, "Storm Spotting and Public Awareness since the First Tornado Forecasts of 1948," *Weather and Forecasting*, 1999, *14* (4).

Dow, Kirstin and Susan L Cutter, "Crying wolf: Repeat responses to hurricane evacuation orders," *Coastal Management*, 1998, *26.*

Duflo, Esther, William Gale, Jeffrey Liebman, Peter Orszag, and Emmanuel Saez, "Saving Incentives for Low-and Middle-Income Families: Evidence from a Field Experiment with H&R Block," *The Quarterly Journal of Economics*, 2006, *121* (4), 1311–1346.

Finkelstein, Amy, "E-ztax: Tax Salience and Tax Rates," *The Quarterly Journal of Economics*, 2009, *124* (3), 969–1010.

Finley, Jno. P., "Tornado Predictions," *American Meteorological Journal*, 1884, *1* (3).

Gallagher, Kelly Sims and Erich Muehlegger, "Giving Green to Get Green? Incentives and Consumer Adoption of Hybrid Vehicle Technology," *Journal of Environmental Economics and Management*, 2011, *61* (1), 1–15.

Golden, Lonnie, "Flexible Work Schedules Which Workers Get Them?," *American Behavioral Scientist*, 2001, *44* (7), 1157–1178.

Grice, G. K., R. J. Trapp, S. F. Corfidi, R. Davies-Jones, C. C. Buonanno, J. P. Craven, K. K. Droegemeier, C. Duchon, J. V. Houghton, R. A. Prentice, G. Romine, K. Schlachter, and K. K. Wagner, "The golden anniversary celebration of the first tornado forecast," *Bulletin of the American Meteorological Society*, 1999, *80* (7), 1341–1348.

Grimard, Franque and Barton Hamilton, "Estimating the elderly's returns on the farm: evidence from Cote d'Ivoire," *Journal of Development Economics*, 1999, *58* (2).

Healy, Andrew and Neil Malhorta, "Myopic Voters and Natural Disaster Policy," *American Political Science Review*, 2009, *103* (3).

Herbst, Chris M, "The Labor Supply Effects of Child Care Costs and Wages in the Presence of Subsidies and the Earned Income Tax Credit," *Review of Economics of the Household*, 2010, *8* (2), 199–230.

Hidalgo, F. Daniel, Suresh Naidu, Simeon Nichter, and Neal Richardson, "Economic Determinants of Land Invasions," *The Review of Economics and Statistics*, 2010, *92* (3).

Hole, Arne Risa, "Calculating Murphy-Topel variance estimates in Stata: A simplified procedure," *The Stata Journal*, 2006, *6* (4).

Hornbeck, Richard, "Nature versus Nurture: The Environments Persistent Influence through the Modernization of American Agriculture," *American Economic Review: Papers & Proceedings*, 2012.

Ito, Koichiro, "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing," *The American Economic Review*, 2014, *104* (2), 537–563.

Jacoby, Hanan G. and Emmanuel Skoufias, "Risk, Financial Markets, and Human Capital in a Developing Country," *Review of Economic Studies*, 1997.

_ **and** _ , "Testing theories of consumption behavior using information on aggregate shocks: Income seasonality and rainfall in rural India," *American Journal of Agricultural Economics*, 1998, *80* (1).

Jappelli, Tullio and Luigi Pistaferri, "The Consumption Response to Income Changes," *Annual Review of Economics*, 2010, *2.*

Jayachandran, Seema, "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries," *Journal of Political Economy*, 2006.

Jensen, Robert, "Agricultural Volatility and Investments in Children," *American Economic Review*, 2000, *90* (2).

Johnson, David S, Jonathan A Parker, and Nicholas S Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *The American Economic Review*, 2006, *96* (5), 1589–1610.

Kochar, Anjini, "Smoothing Consumption by Smoothing Income: Hours-of-Work Responses to Idiosyncratic Agricultural Shocks in Rural India," *The Review of Economics and Statistics*, 1999.

Koopmans, Tjalling C., "Identification Problems in Economic Model Construction," *Econometrica*, 1949, *17* (2).

Kurukulasuriya, Pradeep and Robert Mendelsohn, "Crop switching as a strategy for adapting to climate change," *African Journal of Agricultural and Resource Economics*, 2008.

Liebman, Jeffrey B and Richard J Zeckhauser, "Schmeduling," *Harvard KSG Working Paper*, 2004.

Liu, Simin, Lynn E Quenemoen, Josephine Malilay, Eric Noji, Thomas Sinks, and James Mendlein, "Assessment of a severe-weather warning system and disaster preparedness, Calhoun County, Alabama, 1994.," *American Journal of Public Health*, 1996, *86* (1), 87–89.

Maccini, Sharon and Dean Yang, "Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall," *American Economic Review*, 2009, *99* (3).

Mazzocco, Maurizio and Shiv Saini, "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," *American Economic Review*, 2012, *102* (1).

Mendelsohn, Robert and Ariel Dinar, "Climate Change, Agriculture, and Developing Countries: Does Adaptation Matter?," *World Bank Research Observer*, 1999.

Merrell, David, Kevin M. Simmons, and Daniel Sutter, "The Determinants of Tornado Casualties and the Benefits of Tornado Shelters," *Land Economics*, 2002, *81* (1).

Miguel, Edward, "Poverty and Witch Killing," *Review of Economic Studies*, 2005, *72* (4).

_ **, Shanker Satyanath, and Ernest Sergenti**, "Economic Shocks and Civil Conflict: An Instrumental Variables Approach," *Journal of Political Economy*, 2004, *112* (4).

Mjelde, James W., Troy N. Thompson, and Clair J. Nixon, "Government Institutional Effects on the Value of Seasonal Climate Forecasts," *American Journal of Agricultural Economics*, 1996.

Mohan, S. and N. Arumugam, "CROPES: A Rule-Based Expert System for Crop Selection in India," *Transactions of the American Society of Agricultural Engineers*, 1994, *37* (3).

Moran, Emilio F., Ryan Adams, Bryn Bakoyéma, Stefano T. Fiorini, and Bruce Boucek, "Human strategies for coping with El Niño related drought in Amazônia," *Climate Change*, 2006, *77.*

Munshi, Kaivan, "Networks in the modern economy: Mexican migrants in the US labor market," *Quarterly Journal of Economics*, 2003, *118* (2).

Murphy, Kevin M. and Robert H. Topel, "Estimation and Inference in Two-Step Econometric Models," *Journal of Business & Economic Statistics*, 1985, *3* (4).

Newhouse, David, "The Persistence of Income Shocks: Evidence from Rural Indonesia," *Review of Development Economics*, 2005, *9* (3).

O'Brein, Karen, Robin Leichenko, Ulka Kelkar, Kenry Venema, Guro Aandahl, Heather Tompkins, Akram Javed, Suruchi Bhadwal, Stephan Barg, Lynn Nygaard, and Jennifer West, "Mapping vulnerability to multiple stressors: climate change and globalization in India," *Global Environmental Change*, 2004.

Orlove, Benjamin S., John C. H. Chiang, and Mark A. Cane, "Forecasting Andean rainfall and crop yield from the influence of El Niño on Pleiades visibility," *Nature*, 2000, *403.*

Pagan, Adrian, "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economic Review*, 1984.

Palmer, T. N., "Extended-Range Atmospheric Prediction and the Lorenz Model," *Bulletin American Meteorological Society*, 1993.

Parthasarathy, B., H. F. Diaz, and J. K. Eischeid, "Prediction of All-India Summer Monsoon Rainfall With Regional and Large-Scale Parameters," *Journal of Geophysical Research*, 1988, *93* (D5).

Paxson, Christina H., "Using Weather Variability To Estimate the Response of Savings to Transitory Income in Thailand," *American Economic Review*, 1992, *82* (1).

Rasmusson, Eugene M. and Thomas H. Carpenter, "The Relationship Between Eastern Equatorial Pacific Sea Surface Temperatures and Rainfall over India and Sri Lanka," *Monthly Weather Review*, 1983, *111.*

Rodó, X., E. Baert, and F. A. Comin, "Variations in seasonal rainfall in Southern Europe during the present century: relationships with the North Atlantic Oscillation and the El Niño-Southern Oscillation," 1997.

Roll, Richard, "Orange Juice and Weather," *American Economic Review*, 1984, *74* (5).

Roncoli, Carla, Keith Ingram, and Paul Kirshen, "Reading the Rains: Local Knowledge and Rainfall Forecasting in Burkina Faso," *Society & Natural Resources*, 2002, *15.*

Rose, Elaina, "Consumption Smoothing and Excess Female Mortality in Rural India," *The Review of Economics and Statistics*, 1999, *81* (1).

_ , "Ex ante and ex post labor supply response to risk in a low-income area," *Journal of Development Economics*, 2001, *64.*

Rosenzweig, Cynthia and Martin L. Perry, "Potential Impact of Climate Change on World Food Supply," *Nature*, 1994, *367.*

Rosenzweig, Mark and Christopher Udry, "Forecasting Profitability," 2013. NBER Working Paper No. 19334.

Rosenzweig, Mark R. and Kenneth I. Wolpin, "Natural 'Natural Experiments' in Economics," *Journal of Economic Literature*, 2000, *38* (4).

Seo, S. Niggol and Robert Mendelsohn, "An analysis of crop choice: Adapting to climate change in South American farms," *Ecological Economics*, 2008.

Shapiro, Matthew D and Joel Slemrod, "Consumer Response to Tax Rebates," *American Economic Review*, 2003, *93* (1), 381–396.

Silva, J.M.C. Santos and Silvana Tenreyro, "The Log of Gravity," *The Review of Economics and Statistics*, 2006, *88.*

Silva, JMC Santos and Silvana Tenreyro, "On the existence of the maximum likelihood estimates in Poisson regression," *Economics Letters*, 2010, *107* (2), 310–312.

Simmons, Kevin M. and Daniel Sutter, "WSR-88D Radar, Tornado Warnings, and Tornado Casualties," *Weather and Forecasting*, 2005, *20* (3).

_ **and** _ , "Tornado Warnings, Lead Times, and Tornado Casualties: An Empirical Investigation," *Weather and Forecasting*, 2008, *23* (2).

_ **, Jamie Brown Kruse, and Douglas A. Smith**, "Valuing Mitigation: Real Estate Market Response to Hurricane Loss Reduction Measures," *Southern Economic Journal*, 2002, *68* (3).

Slemrod, Joel and Jon M Bakija, *Taxing Ourselves: A Citizen's Guide to the Debate over Taxes*, MIT Press, 2008.

Smyth, Andrew W, Gulay Altay, George Deodatis, Mustafa Erdik, Guillermo Franco, Polat Gulkan, Howard Kunreuther, Hilmi Lus, Esra Mete, Nano Seeber, and Ozal Yuzugullu, "Probabilistic benefit-cost analysis for earthquake damage mitigation: Evaluating measures for apartment houses in Turkey," *Earthquake Spectra*, 2004, *20* (1), 171–203.

Sovey, Allison J. and Donald P. Green, "Instrumental Variables Estimation in Political Science: A Readers Guide," *American Journal of Political Science*, 2011, *55* (1).

Stockdale, T. N., D. L. T. Anderson, and J. O. S. Alves, "Global seasonal rainfall forecasts using a coupled ocean-atmosphere model," *Nature*, 1998, *392.*

Suri, Tavneet, "Selection and Comparative Advantage in Technology Adoption," *Econometrica*, 2011.

Sutter, Daniel and Kevin M. Simmons, "Preparing for Danger: On the Impact of Tornado Watches on Tornado Casualties," *International Journal of Mass Emergencies and Disasters*, 2014, *32* (1).

Taraz, Vis, "Adaption to Climate Change: Historical Evidence from the Indian Monsoon," 2012. http://www.econ.yale.edu/~vt48/Vis_Taraz_jmp.pdf.

Trombley, Mark A., "Stock Prices and Wall Street Weather: Additional Evidence," *Quarterly Journal of Business and Economics*, 1997, *36* (3).

Tubiello, Francesco N., Jean-François Soussana, and S. Mark Howden, "Crop and pasture response to climate change," *Proceedings of the National Academy of Sciences*, 2007, *104* (50).

Walker, Thomas S. and James G. Ryan, *Village and Household Economies in India's Semi-arid Tropics*, The John Hopkins University Press, 1990.

Wallace, John M. and Shawna Vogel, "El Niño and Climate Prediction - Reports to the Nation on Our Changing Climate," Technical Report, University Corporation for Atmospheric Research 1994. National Oceanic and Atmospheric Administration (NOAA) Award No. NA27GP0232-01.

Wang, Bin, ed., *The Asian Monsoon*, Praxis Publishing, 2006.

Winkelmann, Rainer, *Econometric Analysis of Count Data*, fifth ed., Springer, 2008.

Wolpin, Kenneth I., "A New Test of the Permanent Income Hypothesis: The Impact of Weather on the Income and Consumption of Farm Households in India," *International Economic Review*, 1982, *23* (3).

Wooldridge, J.M., *Econometrics Analysis of Cross Section and Panel Data*, second ed., The MIT Press, 2010.

Working, E. J., "What Do Statistical 'Demand Curves' Show?," *Quarterly Journal of Economics*, 1927, *41* (2).

Yang, Dean and HwaJung Choi, "Are Remittances Insurance? Evidence from Rainfall Shocks in the Philippines," *The World Bank Economic Review*, 2007, *21* (2).