Three Essays in Economics

by

Matthew George Hall

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2016

Doctoral Committee:

Professor George A. Fulton, Co-Chair
Professor Christopher L. House, Co-Chair
Professor Charles C. Brown
Professor Susan M. Collins
Assistant Research Scientist Gabriel M. Ehrlich

to my family

ACKNOWLEDGEMENTS

I owe a debt of gratitude to my advisor, employer, co-chair, and friend George Fulton. Without his help I would never have been able to make the final push necessary to complete this dissertation. George is retiring next year, and the University of Michigan as well as the Research Seminar in Quantitative Economics will be poorer for it. George, thank you so much for everything.

I would also like to thank the other members of my committee, Christopher House, Charles Brown, Susan Collins, and Gabriel Ehrlich. I picked the easiest marks, the people who had never failed to help me when I asked for advice, or to listen when I requested their ear. Their willingness to sign up for this thankless task without hesitation, and on short notice, means the world to me.

For putting up with the aforementioned final push that lasted longer than I had any right to demand, and for the daily encouragement that made it bearable, I thank Jennifer Lambert.

My friends were a crucial support network in Michigan these years, both the ones who left and the one who stayed. I have made lifelong friendships with Collin Raymond, Christina DePasquale, Joshua Montes, and Ozan Jaquette. Josh did extra duty as a co-author, motivator, and informal advisor, for which he has my gratitude. My colleagues at RSQE, especially Joan Crary and Daniil Manaenkov, were a pleasure to work with and to learn from.

A special thank you is owed to my co-author, Aditi Thapar. Her willingness to communicate at all hours to meet my deadlines and to babysit me through learning Eviews were essential to the creation of a quality paper in a tight time frame.

David Ratner, Ryan Nunn, Adam Cowing, Seb Moosapoor, and Michigan resident Justin Gayle - you dragged me to the end. Thank you for everything, and I'll see you in Vegas.

With Gabriel Ehrlich. I don't know if I could have asked any more of you. Not much else to say about that, but RSQE has the right man in the right job, and if it gets too crazy, just ask and I'll be back to help.

Lastly, to my family - Barbara, Brad, Brian, Liz, Sam, Hannah, Shadow, and Maddie - thank you for everything. This was not easy, nor short. Thank you for noticing, and for refraining from making any jokes to that effect after year 7 or so. And thank you for believing this would always happen.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Are Entry Wages Really (Nominally) Flexible?

1.1 Introduction

Downward nominal wage rigidity is often hypothesized to amplify unemployment fluctuations by constraining the responsiveness of wages to negative shocks. There is considerable evidence that the wages of incumbent workers are downwardly rigid, but the wages of new hires appear to be significantly more flexible. Because entry wages determine job creation over the business cycle, a substantial literature argues that downward nominal wage rigidity (hereafter, wage rigidity) is unlikely to explain unemployment dynamics.[1]

In this paper we argue that the apparent flexibility of entry wages is an artifact of selection bias. If unemployed workers are heterogeneous in their ability or willingness to reduce their reservation wages, those with more flexible reservation wages will be more likely to become re-employed. Because new hires will be disproportionately workers with flexible wages, the observed wages of new hires will appear flexible. The unobserved reservation wages of the workers who are not hired, however, may be quite rigid.

We estimate worker wage elasticities with respect to aggregate labor productivity and unemployment in the Panel Study of Income Dynamics and the Current Population Survey. We confirm the

This chapter is co-authored with Gabriel Ehrlich and Joshua Montes. I would like to thank Charles Brown, Susan Collins, George Fulton, Christopher House, Pawel Krowlikowski, Ryan Nunn, David Ratner, and Matthew Shapiro for helpful comments. All errors are our own.

[1] See, for instance, Pissarides (2009), Haefke et al. (2013), and Kudlyak (2014).

consensus in the literature that wages appear to be more elastic for new hires than for incumbents. We contrast this evidence with histograms of nominal wage changes from the same survey data, which exhibit substantial downward nominal wage rigidity in all years for both long-time workers (hereafter job stayers) and for workers with recent spells of non-employment (job finders).

We construct a search and matching model of the labor market and show that hiring wages can appear flexible even if unemployed workers' reservation wages are quite rigid. We estimate the parameters of the model using indirect inference and find substantial nominal rigidity for both job stayers and finders. Using simulated data from the estimated model, we show that the elasticities of observed wages closely resemble those in the data for job stayers and job finders: finders display more elastic wages than stayers. Those elasticities are not targets of the model estimation. The model's ability to generate disparate wage elasticities among job stayers and job finders stems naturally from the selection bias inherent in conditioning the sample on observed wages. Pooling unemployed workers' reservation wages with the observed wages of job finders brings the elasticity of the wages of potential new hires substantially closer to that of job stayers.

Dynamic simulations of the model clarify the mechanisms by which aggregate observed wages appear more responsive to labor market conditions than the underlying levels of wage rigidity would imply. Although the observed wages of job finders fall sharply in response to a negative shock, the reservation wages of unemployed workers remain rigid. This asymmetry highlights the pitfalls of selecting on successful job finding to measure wage responsiveness to aggregate economic conditions among potential new hires. Layoffs rise immediately in response to a negative shock, followed by a persistent decrease in the job finding rate, the majority of which can be attributed to the rigid reservation wages of unemployed workers. The importance of wage rigidity in our model to flows out of unemployment re-establishes its potential as an explanation for observed unemployment volatility.

At least since Shimer's (2005) demonstration that the canonical search and matching model of the labor market with perfectly flexible wages cannot replicate the observed volatility in unemployment, a large literature has explored whether adding some form of wage rigidity can help to reconcile the model to the data. Prominent examples include Hall (2005), who introduces real wage rigidity via a bargaining norm between workers and employers, Gertler and Trigari (2009), who model wage bargaining with staggered multi-period contracts, and Christiano et al. (2015), who endogenously derive wage rigidity from alternating offers in bargaining negotiations.

Several empirical studies, however, show that the wages of new hires are much more responsive to labor market conditions than the wages of longer-tenured workers. Bils (1985), Shin (1994), Solon et al. (1994), Devereux (2001), and Shin and Solon (2006) all find lower elasticities of wages with respect to the unemployment rate for tenured workers than for all workers. Bils (1985), Shin (1994), and Barlevy (2001) find specifically that the elasticity for those in new matches is much higher than the estimates for incumbent workers. Furthermore, Haefke et al. (2013), after correcting for composition bias in worker subgroups, obtain much higher elasticities in the aggregate wages of new hires with respect to average labor market productivity than for the worker population generally. As Pissarides (2009) summarizes the evidence, "Time-series or panel studies on the cyclical volatility of wages show considerable stickiness, but this evidence is dominated by wages in ongoing jobs and is not relevant for job creation in the search and matching model."[2]

Notably, the time series evidence contrasts starkly with the direct survey evidence on unemployed workers' reservation wages reported by Krueger and Mueller (2014). They find that "...self-reported reservation wages decline at a modest rate over the spell of unemployment..." They argue that their evidence suggests that "...many workers persistently misjudge their prospects or anchor their reservation wage on their previous wage." We argue that a model with heterogeneity in the rigidity of unemployed workers' wages resolves the apparent contradiction between these two sources of evidence.

The remainder of the paper proceeds as follows. In section 1.2, we establish within two different data sources the key empirical elasticities regarding entry wage rigidity. In section 1.3 we employ those same data sources to provide evidence in favor of wage rigidity for both job stayers and job finders. In section 1.4 we introduce a labor search model with explicit downward nominal wage

[2]Elsby et al. (2013) are also skeptical of the role that downward nominal wage rigidity plays in unemployment fluctuations. They find a significant number of nominal wage cuts in CPS data and point out that in the Great Recession the most notable distinction from previous contractions, which occurred in times of higher inflation, is not in separations but in the duration of unemployment. This result puts the onus on entry rigidity to explain the data, a hypothesis for which they find little theoretical and no empirical support.

rigidity for all workers. In section 3.5 we estimate the model and illustrate the results. Section 1.6 concludes.

1.2 Elasticities for Job Stayers, Switchers, and New Hires

Our empirical analysis utilizes longitudinal data from two sources, the Panel Study of Income Dynamics (PSID) and the Current Population Survey (CPS). We use the PSID to conduct analyses similar to those in Solon, Barsky, and Parker (1994) and Devereux (2001), in which we estimate the elasticities of the real wages of job stayers and all employed workers with respect to the unemployment rate. We use the CPS to estimate the elasticity of real wages of all workers and job finders with respect to average labor productivity, in the spirit of Haefke et al. (2013). In both cases, we confirm the qualitative patterns in the original studies: the wages of job stayers are less responsive to the unemployment rate than are the wages of all workers, and the wages of job finders are more responsive to labor productivity than are the wages of all workers.

1.2.1 Panel Study of Income Dynamics

The PSID contains data on employment, salary, and hourly wages for household heads and their spouses. We combine the 1980-1997 annual surveys with the 1999-2013 biannual surveys to construct an employment history for respondents that spans from 1980-2013.[3] The number of respondents in these surveys averages about 12,500 per year.[4] The PSID includes occupational codes and industry codes, as well as job start and end dates, which allows us to determine worker tenure over several years.

Table 1.1 provides a summary description of some key variables in the analysis. About 24 percent of those surveyed are salaried workers, while almost 35 percent are hourly employees. The remainder is disproportionately retired or otherwise out of the labor force. Our analysis focuses on hourly workers, for whom we have wage data, and salaried workers, for whom we construct an

[3] We begin the analysis in 1980 because hourly wages are top-coded at very restrictive levels in the 1978 and prior surveys.

[4] We include spouses in the analysis when we consider the entire universe of working adults. The inclusion of spouses is necessary to ensure that the primary earner in a family is present in the analysis because the PSID codes household heads by gender. We also include some results that are restricted to male heads of household to facilitate comparisons with past studies. That restriction does not change the qualitative results.

hourly wage. Most salaried workers do not provide a consistent measure of hours for their primary job, so we assume a fixed number of hours from year to year; this assumption seems reasonable for those who stay at the same job from one survey to the next (about 66 percent of salaried workers), but potentially biases the hourly wage for those who switch jobs.

We categorize the set of employed workers as job stayers, job switchers, and job finders.[5] Job stayers are defined as workers who provided a start date at their current job prior to the last time they were surveyed, when available, or who provided a tenure length at their current employer exceeding the time between survey dates if the start date is unavailable. In addition, job stayers must have had continuous employment between survey dates without spells of unemployment or time out of the labor force. Job switchers are defined as workers who maintained continuous employment, defined by no months in which they were unemployed or out of the labor force, but who provided a start date between survey dates. Job finders are defined as workers who were employed at the time of both the current and prior surveys, but who report having spent time between surveys as either unemployed or out of the labor force.[6]

We estimate regressions of the form:

$$\Delta \ln w_{it} = \beta_0 + \beta_1 t + \beta_2 \Delta U_t + \beta_3 X_{it} + \varepsilon_{it}, \tag{1.1}$$

in which w_{it} is the nominal hourly wage, U_t is the national unemployment rate, and X_{it} is a polynomial measure of work experience and job tenure for worker i at year t. These regressions follow Devereux (2001), who in turn builds on Solon et al. (1994). We use both the dependent variable in Devereux (2001), which is limited to earnings in the worker's primary job, as well as that in Solon et al. (1994), who use all earnings in the surveyed year. In addition, although Solon et al.

[5]Although our categorization involves some subjective judgment, which may induce misclassification, we will classify workers using identical definitions in our estimation of the theoretical model. This method of indirect inference allows us to correct for possible misspecification in these definitions. See section 1.5.1 for more details. Our categorization is a partition of all workers who have valid wages in consecutive surveys.

[6]Note that this definition excludes first-time entrants and re-entrants who spent multiple years unemployed or out of the labor force. Before concluding this section, we want to reiterate that in both datasets our definition of job finders excludes new entrants to the labor market. According to our CPS dataset, over the past 20 years, new entrants constitute between 6-12 percent of the unemployed looking for work and average only 8.6 percent of the total unemployed during that time.

(1994) include only men, we include both men and women in our sample.[7]

Table 1.2 presents the results of our regressions. Although our sample periods differ, and our inclusion of women alters our sample slightly compared to Solon et al.'s, our analysis recovers the same basic fact pattern as the earlier studies. When we confine our sample to the years in which we have annual surveys (1980 to 1997), we estimate that the elasticity of all wages with respect to the annual unemployment rate is -0.83, versus -0.55 for job stayers. When we extend the analysis to include job switchers and job finders, we estimate elasticities of -1.80 and -1.82, respectively, far larger (in absolute value) than for all workers or job stayers. Extending the sample forward in time to incorporate the biannual surveys from 1999 to 2013 yields uniformly smaller elasticities, but the relative magnitudes among all workers, job stayers, and job finders are unaffected.

Our specifications using the average annual wage give similar qualitative results. When we end the sample at 1997, we estimate that the elasticity of real wages with respect to unemployment is -0.7 for all workers and -0.49 for job stayers. When we analyze job finders separately, we estimate an elasticity of -1.65, consistent with the notion that job finders' wages are more responsive to labor market conditions than the wages of other workers. When we extend the sample to include years through 2013, we see somewhat higher elasticities, but the relative magnitudes among all workers, job stayers, and job finders are unaffected.

1.2.2 Current Population Survey

The CPS's outgoing rotation group includes data on wages and weekly hours worked for all members of a household, job codes, and demographic information, along with data from the basic monthly files that include the employment history of each household member. We use the merged datasets of Drew et al. (2014), who build on the methodology of Madrian and Lefgren (2000) to link CPS responses from the same individual longitudinally for up to 16 months, the maximum time that individuals are covered by the survey. The structure of the CPS is such that individuals are surveyed for four consecutive months, are out of the survey for eight months, and then are surveyed again

[7] Both papers use a two-step process to address potential bias in their standard errors due to common time effects across workers. We address this issue by clustering standard errors by year and estimate the regression in equation 1.1 directly.

for four more months. This design allows us to construct a year-over-year measure of wage changes for individuals in the outgoing rotation groups. We combine the 1989-2013 monthly surveys with aggregate labor market data such as the unemployment rate, labor productivity, CPI-U measure of inflation, and the private consumption implicit price deflator.

Following the procedure of Haefke et al. (2013), we restrict the dataset to nonfarm, nonsupervisory, private-sector workers, trim outliers in hours worked and in implied hourly earnings, impute top-coded earnings according to the procedure in Schmitt (2003), and use typical hours worked per week as our divisor in the determination of the hourly wage for salaried workers.[8] We construct a variable for years of school and a standard measure of potential experience (age minus years of school minus 6). Finally, we include dummy variables for female, black, Hispanic, and married with spouse present in our list of covariates.

Haefke et al. (2013) investigate the effect of changes in productivity on aggregate real wages of both job finders and all workers. To account for composition bias in each group, they remove demographically explainable wage determinants for all workers via a Mincer regression, and then analyze the effects of labor market indicators on the respective residualized aggregate wages. They define a job finder, or new hire, as any worker who had an unemployment spell in the prior 3 months. We adopt that definition of a job finder in the CPS, as we do not have employment information for the 8 months prior to those 3.[9] Therefore, our definition of job finder is more restrictive in the CPS than in the PSID, which allows for the measurement of nonemployment spells further from the survey date.

Haefke et al. (2013) estimate the following specification by quarter for each subgroup j of workers from 1984 to 2006:[10]

$$\Delta \ln w_{jt} = \beta_{0,j} + \beta_{1,j} \Delta \ln y_t + \varepsilon_{jt}, \tag{1.2}$$

[8] In contrast, Card and Hyslop (1997) and Elsby et al. (2013) use usual weekly earnings for salaried workers.

[9] See Drew et al. (2014) for a detailed discussion of the structure of the CPS.

[10] They exclude 1995q3 and 1995q4 from analysis because of a change in sample design that makes it difficult to match workers, add quarter dummies to account for residual seasonality, and add a dummy for 2003q1 to reflect the change in occupation classification in 2003 that increases the fraction of supervisory workers.

where w_{jt} is their residualized real wage series for group j and y_t is a measure of labor productivity. Haefke et al.'s preferred specification deflates wages by the BLS private nonfarm business sector implicit deflator and uses aggregate labor productivity as their aggregate labor market indicator of interest.[11]

We follow the spirit of Haefke et al.'s analysis, with the major exception that we use the longitudinal aspect of the CPS to calculate the year-over-year change in wages for individual workers and use those changes as the outcome of interest in our analysis. We choose this approach over the residualization approach of Haefke et al. (2013) because it is more consistent with our analysis using the PSID data and because using within-worker wage changes corrects for unobservable compositional changes among worker groups in addition to observable ones. On the other hand, this approach restricts the sample of workers who could potentially be considered new hires, and requires year-over-year as opposed to quarter-over-quarter comparisons. This restriction also means that we include individuals only once, the second time they reach the outgoing rotation group. Additionally, we use the private consumption deflator as our measure of inflation (using the CPI-U has almost no effect on the results). We therefore run the following regressions for job finders and all workers:

$$\Delta \ln w_{ijt} = \beta_{0,j} + \beta_{1,j} \Delta \ln y_t + \beta_{2,j} X_{it} + \varepsilon_{it}, \tag{1.3}$$

where an observation is a worker i in worker group j and calendar quarter t, and the first difference terms represent year-on-year changes. X_{it} is a cubic polynomial in experience, consistent with our analysis using the PSID data.

Table 1.3 presents the results of our regressions, along with some key results from Haefke et al. (2013) for comparison. Whereas Haefke et al. (2013) estimate that the elasticities of the real wages of all workers and job finders with respect to average labor productivity are 0.24 and 0.79, respectively, we estimate those elasticities as 0.29 and 0.55. Part of the reason that we estimate a lower elasticity for job finders is our inclusion of more recent calendar years, which in our analysis

[11]The aggregate labor productivity series published by the Bureau of Labor Statistics has been revised subsequent to the publication of their paper, which complicates the comparison of our results to theirs.

of the PSID data led to less estimated responsiveness of wages to labor market conditions.

Overall, we view our results as qualitatively similar to Haefke et al.'s: the wages of new hires are more responsive to productivity changes than are the wages of all workers.[12] More broadly, the estimates in this section confirm the key patterns in the literature: the observed real wages of job finders are roughly twice as responsive to changes in labor market conditions as the real wages of all workers, which are themselves more responsive than the wages of job stayers.

1.3 Estimating Downward Nominal Wage Rigidity for Job Stayers, Job Switchers, and New Hires

In this section, we employ a complementary approach to those in section 1.2 to measuring wage rigidity. We focus particularly here on downward nominal wage rigidity by attempting to measure the proportion of wage cuts that would have occurred in a counterfactual environment with perfectly flexible wages that are "missing" from the observed data. We estimate this proportion using the method in Ehrlich and Montes (2014), which builds on the method of Card and Hyslop (1997). The essence of the method is to extrapolate from the upper half of the observed wage change distribution to what the nominally negative portion of the distribution would look like in the absence of rigidity, and calculate the proportion of counterfactual mass that is missing in the observed distribution.[13] Although such analyses have typically focused on the wages of job stayers, the same method can be extended to job switchers and job finders, as shown in the subsections below.

The key results that emerge from performing this analysis in the PSID and CPS are that the observed wages of job finders exhibit less downward nominal wage rigidity than the wages of job stayers, but the degree of rigidity in job finders' wages is nonetheless substantial. We begin this section by describing the visual evidence for wage rigidity in the PSID and CPS, before providing formal estimates of the degrees of nominal rigidity.

[12]Again, we note that our definition of new hires is restricted to workers who were employed one year previously, which is not the case for Haefke et al. (2013).
[13]The appendix describes the method in detail.

1.3.1 PSID

Our longitudinal dataset using survey data from the PSID includes 144,047 observations of (constructed) hourly wage data matched to the same worker in consecutive surveys over 25 surveys, the majority of the observations being job stayers. Figure 1.1 illustrates these data in histograms of one-year and two-year percent wage changes for job stayers, job switchers, and job finders for survey years 1980-2013. The histograms are truncated at -50 and 50 percent, with a dotted vertical line to indicate a zero percent nominal wage change.[14]

The first row of histograms in figure 1.1 contain a spike in the proportion of reported wage changes at nominal zero, which we interpret as one of the hallmarks of downward nominal wage rigidity. There is also a visually evident asymmetry between the nominally positive and nominally negative portions of the distribution, as the proportions of nominally negative wage changes are smaller than a simple extrapolation from the nominally positive portion of the distribution would indicate. This asymmetry is especially evident in the 2-year changes in the histogram for job stayers, which contain sample years with lower inflation, among other factors, leading to smaller nominal wage increases. Negative wage changes have nevertheless remained relatively uncommon; thus, the distribution of wages has "piled up" against the barrier at nominal zero.

The second and third rows of figure 1.1 show similar histograms for job switchers and job finders, respectively. The dispersion of wage changes is higher for both job switchers and job finders than for job stayers. Furthermore, the median wage change for job finders, at 3.7%, is lower than for job stayers, which is 4.6%. Nevertheless, the histograms share significant similarities with those for stayers: 1) a spike at nominal zero, and 2) asymmetry between the nominally positive and nominally negative portions of the wage change distribution. There is less mass in the nominally negative portion than would be implied by a symmetrical wage change histogram, consistent with the idea that the wages of job switchers and job finders exhibit some degree of downward nominal rigidity.

[14]Individual-year histograms of wage changes for stayers, switchers, and finders are also shown in the appendix. Each year exhibits the same basic pattern, with a spike at nominal 0 wage change.

1.3.2 CPS

Figure 1.2 provides wage change histograms for all workers, all workers excluding job finders, and job finders only in the CPS outgoing rotation groups for the years 1989 to 2013. We censor the histograms at -80% and 80% because the sample sizes are larger than in the PSID.[15] The sample is limited to those workers who can be matched between their 4th and 16th months in the survey. We divide these workers into job finders and non-job finders, a combination of job stayers and job switchers, as described in section 1.2.2. There are 23,600 total wage changes on average per surveyed year, of which approximately 900 per year come from workers we classify as job finders.

The histograms in figure 1.2 are qualitatively similar to those in figure 1.1. There are large spikes at nominal zero and an asymmetry between the nominally positive and negative portions of the distributions, with the latter displaying "missing" mass. Again, the wage change distributions of job finders display weaker, but still suggestive, evidence of downward nominal rigidities than the distributions for other workers.

1.3.3 Systematic Measurement of Wage Rigidity

In this section we provide formal estimates of the proportion of nominal wage cuts that would have occurred in an environment with no nominal wage rigidity that are instead prevented by downward nominal wage rigidity. The basic approach, described in detail in the appendix, is to construct an empirical distribution of log nominal wage changes and reflect the 50-100th percentile of changes back on the 0-50th percentiles. The implied share of nominal wage cuts that would be expected based on the upper half of the wage change distribution is compared with the actual share of nominal wage cuts. The statistic:

$$\widehat{wr} = 1 - \frac{\hat{F}^{obs}(0^{-})}{\hat{F}^{cf}(0^{-})} \tag{1.4}$$

[15]Year-by-year histograms for each category of worker are found in the appendix. The histograms for 1995 are omitted because a change in sampling design does not permit matches of worker wages to their employment history.

represents the fraction of wage changes that are "missing", where $\hat{F}^{cf}(0^{-})$ is the estimated counterfactual distribution of log wage changes and $\hat{F}^{obs}(0^{-})$ is the empirical distribution of log wage changes.

This statistic reflects the combination of two phenomena associated with downward nominal wage rigidity. First, it captures the extent to which slightly negative nominal wage changes are "swept up" to nominal 0. Second, it may also capture the share of workers who would have received a wage cut in an environment with flexible wages but who instead separated from their employer either through a layoff or a quit.

Table 1.4 displays our estimates of the proportion of nominal wage cuts prevented by wage rigidity both in the PSID and in the CPS. The table presents estimates for the PSID for job stayers, job switchers, and job finders for the years 1980 to 2013, and for the CPS for all workers, non-job finders, and job-finders for the years 1989 to 2013. The table also displays separate estimates for salaried and hourly workers.

We estimate that 52.7% of counterfactual nominal wage cuts are missing among job stayers in the PSID, versus 56.1% and 36.4%, respectively, for job switchers and job finders. The estimates for salaried and hourly workers do not differ systematically: salaried job stayers display less wage rigidity than hourly job stayers, but salaried job switchers and finders display more wage rigidity than their hourly counterparts.

The estimates using the CPS data are qualitatively similar. We estimate that among all workers, 47.4% of counterfactual nominal wage cuts were prevented by wage rigidity, versus 38.6% for job finders. Again, the estimate for salaried and hourly workers do not vary in a consistent fashion.

One potential limitation of using reported survey data on nominal wages over time to estimate wage changes is that respondents may round their hourly wage to the nearest dollar or half dollar, or their salary to the nearest 1000 dollar value.[16] This rounding could lead to an overstatement of the number of unchanged nominal wages from year to year. An inflated number of unchanged wages could bias our measure of rigidity if small wage cuts disappear due to rounding. To examine the

[16] See, for instance, Altonji and Williams (1997).

potential effect such rounding has on our results, we re-estimate wage rigidity for hourly workers after excluding all round-dollar results (about one third of the sample for incumbents and finders, slightly less for switchers). Reassuringly, the percent of counterfactual wage cuts prevented by wage rigidity decreases only from 55.1% (see table 1.4) among job stayers to 53.6%, with that third of respondents excluded from the analysis, while rigidity among job switchers declined from 50.2% to 48.4%. Finders registered a wage rigidity measure of 35.4%, versus 36.6% including the entire sample of hourly employees.

We interpret these estimates as indicating a substantial amount of downward nominal wage rigidity for workers in the United States economy. Importantly, although the wages of job finders appear to exhibit less rigidity than the wages of other workers, they are by no means perfectly flexible. Therefore, these results stand in some contrast to the results in section 1.2, which indicated much more responsiveness of the wages of job finders to labor market conditions than the wages of job stayers. In the next section, we build a search and matching model of the labor market that attempts to reconcile these results, and implies that the apparent flexibility of the wages of job finders stems from composition bias in the pool of newly hired workers.

1.4 Model

We consider a general equilibrium model with search and matching in the labor market that is closely related to the canonical model of Mortensen and Pissarides (1994). However, we model wage setting differently than those authors do. We assume, as in Barattieri et al. (2010) and Daly and Hobijn (2014), that workers set their wages unilaterally. We further follow Daly and Hobijn (2014) by adopting a Calvo-style (1983) process: we assume that in any given period, a fraction of workers are constrained from reducing their nominal wages. The model allows the fractions of employed and unemployed workers who are prevented from reducing their wage demands to differ.

1.4.1 Model Environment

We consider the stationary equilibrium of a discrete time model with no aggregate shocks but with shocks to a worker's idiosyncratic productivity and their ability to reduce their nominal wage demands each period. Each firm has one job, which can either be vacant or filled and producing output. There is a unit mass of workers who can be either employed in a job or unemployed and searching for a job.

Firms and workers are infinitely lived with a common discount rate β and have linear preferences over profits and consumption, respectively. Workers and firms cannot store goods, thus workers consume their entire incomes each period. There is also no intensive margin of labor supply: workers in a filled job supply exactly one unit of labor, L, each period. Unemployed workers receive an unemployment benefit b each period.

Firms in a match with a worker can decide whether to continue to employ the worker at the worker's demanded wage or to terminate the job. Labor is the only input into production, and the output of a filled job is given by:

$$Y = pL = p \tag{1.5}$$

where p is stochastic and can be conceptualized either as the productivity of a worker or the price of a job's output. We refer to p as productivity in this paper. The per-period profits π of a firm with a worker with productivity p and paying wage w are then:

$$\pi(p, w) = p - w. \tag{1.6}$$

Firms that are not in a match and wish to meet with a worker must post a vacancy at per-period cost c, expressed in units of output. There is free entry into vacancy posting.

Unemployed workers and firms with vacant jobs form matches according to a matching function $m(v, u)$, where v is the number of vacancies and u is the number of workers who are unemployed.[17]

[17] Because we have normalized the number of workers to 1, the number of unemployed is synonymous with the

We assume that the matching function has the Cobb-Douglas form:

$$m(v, u) = Av^{\phi}u^{1-\phi} \tag{1.7}$$

where A is a parameter that governs matching efficiency and ϕ is the elasticity of the matching function with respect to the number of vacancies. Denoting 'labor market tightness' v/u as θ, the probability f that a worker meets a vacancy is $f(\theta) = m(v, u)/u = A\theta^{\phi}$. The probability q that a firm with a vacant job meets an unemployed worker is $q(\theta) = m(v, u)/v = A\theta^{\phi-1}$.

There is no on-the-job search, and matches end with exogenous probability s_x each period. Endogenous separations occur in two ways. First, matches end when the productivity level of the match falls to a low enough level that the match surplus between the worker and firm is exhausted. Those separations are bilaterally efficient. Second, bilaterally inefficient separations occur when the worker is unable to cut his or her nominal wage demand below the maximum level that the firm is willing to pay, but would have been willing to do so in an environment with flexible wages.

We model wage rigidity according to the process in Calvo (1983). Employed and unemployed workers set their reservation wages unilaterally. We assume employed and unemployed workers are unable to reduce their wage demand in any given period, with probabilities λ_E and λ_U, respectively. Firms then decide whether to continue or to terminate matches given workers' wage demands.

The timing of each period is as follows:

1. The period begins and employed and unemployed workers draw realizations of whether they can reduce their reservation wages in the period.
2. Workers draw their idiosyncratic productivity levels, and firms and workers observe workers' productivity levels.
3. Firms post vacancies and matching between vacancies and unemployed workers occurs.

   Not every match between an unemployed worker and a vacancy will result in the formation of a new job both because of exogenous separations and because the worker's wage demand may be higher than the firm will accept. To distinguish between a match and a new employment relationship that enters production, we will call a match between an unemployed worker and a vacant job an *interview*. The probabilities f and q are the likelihoods of an unemployed worker receiving an interview in a given period and of a firm that has posted a vacancy interviewing a worker, respectively.
4. Exogenous separations occur.

unemployment rate, and we will use the two interchangeably.

Note that exogenous separations can occur even in new interviews, such that the worker is never employed by the firm regardless of productivity levels or wage demands.

5. Workers in a negotiation set their wage demands, while unemployed workers re-set their reservation wages.
6. Firms decide whether to accept matched workers' wage demands and proceed to production, or to terminate the relationship.

   We will refer to the process of workers setting their wage demands and firms deciding whether to accept them as a *negotiation*, although there is no actual bargaining involved. Note that from the firm's perspective, there is no difference between a negotiation with a previously unemployed worker and a worker in an ongoing employment relationship. Therefore, we will not generally distinguish between the two.[18]
7. Production occurs, wages and unemployment benefits are paid, profits are earned, and consumption occurs.
8. The period ends.

Firms can therefore be in two different states, with an unfilled vacancy or in a match with a worker. We will denote the values to the firm of being in these states as V and J, respectively. Workers can find themselves in four possible states: unemployed with a flexible wage, unemployed with a rigid wage, employed with a flexible wage, and employed with a rigid wage. We will denote the values of the worker to being in these states as U^F, U^R, W^R, and W^F, respectively. We define the value functions for these states in sections 1.4.2 and 1.4.3.

We assume that a worker's log productivity follows the $AR(1)$ process:

$$\ln p = (1 - \psi_p) \ln \bar{p} + \psi_p \ln p_{-1} + \varepsilon_p, \ \ \varepsilon_p \sim N\left(0, \sigma_p^2\right). \tag{1.8}$$

Productivity is a time-varying, mean-reverting characteristic of the individual worker. Further, a worker's productivity process persists in unemployment. The productivity distribution of employed workers will differ from the distribution for all workers because firms will lay off workers when their reservation wages exceed the cutoff value associated with the worker's productivity.[19]

1.4.2 Firm's Problem

The value to the firm of posting a vacancy, denoted V, is defined in step 3 in the timeline and given as:

$$V = -c + q(\theta)(1 - s_x) \iint J(p, w)\, \mathrm{d}G(p, w) + (1 - q(\theta)(1 - s_x))\beta \mathbb{E}[V'], \tag{1.9}$$

[18] The distinction does matter for calculating employment flows such as job creation and job destruction.

[19] It is unnecessary that a worker's productivity level be higher than the worker's wage in every period due to the associated option value of a match.

where $G(p, w)$ is the stationary joint cumulative distribution of productivity levels and wage demands from unemployed workers. The firm incurs the flow cost c of posting a vacancy and gains the expected value of a negotiation with probability $q(\theta)(1-s_x)$, which accounts for both the likelihood of a match and its survival to become a negotiation, as well as the continuation value of the vacancy conditional on not matching.

The value to the firm of being in a negotiation with a match of productivity p and worker reservation wage w, denoted by J and defined in step 6 in the timeline, is given by:

$$J(p,w) = \max_{\text{discontinue,continue}} \Big\{ \beta \mathbb{E}[V'], \\ p - w + \beta(1-s_x) \iint J(p',w')\, \mathrm{d}F(p'|p)\mathrm{d}H(w'|p',w) \Big\}. \tag{1.10}$$

The firm decides between terminating the match or entering into production with the matched worker. In production the firm receives the flow surplus $p - w$ and the expected continuation value of a filled job conditional on the current period's wage and productivity (inclusive of the risk of an exogenous separation during negotiation next period). $F(p'|p)$ is the cumulative distribution function of next period's productivity level given this period's productivity, and $H(w'|p', w)$ is the cumulative distribution function of next period's wage demands for a worker in a filled job given current wage w and next-period's productivity level p'.[20]

Given equation 1.10, we can define the wage for which the firm is indifferent between continuing and terminating the employment relationship. That cutoff wage, denoted as $\tilde{w}(p)$, is a function of productivity and solves the equation:

$$\beta \mathbb{E}[V'] = p - w + \beta(1-s_x) \iint J(p',w')\, \mathrm{d}F(p'|p)\mathrm{d}H(w'|p',w). \tag{1.11}$$

1.4.3 Worker's Problem

We define the worker's value functions as of step 5 in the model, after matching and exogenous separations have occurred, when the worker must decide his or her reservation wage. The value of being in a negotiation with a flexible wage is a function of this period's productivity level, and we denote it $W^F(p)$. The value of being in a negotiation with a rigid wage depends on both this period's productivity and last period's wage demand, and we denote it $W^R(p, w_{-1})$. It is sometimes convenient to represent expectations of next period's value of being employed, without knowing

[20] The expected value of a match next period can be further decomposed based on the probability of wage adjustment next period, but we omit that characterization here.

whether the worker's wages will be flexible or rigid. We denote this expectation as $\mathbb{E}[W(p', w)] = \mathbb{E}[(1-\lambda_E)W^F(p') + \lambda_E W^R(p', w)]$.

Likewise, the value of being unemployed in step 5 with a flexible reservation wage is a function of this period's productivity only, so we denote it $U^F(p)$. The value of being unemployed with a rigid reservation wage is a function of this period's productivity and last period's wage demand, so we denote it $U^R(p, w_{-1})$. When we wish to denote the expected value of being unemployed next period, we use the notation $\mathbb{E}[U(p', w)] = \mathbb{E}[(1-\lambda_U)U^F(p') + \lambda_U U^R(p', w)]$.

The value to the worker of being in a negotiation with a flexible wage and productivity p is given by:[21]

$$\begin{aligned} W^F(p) = \max_w \Bigg\{ & \mathbb{1}(w \leq \tilde{w}(p))\Big(w + \beta \int \left\{(1-s_x)\mathbb{E}[W(p', w)] + s_x \mathbb{E}[U(p', w)]\right\} \mathrm{d}F(p'|p)\Big) \\ & + \mathbb{1}(w > \tilde{w}(p))\Bigg(b + \beta \int \Big(f(\theta)(1-s_x)\mathbb{E}[W(p', w)] \\ & + (1 - f(\theta)(1-s_x))\mathbb{E}[U(p', w)]\Big) \mathrm{d}F(p'|p)\Bigg)\Bigg\}. \end{aligned} \tag{1.12}$$

The worker chooses his or her wage demand w knowing the firm's cutoff wage given the current productivity level $\tilde{w}(p)$. Choosing a wage demand lower than that cutoff, as in the first term of the value function, yields the demanded wage this period and a continuation value associated with starting the next period in an ongoing match with the firm. Choosing a higher wage leads to a termination of the match, yielding the worker flow payoff b this period and a continuation value associated with starting the next period unmatched with a firm. We denote the wage schedule that solves this maximization problem as $w^\star_{EF}(p)$.

The value to the worker of being in a negotiation with productivity p and downwardly rigid wage w_{-1} is given by:

$$\begin{aligned} W^R(p, w_{-1}) = {} & \mathbb{1}(\frac{w_{-1}}{1+\pi} \leq w^\star_{EF}(p))W^F(p) + \mathbb{1}(\frac{w_{-1}}{1+\pi} > w^\star_{EF}(p)) \times \cdots \\ & \Bigg\{\mathbb{1}(\frac{w_{-1}}{1+\pi} \leq \tilde{w}(p))\Big(\frac{w_{-1}}{1+\pi} + \beta \int \Big[(1-s_x)\mathbb{E}[W(p', \frac{w_{-1}}{1+\pi})] + s_x \mathbb{E}[U(p', \frac{w_{-1}}{1+\pi})]\Big)\Big] \mathrm{d}F(p'|p)\Big) \\ & + \mathbb{1}(\frac{w_{-1}}{1+\pi} > \tilde{w}(p))U^R(p, w_{-1})\Bigg\}. \end{aligned} \tag{1.13}$$

The previous period's real wage, w_{-1}, divided by $1+\pi$, where π is the rate of inflation, represents

[21] We distinguish between being employed and being in a negotiation because an unemployed worker could receive a job match in a period but have a reservation wage higher than the firm will accept. In that case, the worker will not actually be employed in the period.

the new real wage that corresponds with a downwardly rigid nominal wage. The first term in the value function represents the case in which the optimal wage demand is below last period's wage, in which case the problem reduces to the problem of the worker with a flexible wage. In the case that the previous period's wage demand is binding, it may or may not be acceptable to the firm. If the wage demand is acceptable to the firm, the worker receives that wage plus next period's continuation value. If it is not acceptable, the worker receives the payoffs associated with unemployment, defined below. Implicit in $W^R(p, w_{-1})$ is an optimization problem, because workers have freedom to raise their reservation wages. As written, this choice is subsumed in $W^F(p)$. The wage schedule associated with this value function is $w^\star_{ER}(p)$. Note that $w^\star_{ER}(p, w_{-1})$ is the minimum of w_{-1} and $w^\star_{EF}(p)$.

The value to the worker of being unemployed with productivity p and a flexible reservation wage is given by:

$$\begin{aligned} U^F(p) = \max_w \Big\{ & b + \mathbb{E}[f(\theta')](1 - s_x)\beta \int \mathbb{E}[W(p', w)] \, \mathrm{d}F(p'|p) \\ & + (1 - \mathbb{E}[f(\theta')](1 - s_x))\beta \int \mathbb{E}[U(p', w)] \, \mathrm{d}F(p'|p) \Big\}. \end{aligned} \tag{1.14}$$

The unemployed worker receives the unemployment benefit b this period and a continuation value that reflects the probabilities of matching or failing to match next period. Since wages are always flexible upwards, it is optimal for an unemployed worker to set their reservation wage at the minimum possible value, for instance the minimum wage.[22] The reservation wage that solves the unemployed worker's degenerate maximization problem is denoted as $w^\star_{UF}(p)$.

The value function for unemployed workers with productivity p and a downwardly rigid reservation wage w_{-1} is:

$$\begin{aligned} U^R(p, w_{-1}) = b & + \mathbb{E}[f(\theta')](1 - s_x)\beta \int \mathbb{E}[W(p', \frac{w_{-1}}{1+\pi})] \, \mathrm{d}F(p'|p) \\ & + (1 - \mathbb{E}[f(\theta')](1 - s_x))\beta \int \mathbb{E}[U(p', \frac{w_{-1}}{1+\pi})] \, \mathrm{d}F(p'|p). \end{aligned} \tag{1.15}$$

This function follows the pattern of the function in the flexible wage case closely, except that it must

[22] As a result the value function can be expressed using only $W^F(p')$ and $U^F(p')$: even if next period's wage is rigid, the rigidity will never bind strictly. The simplified value function is

$$U^F(p) = b + \mathbb{E}[f(\theta')](1 - s_x)\beta \int W^F(p') \, \mathrm{d}F(p'|p) + (1 - \mathbb{E}[f(\theta')](1 - s_x))\beta \int U^F(p') \, \mathrm{d}F(p'|p).$$

account for the probability that wage rigidity will again be binding next period. Again, note that $w^{\star}_{UR}(p, w_{-1})$ is the larger of w_{-1} and $w^{\star}_{UF}(p)$. Because the latter term is the lowest possible wage, an unemployed worker with a rigid wage will always set this period's reservation wage to equal last period's reservation wage. The wage schedule associated with this value function is $w^{\star}_{UR}(p, w_{-1})$.

1.4.4 Stationary Equilibrium

To define an equilibrium of the model, we must first derive the equations for flows into and out of employment. The matching function dictates the number of unemployed workers who are matched to a vacant job each period, but not all matches will result in a flow into employment, because of exogenous separations and negotiation failures. Given the cumulative distribution of productivity and reservation wages across unemployed workers $G(p, w)$, with corresponding marginal distribution $G_w(w)$, the job creation flow of workers from unemployment to employment is determined by:

$$\text{Jobs Created} = \left[f(\theta)(1 - s_x) \iint^{\tilde{w}(p)} \mathrm{d}G(p, w) \right] u = f(\theta)(1 - s_x)\mathbb{E}[G_w(\tilde{w}(p))]u \tag{1.16}$$

The number of jobs created is the number of unemployed workers times the matching rate for an interview f, the likelihood of continuation into negotiation $(1 - s_x)$, and the likelihood of a successful negotiation $\mathbb{E}[G_w(\tilde{w}(p))]$.

In order to determine the flow into unemployment, we define the stationary joint cumulative distribution of productivity levels and reservation wages across employed workers, $\Lambda(p, w)$. The mass of $\Lambda(p, w)$ less than the cutoff wage function $\tilde{w}(p)$, expressed as $\Lambda_w(\tilde{w}(p))$, represents the number of workers who will continue in employment next period. Then the job destruction flow of workers from employment to unemployment is given by:

$$\begin{aligned} \text{Jobs Destroyed} &= \left[(1 - s_x) \left(1 - \iint^{\tilde{w}(p)} \mathrm{d}\Lambda(p, w) \right) + s_x \right] (1 - u) \\ &= \Big((1 - s_x)\mathbb{E}[1 - \Lambda_w(\tilde{w}(p))] + s_x \Big) (1 - u)\,, \end{aligned} \tag{1.17}$$

where the stock of employed workers $1 - u$ separates exogenously at rate s_x and endogenously due to wage demands exceeding the firms' cutoff wage function.

The stationary unemployment rate that is consistent with these flows is therefore implicitly defined as the u that equalizes the number of jobs created (equation 1.16) with the number of jobs destroyed (equation 1.17) :

$$u^{\star} = \frac{(1-s_x)\mathbb{E}[1-\Lambda_w(\tilde{w}(p))]+s_x}{f(\theta^{\star})(1-s_x)\mathbb{E}[G_w(\tilde{w}(p))]+(1-s_x)\mathbb{E}[1-\Lambda_w(\tilde{w}(p))]+s_x} \tag{1.18}$$

where stationary labor market tightness $\theta^{\star}$ is defined as the ratio of the stationary vacancy level $v^{\star}$ to the stationary unemployment level $u^{\star}$.

Thus, a *recursive stationary equilibrium* of the model is a collection of value functions $\{V, J, W^F, W^R, U^F, U^R\}$, a collection of policy functions $\{\tilde{w}(p), w^{\star}_{EF}(p), w^{\star}_{ER}(p, w_{-1}), w^{\star}_{UF}(p), w^{\star}_{UR}(p, w_{-1})\}$, an unemployment level $u^{\star}$, and a vacancy level $v^{\star}$ such that:

- Firms maximize expected profits;
- Workers maximize their expected value functions taking firms' policies as given;
- Posting a vacancy has an expected value of zero; and
- Employment flows are consistent with firm and worker policy functions.

The appendix describes our procedure for solving the model.

1.5 Model Estimation and Results

In this section we estimate the parameters of the model described in section 1.4 and describe some implications of the results. We also examine the model's response to one-time permanent shocks to aggregate productivity. We compare the results of both the steady-state model and the simulations with aggregate shocks to the empirical results from the literature and in this paper, and argue that the model is able to reconcile the observed facts.

1.5.1 Target Moments and Estimation

The theoretical model has 11 parameters: $\beta, \pi, \phi, A, \psi_p, \sigma_p, b, c, \lambda_E, \lambda_U,$ and s_x. We set β to 5 percent annually, as in Shimer (2005) and Hall (2005). Our model does not explicitly feature real productivity growth, so we set π to 4 percent annually to capture both price inflation and productivity growth. Implicitly, β represents a discount factor that encompasses both pure time preference and trend growth in consumption.

We estimate eight of the nine remaining parameters, $\Theta = \{A, \psi_p, \sigma_p, b, c, \lambda_E, \lambda_U, s_x\}$, via indirect inference, in order to match a set of simulated moments, $\hat{\mu}^s(\Theta)$, to a set of real-world target moments μ. The elasticity of the matching function with respect to the unemployment rate, $1-\phi$, is set to equal the resulting share of job surplus accruing to the worker (Hosios 1990).[23] The estimated parameters are the values that minimize $\hat{\Theta} = \arg\min_{\Theta}[\hat{\mu}^s(\Theta) - \mu]' W^{-1} [\hat{\mu}^s(\Theta) - \mu]$, where the weighting function W is a diagonal matrix of the squares of the target moments.

Estimating the model via indirect inference helps to correct for potential sources of error in our empirical approach. The first is mis-classification of job stayers and job finders. The second is bias in our measurement of the fraction of counterfactual wage cuts prevented by downward nominal wage rigidity. Besides helping to correct for potential measurement error, the indirect inference procedure provides a tight link between the reduced form empirical estimates and their counterparts in the simulated data.

Table 1.5 displays the empirical moments that we target, along with the simulated moments that result from the estimated model. All moments were calculated over the years 1980 to 1997. The unemployment rate, $u = .061$, job-finding rate, $f = .42$, and median duration of unemployment, $D = 3.9$ months, are targeted to CPS quarterly averages.[24]

The remainder of the target moments are derived from the PSID data.[25] Wage rigidity for incumbent workers ($\widehat{wr}_s = 0.53$ for job stayers) and new hires ($\widehat{wr}_f = 0.36$ for job finders), outlined in section 3 above, are derived from the PSID data for 1980-1997. The difference between the 50th and 80th percentile of annual real log wage changes, $\Phi_{80} - \Phi_{50}$, is estimated from the same dataset to be about 11 percentage points. The moments $\hat{\alpha}$ and $\hat{\sigma}_{\ln w}$, are taken from the regression $\ln w_{it} = \delta_0 + \delta_t + \alpha \ln w_{it-1} + u_{it}$, and are estimated as $\hat{\alpha} = 0.88$ and $\hat{\sigma}_{\ln w} = 0.19$.

Although in principle all of the target moments can influence all of the estimated parameters, in practice some of the target moments have a larger influence on some parameters than on others. The Calvo parameters λ_E and λ_U are primarily determined by the wage rigidity target moments (λ_U is also influenced by D). The parameters governing the productivity process, ψ_p and σ_p, are heavily influenced by many of the targets, but ψ_p is directly characterized by $\hat{\alpha}$, while σ_p is influenced more by $\hat{\sigma}_{\ln w}$ and D. The matching efficiency parameter A and the exogenous separations rate s_x are jointly determined by u and f. The cost of vacancy creation c and the flow unemployment benefit b are characterized in part by the share of job surplus accruing to the worker.

The model matches the target moments u, f, $\hat{\alpha}$, $\Phi_{80} - \Phi_{50}$, and D reasonably closely. $\hat{\sigma}_{\ln w}$ was more difficult to match, which is perhaps unsurprising given the model's lack of heterogeneity in job amenities, investment in human capital, or match quality between jobs and workers. The estimated

[23] This condition ensures that the number of vacancies is efficient in an environment with flexible wages.

[24] The first two are from data constructed by Robert Shimer using the CPS. For more details, see Shimer (2012).

[25] We have also estimated this model using CPS-derived moments, with qualitatively similar results, but we highlight the PSID data because it is a superior dataset for multi-year analysis of the evolution of wages (as opposed to a single one-year change for each worker in the CPS).

model matches the proportion of wage cuts prevented by wage rigidity for job finders $\widehat{wr}_f$ quite closely, but produces an estimate for job stayers $\widehat{wr}_s$ that is lower than in the data.

Table 1.6 lists the estimated parameter values and their standard errors. The values for λ_E and λ_U could not be determined to be different from each other at the 95th percent significance level, and they correspond to a probability of experiencing rigid wages of greater than 90 percent per month. The exogenous separations rate s_x corresponds to a little more than half of the total separations implied by the stationary unemployment and finding rates. The elasticity of the matching function with respect to unemployed workers $1 - \phi$ is 0.73, reflecting the "take-it-or-leave-it" nature of the wage demands of workers in our model, which causes the majority of the match surplus to accrue to the workers. Nonetheless, because workers realize that they may be unable to cut their wage demands in the future, they moderate their wage demands in the present, leaving a non-trivial share of the match surplus to the firm.[26] The flow benefit of unemployment, b, is estimated to be quite low at 0.29, or approximately 30% of the average wage, which helps to moderate their wage demands.

We estimate the following equation on our simulated data using individual-specific productivity levels and wages:

$$\Delta \ln w_{ijt} = \beta_{0,j} + \beta_{1,j} \Delta \ln y_{ijt} + \varepsilon_{ijt}, \tag{1.19}$$

We divide workers into seven groups j: 1) all employed workers, 2) incumbent workers who were employed in the previous month, 3) incumbent workers who have had continuous employment for at least a year, 4) newly hired workers, 5) workers who were hired in the previous 3 months, 6) unemployed workers, and 7) unemployed workers plus new hires. For each individual in each group we record the most recent observed wage and productivity data in the event that they were not employed in the prior period (e.g. newly hired workers). We use reservation wages in place of actual wages for unemployed workers.

Table 1.7 shows the results of these regressions. Even with the high level of wage rigidity for job finders estimated in our model, the regressions show that their wages are very responsive to their productivity levels. The estimated elasticity of wages with respect to individual productivity is 0.86 for those who have been working for 3 months or less and 1.06 for new hires. Job stayers exhibit a much lower wage elasticity of 0.36. The reservation wages of unemployed workers exhibit a wage elasticity with respect to productivity of 0.33, much lower than for job finders, consistent with the survey data of Krueger and Mueller (2014). When we estimate the regression combining unemployed

[26] This result is reminiscent of Elsby (2009), who argues that firms will respond to downward nominal wage rigidity by compressing wage increases in good times.

workers' reservation wages with the observed wages of new hires, we estimate an elasticity of 0.42, similar to the elasticity for all employed workers. We interpret these results as supporting our contention that composition bias drives the apparent flexibility of the wages of new hires.

1.5.2 Dynamic Simulation

We use our estimated model to simulate a dynamic economy in which the aggregate productivity level changes. We run 1,000 simulations of the economy containing 10,000 workers over 14 years (168 months), discarding the first 10, and implement a permanent unanticipated productivity shock (with equal probability of a 1 percent increase or a 1 percent decrease) in the first month of year 12. The resulting simulated data represents 1,000 unique simulations of 12 months of base data and 36 months of data responding to the new permanent change in aggregate productivity. It is common knowledge among all agents in the model that the shock will be permanent.

To facilitate computation, we use policy functions of firms and workers that correspond to the eventual new steady state of the model immediately after the aggregate productivity shock hits. Although the model is not fully dynamic in that sense, we argue that it approximates the transition from one steady-state to another reasonably well for two reasons. First, following the logic of Shimer (2005) and Pissarides (2009), because the hiring and job separation probabilities are quite large in the model, the model will converge quickly to the new equilibrium after the aggregate shock.[27] Additionally, employment relationships are on average long-lasting (46 months in the steady-state model). Thus, firms' hiring and workers' wage-posting decisions in the new steady state are likely to be a close proxy for their behavior in transition between states. We allow free entry of vacancies to adjust to period-specific labor market flows.

Figures 1.3 and 1.4 show the impulse-response functions of unemployment, job finding, separations, wage demands, and productivity after positive and negative shocks to aggregate productivity. The unemployment rate behaves asymmetrically with regards to the productivity shocks, as a positive productivity shock leads to a smaller change in unemployment and a larger initial change in wages than a negative productivity shock causes. This asymmetry results from wages being downwardly rigid but upwardly flexible in our model.

Figure 1.3 also decomposes unemployment into its various sources in the model. Exogenous separations and rigidity-based separations are the proximate causes of a little less than half the total unemployment in any given period, while rigid wages for unemployed workers and failures to

[27] In the baseline steady state, the job finding probability is 0.395 and the separation probability is .022, implying a half-life for the deviation of unemployment from its steady state value of approximately 1.66 months.

match with a vacancy account for the remainder. In the event of a negative shock, unemployment due to rigid wages for both the employed and the unemployed increases significantly, but after the initial period (during which many workers flow from the employed to unemployed states) the effect of rigidity for unemployed workers dominates the effect of rigidity for employed ones. Therefore, the finding rate has a slower adjustment process after a negative shock than the separation rate and contributes more overall to the volatility of unemployment, as seen in the bottom panels of figure 1.4.

The observed labor productivity of job finders also responds asymmetrically to positive and negative productivity shocks. Because idiosyncratic productivity evolves as a Markov process independent of firm and employee behavior, these changes are driven entirely by compositional changes. The new hiring spurred by a positive productivity shock brings more marginal employees into employment, leading to a gradual rise in average observed productivity. The wage rigidity for entrants binds much more strongly for hires in the event of a negative shock than for a positive one, leading to the composition of new hires with flexible wages rising dramatically in recessions. Thus, a negative shock leads to a rapid decline in average productivity among finders, from which it gradually recovers.

This pattern is also evident in the wage demands of job finders versus all workers, shown in the second row of figure 1.4. Real wages track the path of productivity closely, illustrating the importance of wage flexibility in achieving full employment in an economic downturn. In the model, unemployed workers, from whom the pool of new hires is selected, do not exhibit this kind of flexibility on average. The ratio of the reservation wage to the idiosyncratic productivity level of the unemployed, shown in the third row of figure 1.4, does not respond any more dramatically than the wage-to-productivity ratio of employed workers. The gap between these ratios constitutes a major barrier to unemployed workers' chances of becoming re-employed.

We use the dynamic responses of the model to the aggregate productivity shocks to estimate aggregate elasticities in the spirit of the empirical work in section 1.2. To measure the responsiveness of aggregate wages in various subgroups of workers (all workers, new hires, and incumbent workers), we randomly select a "survey month" from the quarter before the productivity shock and measure the log change in average wages three months later. The upper panel of table 1.8 displays the results of regressing these log changes in average wages on equivalent economy-wide changes in average labor productivity or the unemployment rate. The elasticity of average wages with respect to productivity is 0.68 for new hires (defined as workers who have found a job in the past three months), significantly higher than it is for the working population as a whole, 0.30. Both of these elasticities compare closely with the elasticities found in Haefke et al. (2013). A similar pattern holds for the elasticities of average wages with respect to the unemployment rate, as all workers exhibit an elasticity of -0.16, versus -0.38 for job finders.

The bottom panel of table 1.8 demonstrates the effect of aggregate productivity and the unemployment rate on individual wage changes by worker category. In this analysis a job finder is defined as in the PSID data to be a worker who was unemployed or out of the labor force at any point between surveys (typically 12 months). Job finders appear to be more responsive to changes in labor market conditions, as indicated by the greater responsiveness of their wages to labor market conditions. Including unemployed workers with job finders, as in the rightmost column, indicates that part of that finding stems from composition bias. Workers who remain unemployed actually exhibit more rigid wages, as measured by less elastic wage demands with respect to labor market conditions, than workers who have maintained a job over the course of a year. This result again stems from the negative selection with regard to wage flexibility associated with unemployed workers in the model: unemployed workers are much more likely to have experienced rigid wages than the working population.

1.6 Conclusion

This paper has demonstrated that composition bias potentially accounts for the apparent flexibility of the wages of newly hired workers. Newly hired workers are disproportionately likely to have flexible wages relative to the pool of unemployed workers. Empirical analyses that omit the reservation wages of the unemployed are prone to over-estimate the responsiveness of potential new hires to macroeconomic conditions. Of course, because reservation wages are not regularly measured in most economic datasets, this problem is inherently difficult to solve.

Using a model that explicitly ascribes high degrees of nominal wage rigidity to both incumbent workers and new hires, we are able to reconcile the dissonant statistics in the empirical data. Furthermore, when we re-estimate the key empirical elasticities including the reservation wages of unemployed workers along with the observed wages of new hires, the combined wages are noticeably less responsive to labor market conditions.

Dynamic simulations point to the same conclusion. The simulated elasticities of wages with respect to productivity closely match those in the U.S. data. A negative productivity shock leads to a persistent increase in unemployment attributable to wage rigidity among the unemployed, but the aggregate wages of new hires appear flexible because of selection effects. Therefore, we argue that nominal wage rigidity may account for a substantial share of the lower job finding rates during recessions.

1.7 Appendix

1.7.1 Technical Appendix on Data

This section of the appendix details the methods employed in Sections 2 and 3 to analyze nominal wage rigidity and several real wage elasticities in PSID and CPS data.

1.7.1.1 PSID

We extracted employment, wage, and demographic information from each of the 1980-2013 family-level surveys for both the head of household and, where available, spouse. Because the PSID uses primarily gender-based assignment of "head of household," the primary earner of each family might be either the head or the spouse according to that family's particular economic situation.

These data were used to create an individual dataset, whereby each family-level entry was merged onto the individual file to obtain an individual weight corresponding to each record. The resulting file contained an average of 12,513 records per year, of which 69.1% were in the labor force and 5.9% within that category were unemployed.

The hourly nominal wage for each respondent was calculated using their primary job only, as this best suited our understanding of where wage rigidity manifests and is most economically relevant. For hourly employees, we used their current reported hourly wage; for salaried workers, we used their hourly wage when they reported it, but their more common response was to report pay over a longer horizon such as "per month" or "per year" values. In these cases, we assumed 52 working weeks per year, 40 hours per week to assign an hourly wage. We excluded all top-coded data, which applies to hourly wages at or above 100 dollars per hour before 1993 and 1,000 dollars per hour from 1993 to 2013. Salaried employees are top-coded above 1 million dollars from 1980 to 1993 and 10 million dollars thereafter. We exclude from our analysis workers who earn significant money from bonus or incentive-laden schemes.

In alternate specifications we use the PSID-generated hourly imputed wage, which adds the earnings of the surveyed year (not the wages at the time of the survey) and divides by the imputed hours spent working. This measure is not our preferred specification because it potentially involves a host of relationships between each worker and his or her various employers, but as expected the level of wage rigidity decreased modestly using the imputed wage relative to our preferred hourly wage measure.

We calculate a tenure measure for all respondents who are currently employed. This measure is used both to categorize workers as "job stayers" versus "job switchers" and as a regressor to control for job-specific productivity gains in our elasticity measures. The tenure variable is constructed from a series of questions that ask how long the respondent has been with their current employer, with the answers converted into years. Job stayers are then distinguished from switchers by the time between surveys for each respondent; if the time elapsed between surveys is longer than the reported tenure, the respondent is deemed a switcher (unless they have experienced a month or more of non-employment). Otherwise, they are deemed a stayer. We clean the tenure variable according to the procedure in Altonji and Williams (1997) by allowing tenure to increment by only one year at a time for each respondent and re-setting it whenever the raw tenure measure indicates that the respondent is either a switcher or a finder.

It is important to note that our definition of a job finder excludes some respondents who might nonetheless be best described as such. In both our direct measure of wage rigidity and our estimates of the various wage elasticities for job finders, we require that finders have a wage in both the previous and current survey dates to be included in the sample. A respondent who was not working in the prior year at the time of their survey date, but who nonetheless found a job in time for the current year's survey date, would be excluded from the analysis. Therefore, the job finder sample skews towards those who had shorter spells of unemployment. In addition, the worker who experiences multiple rounds of unemployment will find their current wages compared not against their most recent job, but instead against the job held at the time of the previous survey.

We also attempted two other methods to categorize workers. In the first, we augmented the tenure versus time elapsed between surveys comparison with the additional restriction that occupational codes and industry codes could not switch from one year to the next. In the second, we directly compared start dates with previous survey dates. We prefer the original method because a) responses to length of time on the job were more frequently given (or perhaps known) than start dates, and b) occupational codes and industry codes were not always comparable between years.

1.7.1.2 CPS

We use the Integrated Public Use Microdata Series from 1989-2014. We restrict our analysis to nonfarm, nonsupervisory workers between the ages of 25 and 60, so we exclude those whose occupational codes indicated they were either in "Executive, Administrative, and Managerial Occupations" or "Management Related Occupations." In addition, we exclude "Agriculture, Forestry, and Fisheries" occupational codes and evaluate only workers who were either paid by a flat salary or hourly wage.

The top-coding of hourly and salaried workers is more severe in the CPS than in the PSID, so we impute earnings for top-coded workers with a Tobit distribution using year-specific Mincer regressions that include a host of demographic attributes, including years of schooling, a quartic in potential experience (age minus years of schooling minus 6), and race, gender, and married indicator variables. Hourly employees are top-coded at 100 dollars per hour. For salaried workers, weekly earnings are top-coded at 1,923 dollars per hour from 1989-2007 and 2,885 dollars thereafter.

We construct a measurement of hours for salaried workers that is equal to their typical hours worked per week when available, substituted for hours worked in the past week if it is not. We then trim the sample of outliers (totaling 1 percent of workers) to address concerns that the previous week may have been atypical. To create an hourly wage for salaried workers, we divide weekly earnings by weekly hours worked. Employees who were paid an hourly wage have it reported as such. Real hourly wages are trimmed symmetrically at the 0.5 and 99.5 percentiles as well.

We exclude intervals that span 1995q3 and 1995q4 from our analysis because a change in sample design renders us unable to match workers across that break.

We are unable to create an analog to job stayers and job switchers in CPS data, as we have an incomplete employment history of each respondent. The limited employment history allows us to categorize a worker as a job finder, but this definition is more restrictive than the one used in the PSID. A worker who has a job in the outgoing rotation group but reports a month of non-employment at any point in the three months prior is called a job finder in analysis performed with CPS data.

As in the PSID, wage changes are the unit of analysis. As such, workers who were unemployed in the outgoing rotation group either the first or second time do not have measurable wage changes and are excluded from the analysis.

We weight each record in our sample according to the earnings weight variable EARNWT.

1.7.2 Measuring Wage Rigidity

This paper measures the fraction of counterfactual nominal wage cuts prevented by downward nominal wage rigidity using the approach in Ehrlich and Montes (2015), which builds on the approach of Card and Hyslop (1997). We present a brief overview here.

For each year t, estimate the distribution of observed wage changes using kernel density estimation.[28] The estimate of the density at a point x is

$$\hat{f}_t(x) = \frac{1}{n}\sum_{j=1}^{n}\frac{1}{h_j}K\left(\frac{x-x_j}{h_j}\right) \tag{1.20}$$

where n is the number of observations, x_j for $j \in \{1, ..., n\}$ denotes a point in the observed distribution, h_j is an adaptive bandwidth following the procedure of Van Kerm (2003), and K is a kernel function.[29] The specific kernel function used in the estimation is an Epanechnikov kernel of the form

$$K(z) = \begin{cases} \frac{3}{4}\left(1-z^2\right) & \text{if } |z| < 1 \\ 0 & \text{otherwise.} \end{cases} \tag{1.21}$$

Denote the estimated distribution of observed wage changes as $\hat{f}_t^{obs}$, and let m_t represent the median wage change from year $t-1$ to year t expressed in percentage points.

Next, construct a counterfactual wage change distribution $\hat{f}^{cf}$ for establishment i by averaging the upper tails of the estimated observed distributions $\hat{f}_t^{obs}$ across each year. In constructing the average, first normalize the observed distribution for each year around its median.[30] Then, reflect the averaged distribution of the upper tails around the median each year.

The estimated proportion of wage cuts prevented by wage rigidity is then calculated by comparing the implied proportion of counterfactual wage cuts with the number observed. For year t, denote the proportion of wage cuts in the estimated observed wage change distribution as $\hat{F}_t^{obs}(0^-)$.[31] Denote the proportion of wage cuts in the estimated counterfactual distribution as $\hat{F}_t^{cf}(0^-)$. Let the sum across years of these proportions be denoted $\hat{F}^{obs}(0^-)$ and $\hat{F}^{cf}(0^-)$. The measure of wage rigidity is then the proportion of counterfactual wage cuts that are "missing" from the data and is calculated as

$$\widehat{wr} = 1 - \frac{\hat{F}^{obs}(0^-)}{\hat{F}^{cf}(0^-)}. \tag{1.22}$$

[28] The estimation procedure focuses on wage changes within 15 percentage points of the median wage change each year to avoid the influence of outliers.

[29] The global bandwidth is set to be 0.005. The adaptive bandwidths are calculated as the product of the global bandwidth and a local bandwidth factor that is proportional to the square root of the underlying density function at the sample points. The adaptive bandwidths have the property that their geometric average equals the global bandwidth.

[30] In practice, in situations in which the observed median is negative and there are more observed wage cuts than wage increases, recalculating the median by excluding observed wage changes between -0.25% and 0.25% helps to correct for the "sweep-up" of counterfactual wage cuts to zero. This adjustment improves the accuracy of the procedure in the Monte Carlo simulations discussed in Ehrlich and Montes (2015). Those years are then excluded when averaging the upper tails, but are included when calculating the counterfactual wage cuts prevented by wage rigidity.

[31] The notation 0^- indicates that the measured proportion does not include wage changes of exactly zero.

Therefore, the wage rigidity estimate in equation (1.22) is time-invariant. $\widehat{wr}$ has the natural interpretation that a value of 0.25 implies that 25 percent of counterfactual nominal wage cuts were prevented by downward nominal wage rigidity over the sample period.[32]

1.7.3 Computational Methods

We approximate the value functions for firms and workers using standard value function iteration techniques. We approximate the productivity process using the method of Tauchen (1986), using a productivity grid with 200 nodes. The nodes are spaced evenly in log terms, ranging from two standard errors below the mean to two standard errors above. At the estimated values for persistence of the productivity process ψ_p and innovation σ_p, the minimum grid value for p is 0.7469 and the maximum value is 1.3387. We allow workers to choose reservation wages along an evenly-spaced 250-point grid with a minimum value of 0.7095 and a maximum value of 1.4057; these extrema represent a range that encompasses that of p and extends an additional 5 percent in either direction.

The period for the model simulations is taken to be one month. We draw one set of random shocks to use in every simulation. We simulate 2500 workers for 60 years, or 720 periods, discarding the first 10 years (120 periods) for burn-in. We sample workers' simulated wages annually, except where noted otherwise in the text, for the purpose of measuring individual and aggregate elasticities and wage rigidities.

In order to reduce the numerical error associated with the calculation of the standard errors in table 6, we calculate the derivatives of the simulated moments $\hat{\mu}^s(\Theta)$ with respect to the model parameters Θ using three different step sizes, 1%, 3%, and 5%, and take the average derivatives.

[32]Nothing in this procedure prevents $\widehat{wr}_i$ from being negative. A value for $\widehat{wr}_i$ of -0.25 would imply that there are 25 percent more wage cuts in the data than would be predicted by the distribution of nominally positive wage changes.

Table 1.1: Descriptive Statistics
PSID, 1980 to 2013

Variable	All Respondents	Salaried Workers	Hourly Workers	Other Workers	Not Employed
Average Respondents per Survey Year	12,513	2,828	3,974	603	4,604
Average Age	43.0	39.8	37.7	41.6	50.0
Proportion Male	0.50	0.54	0.50	0.50	0.31
Average Job Tenure	6.7	7.3	6.3		
Proportion of Job Stayers	0.64	0.69	0.60		
Hourly Wage and Salary Income ($/hr.)	14.0	18.6	11.0		
Labor Force Participation Rate	0.69				
Unemployment Rate	0.059				
Average Unemployment Duration (Weeks)					29.1

Notes: Job tenure, proportion of job stayers, and hourly wage and salary income under the category "All Respondents" include salaried and hourly workers. Unemployment duration average among not employed who are in the labor force.

Table 1.2: Elasticity of Real Wages with Respect to Unemployment in the PSID

All Workers	-0.83	-0.50	-0.70	-1.07
	(0.46)	(0.48)	(0.19)	(0.58)
All Job Stayers	-0.55	-0.23	-0.49	-0.76
	(0.53)	(0.50)	(0.13)	(0.49)
All Job Switchers	-1.80	-1.70	-0.59	-3.35
	(1.16)	(0.93)	(0.70)	(0.96)
All Job Finders	-1.82	-1.22	-1.65	-1.91
	(0.36)	(0.62)	(0.79)	(1.32)
Wage Measure	current job	current job	average wage	average wage
Sample Years	1980-1997	1980-2013	1980-1997	1980-2013

Notes: standard errors are clustered by the year level. We use two measures of wages: 1) hourly wage in primary job at time of survey, and 2) average hourly wage in surveyed year across all jobs. The job tenure variable is adjusted using the procedure in Altonji and Williams (1997). We define finders as workers who are employed at the time of the survey date, but who responded that they had experienced unemployment at some point during the surveyed year. Stayers are defined similarly as continuously employed workers with the same employer between surveys. Switchers are defined as continuously employed workers who are not with the same employer as the prior survey. All includes stayers, finders, and workers who switched jobs between surveys. Analysis includes only workers who are primarily hourly employees or salaried employees.

Table 1.3: Elasticity of Real Wages with Respect to Productivity in the CPS

	Group Wage Changes	Individual Wage Changes	
All Workers	0.18	0.33	0.33
	(0.15)	(0.20)	(0.20)
All Workers Less Finders		0.32	0.31
		(0.20)	(0.20)
All Job Finders	0.94	0.59	0.75
	(0.40)	(0.63)	(0.46)
Sample Years	1989-2006	1989-2006	1989-2013

Notes: Group wage changes are imputed as in Haefke, Sontag, and van Rens (2014). Standard errors are clustered by the year level for individual wage changes. We define finders as workers who are employed at the time of the survey date, but who responded that they had experienced unemployment at some point during the prior 3 months. Analysis includes only nonfarm, non-management workers between the ages 25 and 60 who are primarily hourly employees or salaried employees.

Table 1.4: Measured Wage Rigidity in the PSID and CPS

Category of Worker	PSID, 1980-2013 Job Stayers	PSID, 1980-2013 Job Switchers	PSID, 1980-2013 Job Finders	CPS, 1989-2013 All Workers	CPS, 1989-2013 All Non-Finders	CPS, 1989-2013 Job Finders
All Workers (Salaried & Hourly)	52.7%	56.1%	36.4%	47.4%	47.5%	38.6%
(Std Dev)	(0.4%)	(2.0%)	(1.2%)	(0.3%)	(0.3%)	(1.5%)
Salaried Workers	48.7%	62.7%	40.0%	41.0%	40.8%	43.9%
(Std Dev)	(0.7%)	(3.7%)	(4.8%)	(0.7%)	(0.7%)	(6.8%)
Hourly Workers	55.1%	50.2%	36.6%	51.9%	52.2%	41.6%
(Std Dev)	(0.5%)	(1.8%)	(1.3%)	(0.4%)	(0.4%)	(1.6%)

Notes: Estimates of wage rigidity measures the proportion of counterfactual wage cuts missing from the observed wage change distribution, measured as described in Appendix A. Standard errors from 100 bootstrap replications in parentheses.

Table 1.5: Empirical and Simulated Moments

Target Moments	Description/Source	Target Values	Model Values
u	Average Monthly Unemployment Rate, CPS 1980-2007	0.061	0.056
f	Average Monthly Job Finding Hazard Rate, CPS 1980-2007	0.432	0.395
$\Phi(.8) - \Phi(.5)$	80th Percentile-50th Percentile Real Log Wage Changes, PSID 1980-2013	0.112	0.094
$\hat{wr}_f$	Wage Rigidity for Finders, PSID 1980-1997	0.364	0.394
$\hat{wr}_s$	Wage Rigidity for Stayers, PSID 1980-1997	0.527	0.435
$\hat{\beta}_{\Delta \ln w}$	AR(1) Coefficient on Log Wages, PSID 1980-1997	0.881	0.915
$\sigma_{\ln w}$	Std. Dev. Of AR(1) Log Wage Innovations, PSID 1980-1997	0.188	0.045
D	Mean Duration (months) unemployed, CPS 1980-2007	3.900	3.861

Table 1.6: Model Parameters

Parameter	Description	Values (annual)	Values (monthly)	Standard Error
β	Time Preference	0.950	0.996	–
π	Trend Nominal Wage Growth	0.040	0.003	–
$\bar{p}$	Average Productivity Level (Normalized to 1)	1.000	1.000	–
$1-\phi$	Elasticity of Matching Function w.r.t. unemployment	0.730	–	–
A	Efficiency of Matching Function	–	0.768	0.111
s_x	Exogenous Separations Rate	–	0.014	0.002
ψ_p	Persistence of Productivity Process	0.925	0.994	0.002
σ_p	Standard Deviation of Productivity Shock	–	0.017	0.001
b	Flow Benefit of Unemployment	–	0.290	0.410
c	Flow Cost of Vacancy Posting	–	0.158	0.216
λ_U	Probability of Rigid Wages - Unemployed Worker	0.314	0.908	0.069
λ_E	Probability of Rigid Wages - Employed Worker	0.414	0.929	0.291

Note: The standard errors for ψ_p, λ_U, and λ_E are reported for the annualized values.

Table 1.7: Elasticity of Real Wages with Respect to Individual Productivity in Simulated Data

Worker Subgroup	Elasticity
All Employed	0.424
	(0.001)
Job Stayers	0.361
	(0.001)
Job Finders	0.860
	(0.003)
Unemployed Workers	0.331
	(0.002)
New Hires and Unemployed Workers	0.423
	(0.002)

Notes: A stayer is someone who was either employed in the previous month (Stayers - 1 month) or continuously employed during the past year (Stayers - 12 months). A finder is anyone who is currently employed but was unemployed in either the previous month (Finders - 1 month) or previous 3 month (Finders - 3 months). Unemployed wages reflect wage demands rather than observed wages.

Table 1.8: Elasticity of Real Wages with Respect to Productivity, Unemployment (Simulated Data)

Panel A: Aggregate Wages	All	Stayers	Finders			
elasticity of wage wrt productivity	0.30	0.28	0.68			
standard error	(0.00)	(0.00)	(0.03)			
elasticity of wage wrt unemployment	-0.16	-0.16	-0.38			
standard error	(0.00)	(0.00)	(0.02)			
					Including Unemployed	
Panel B: Individual Wages	All	Stayers	Finders	Unemployed	All	Finders
elasticity of wage wrt annual productivity	0.62	0.55	0.77	0.49	0.63	0.66
standard error	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)
elasticity of wage wrt annual unemployment	-0.45	-0.40	-0.56	-0.35	-0.46	-0.48
standard error	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Notes: a stayer is someone who was continuously employed during the past year. A finder is anyone who is currently employed but was unemployed at some previous point. For the aggregate wages, a Finder was unemployed in the previous 3 months, as in Haefke, Sonntag, and ven Rens. For individual wages, a finder was unemployed at any point in the previous year.

Figure 1.1: Nominal Wage Changes in the PSID by Worker Category

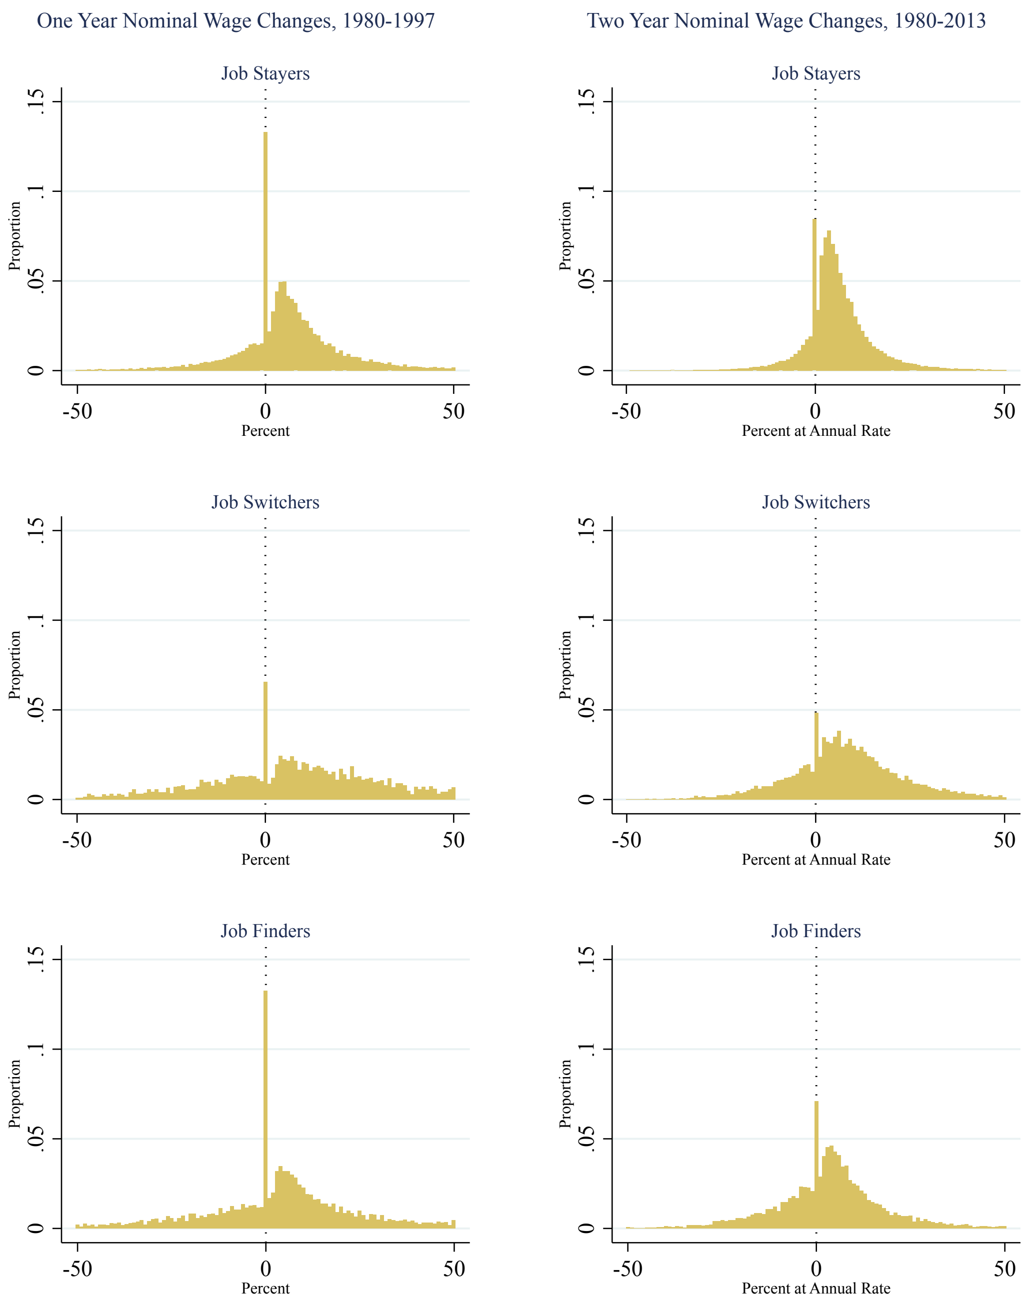

Figure 1.2: Nominal Wage Changes in the CPS by Worker Category

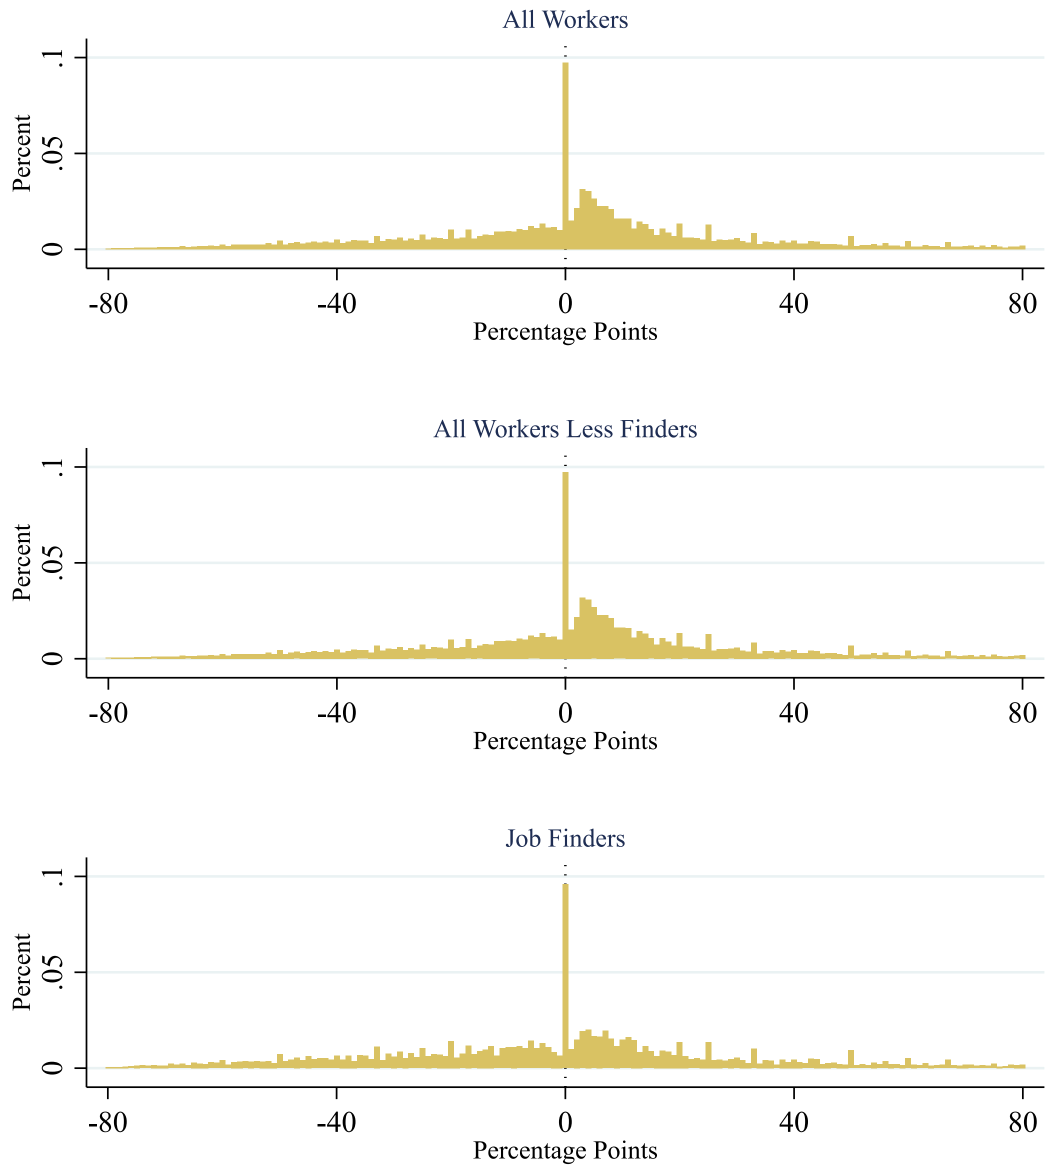

Notes:
Data range from 1989 to 2013.
Years 1995-1996 are excluded because of a sample design change in 1995 that hinders matching.
Graphs are truncated at -80 and 80 percent.

Figure 1.3: Sources of Unemployment

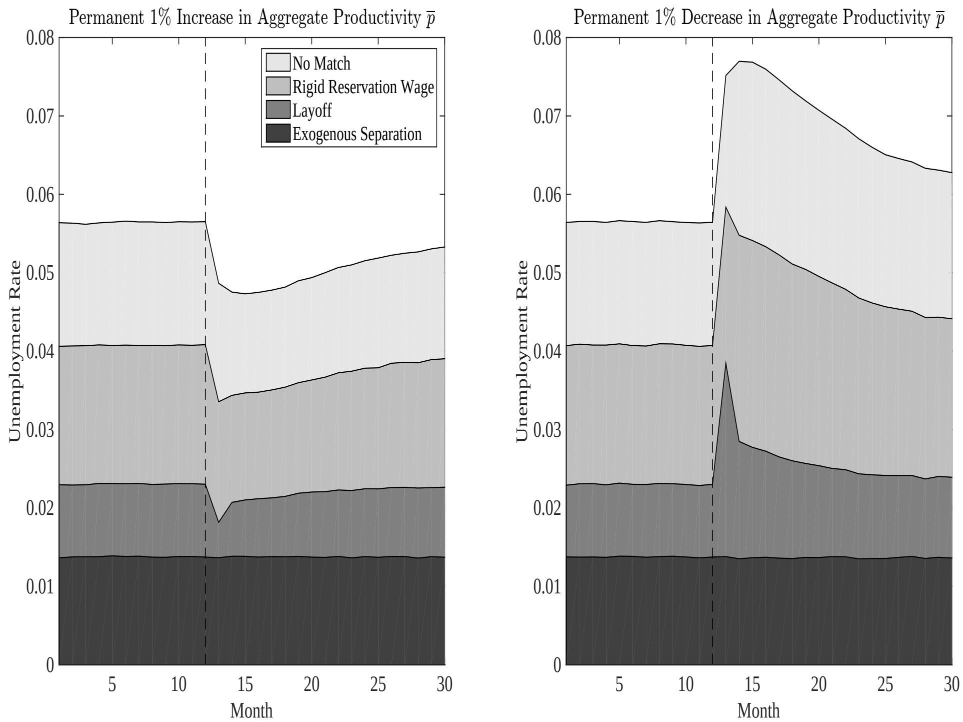

Figure 1.4: Labor Market Responses to Aggregate Productivity Shocks

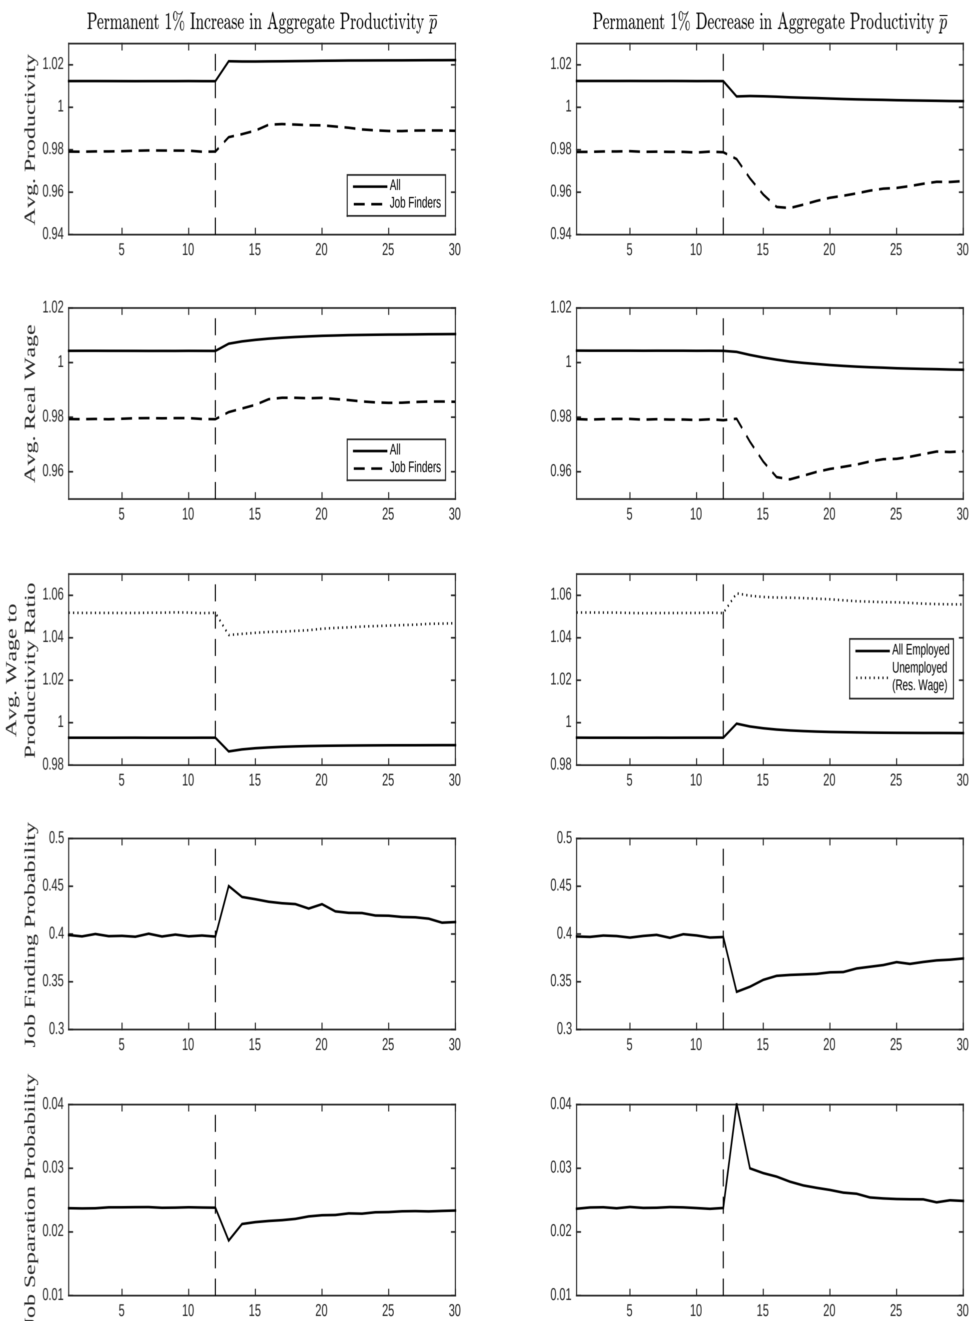

Figure 1.5: Nominal Wage Growth Among Job Stayers in the PSID by Year

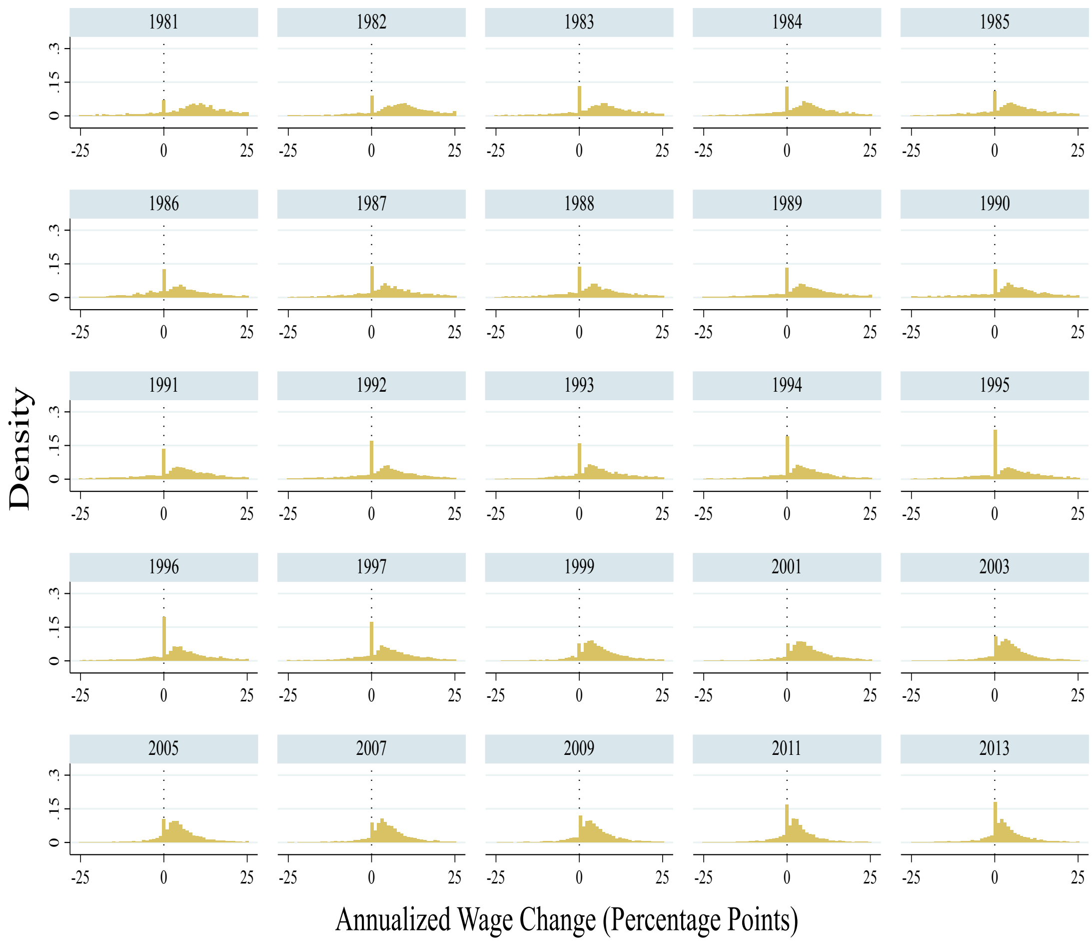

Notes:
Years 1981-1997 show the annual wage change from the previous year.
Years 1999-2013 show the annualized two-year wage change from two years previously.
Graphs are truncated at -25 and 25 percent.

Figure 1.6: Nominal Wage Growth Among Job Switchers in the PSID by Year

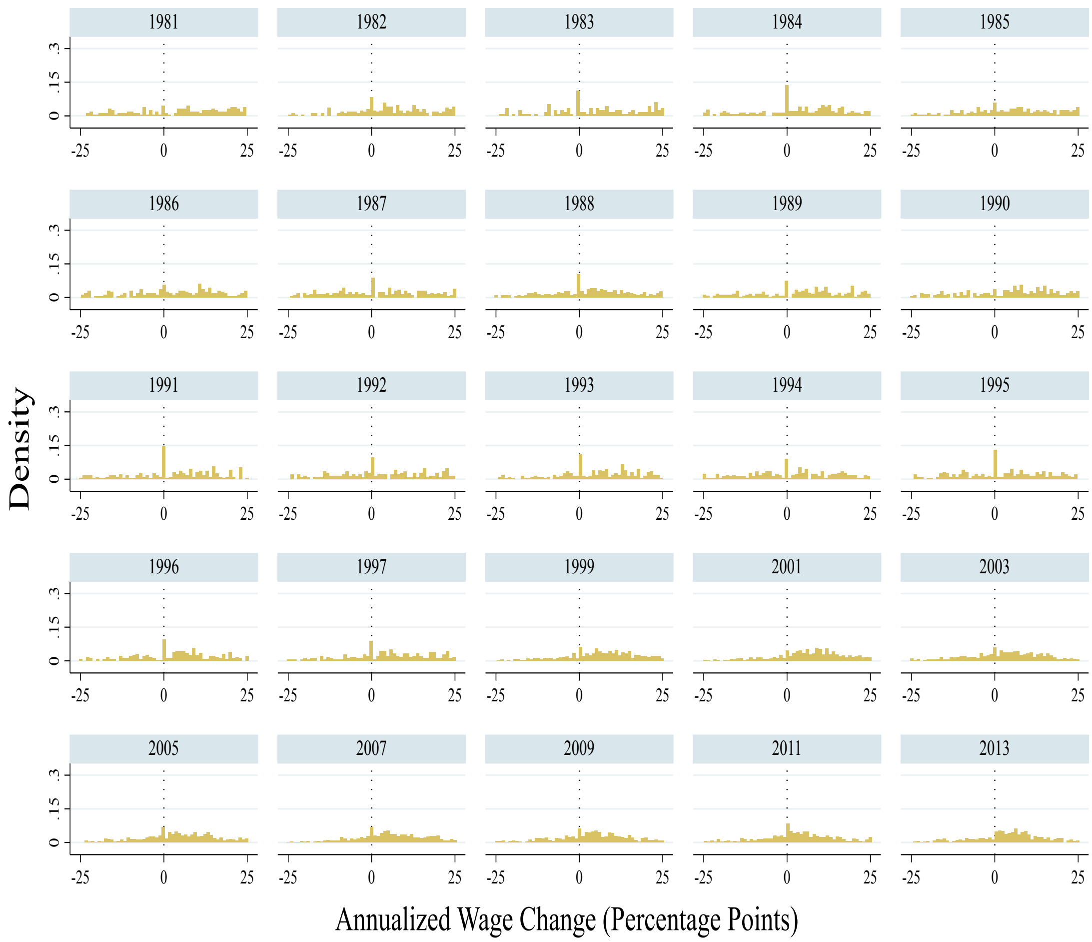

Notes:
Years 1981-1997 show the annual wage change from the previous year.
Years 1999-2013 show the annualized two-year wage change from two years previously.
Graphs are truncated at -25 and 25 percent.

Figure 1.7: Nominal Wage Growth Among Job Finders in the PSID by Year

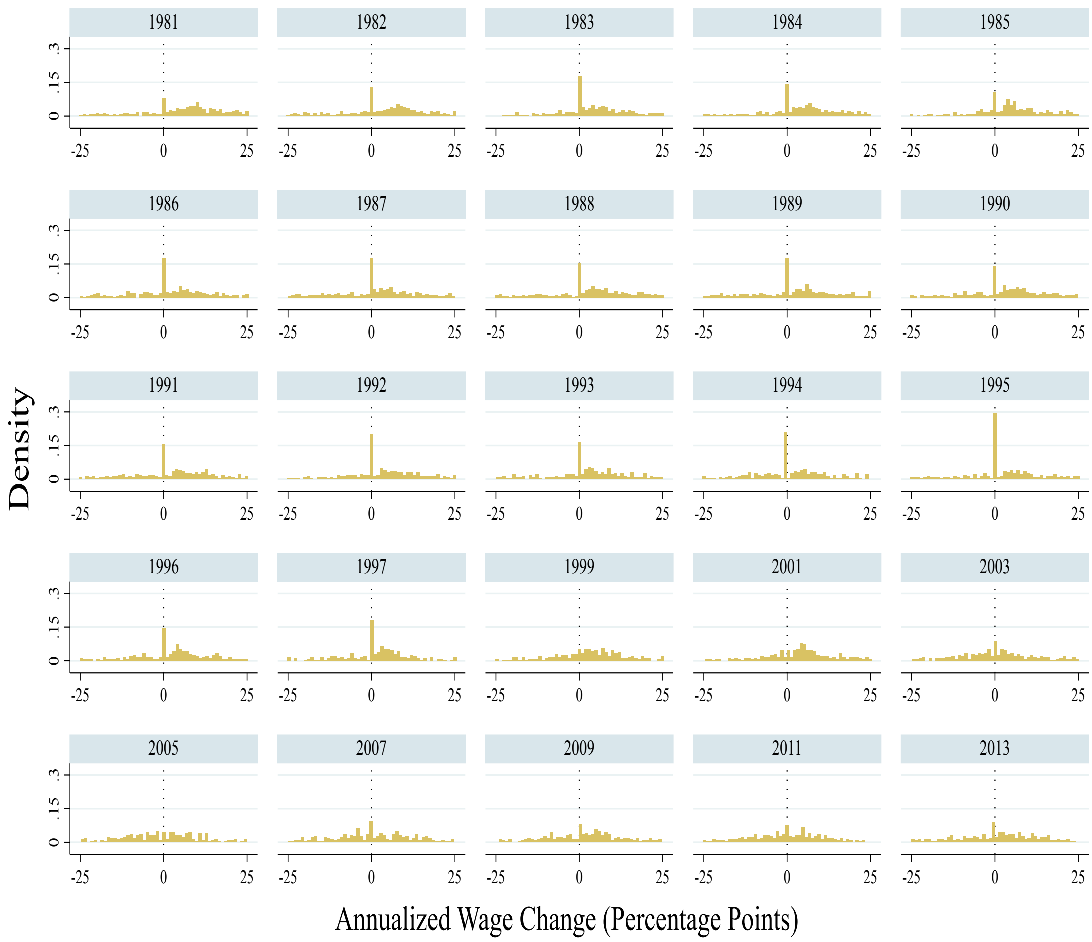

Notes:
Years 1981-1997 show the annual wage change from the previous year.
Years 1999-2013 show the annualized two-year wage change from two years previously.
Graphs are truncated at -25 and 25 percent.

Figure 1.8: Nominal Wage Growth Among All Workers in the CPS by Year

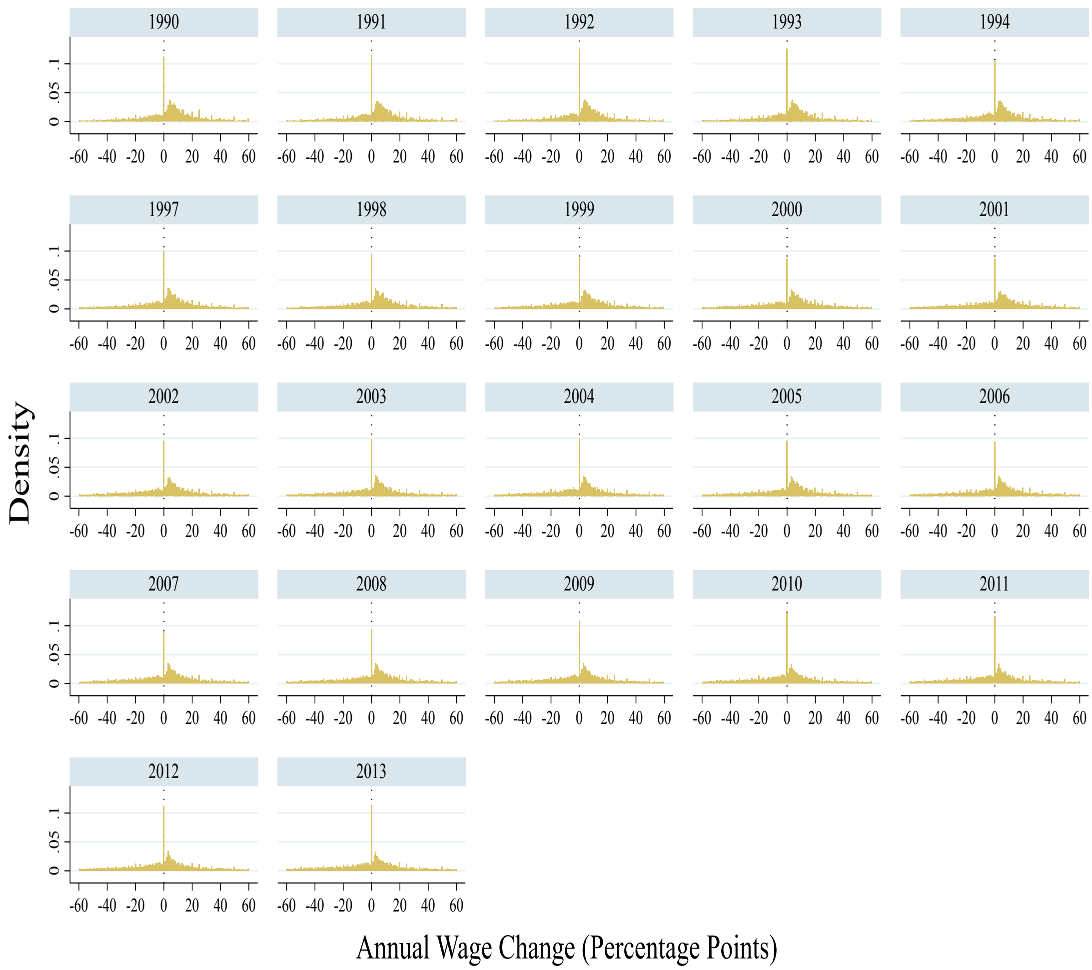

Notes:
Years 1995-1996 are exluded because of a sample design change in 1995 that hinders matching.
Graphs are truncated at -60 and 60 percent.

Figure 1.9: Nominal Wage Growth Among Workers, Excluding Finders, in the CPS by Year

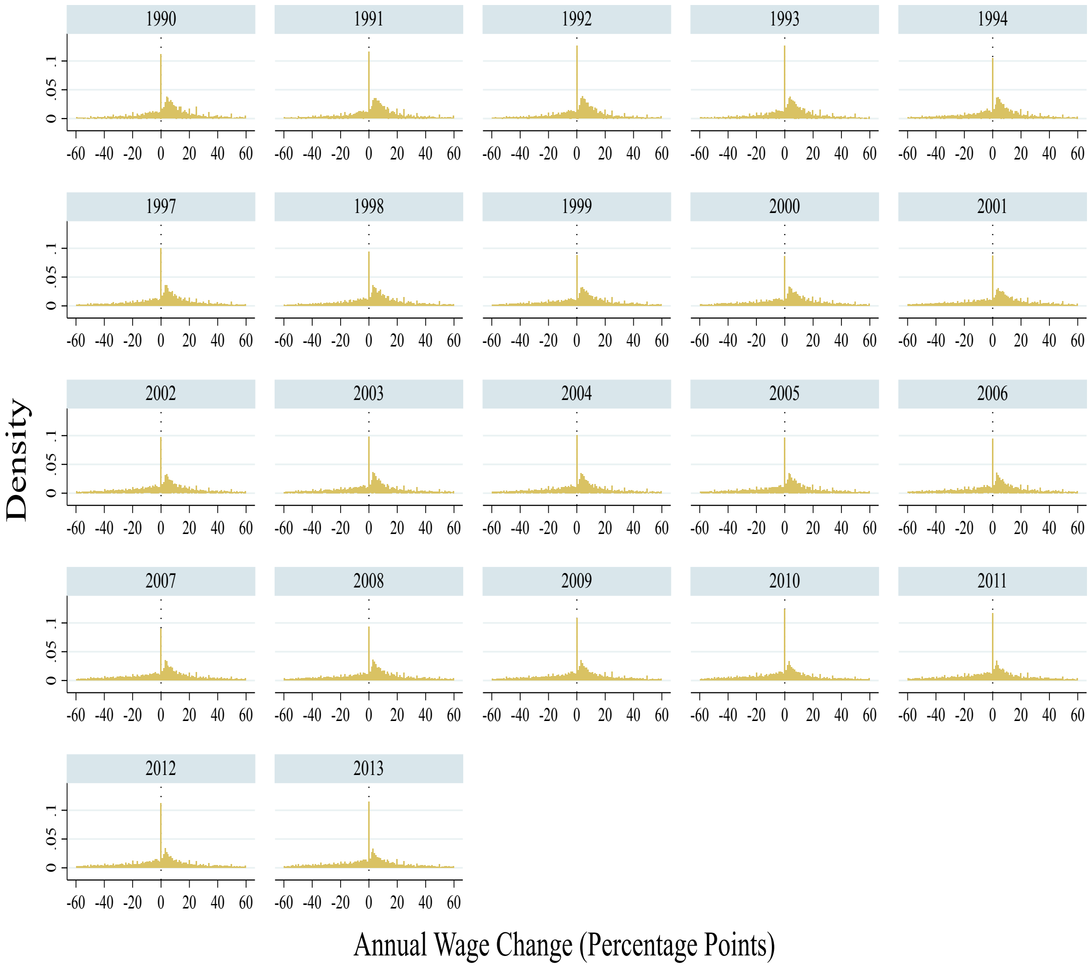

Notes:
Years 1995-1996 are exluded because of a sample design change in 1995 that hinders matching.
Graphs are truncated at -60 and 60 percent.

Figure 1.10: Nominal Wage Growth Among Job Finders in the CPS by Year

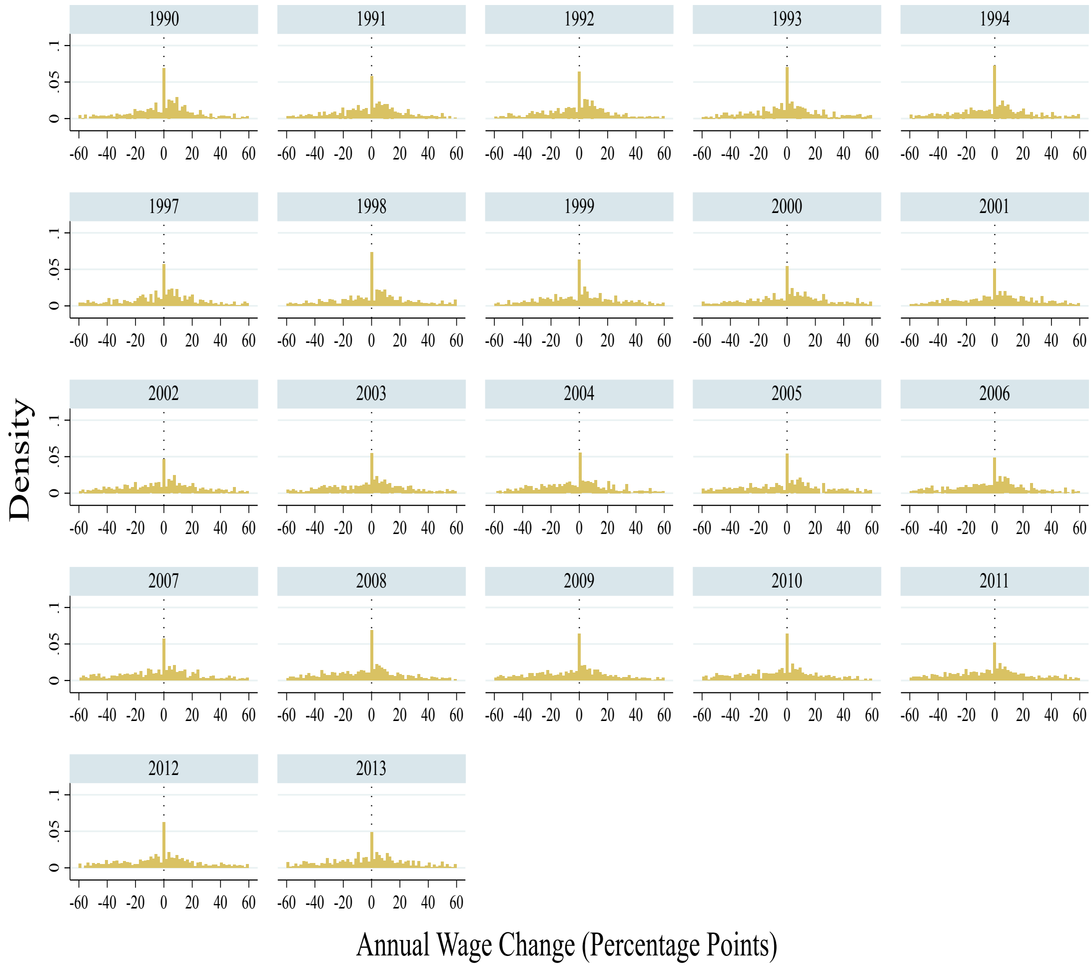

Notes:
Years 1995-1996 are exluded because of a sample design change in 1995 that hinders matching.
Graphs are truncated at -60 and 60 percent.

Figure 1.11: 2-Year Nominal Wage Growth Among Job Stayers in the PSID by Year

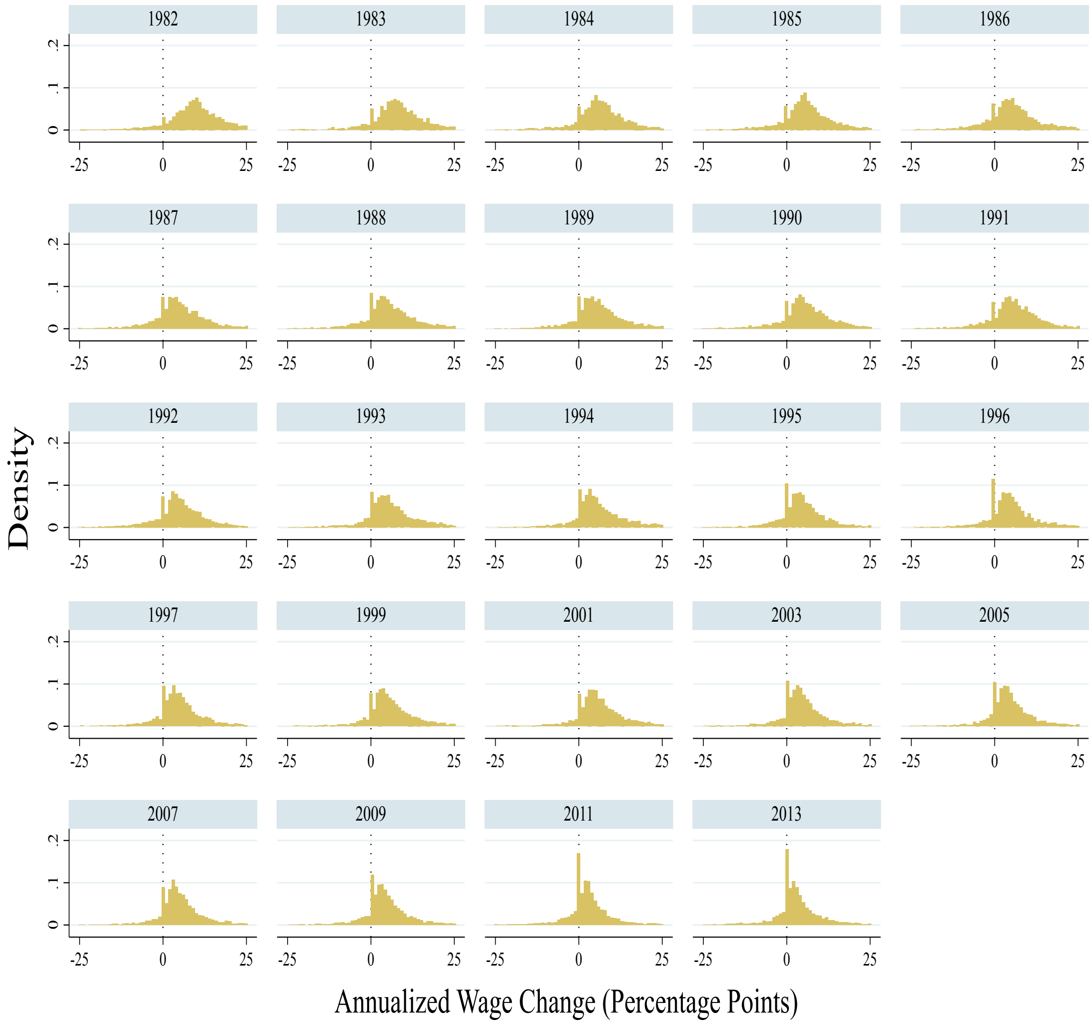

Note: Years show the annualized two-year wage growth rate from two years previously.

Figure 1.12: 2-Year Nominal Wage Growth Among Job Switchers in the PSID by Year

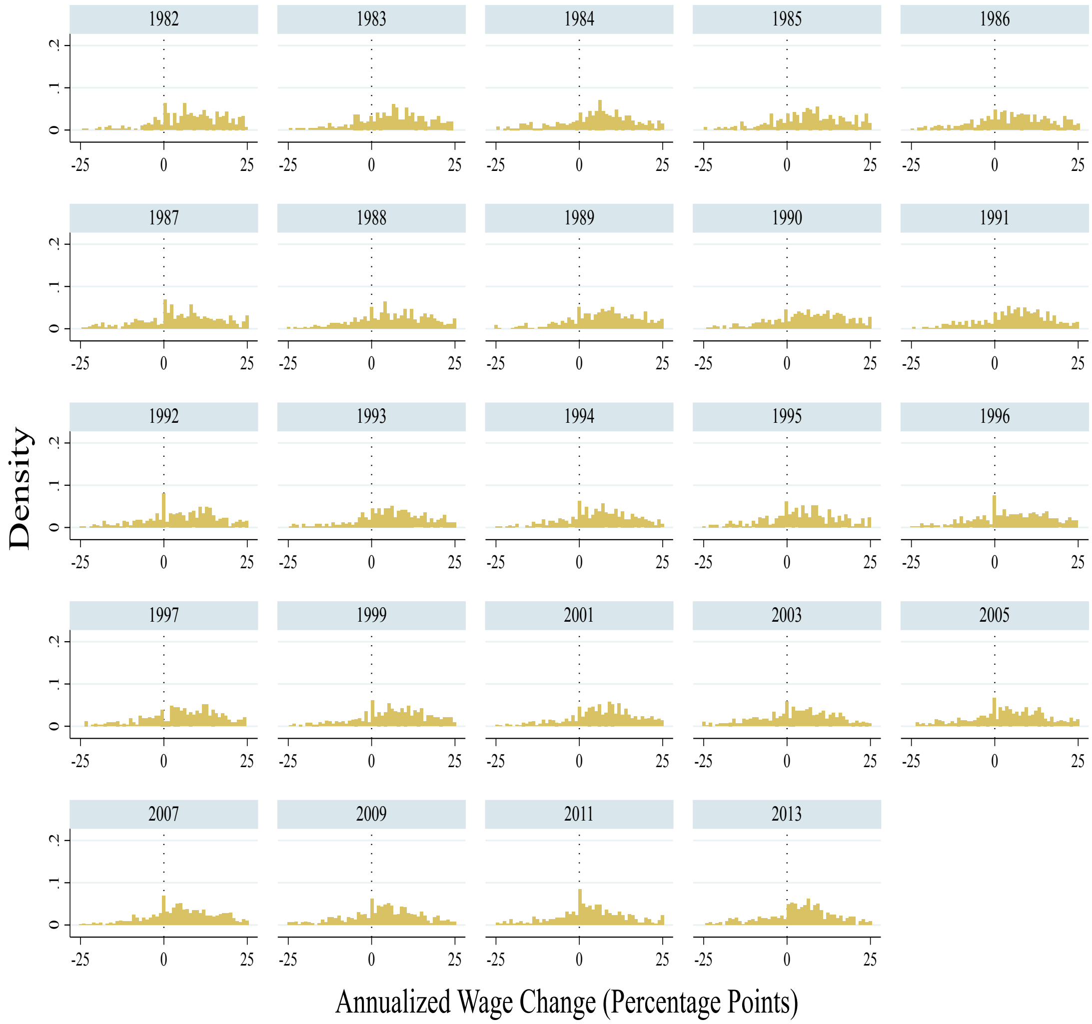

Note: Years show the annualized two-year wage growth rate from two years previously.

Figure 1.13: 2-Year Nominal Wage Growth Among Job Finders in the PSID by Year

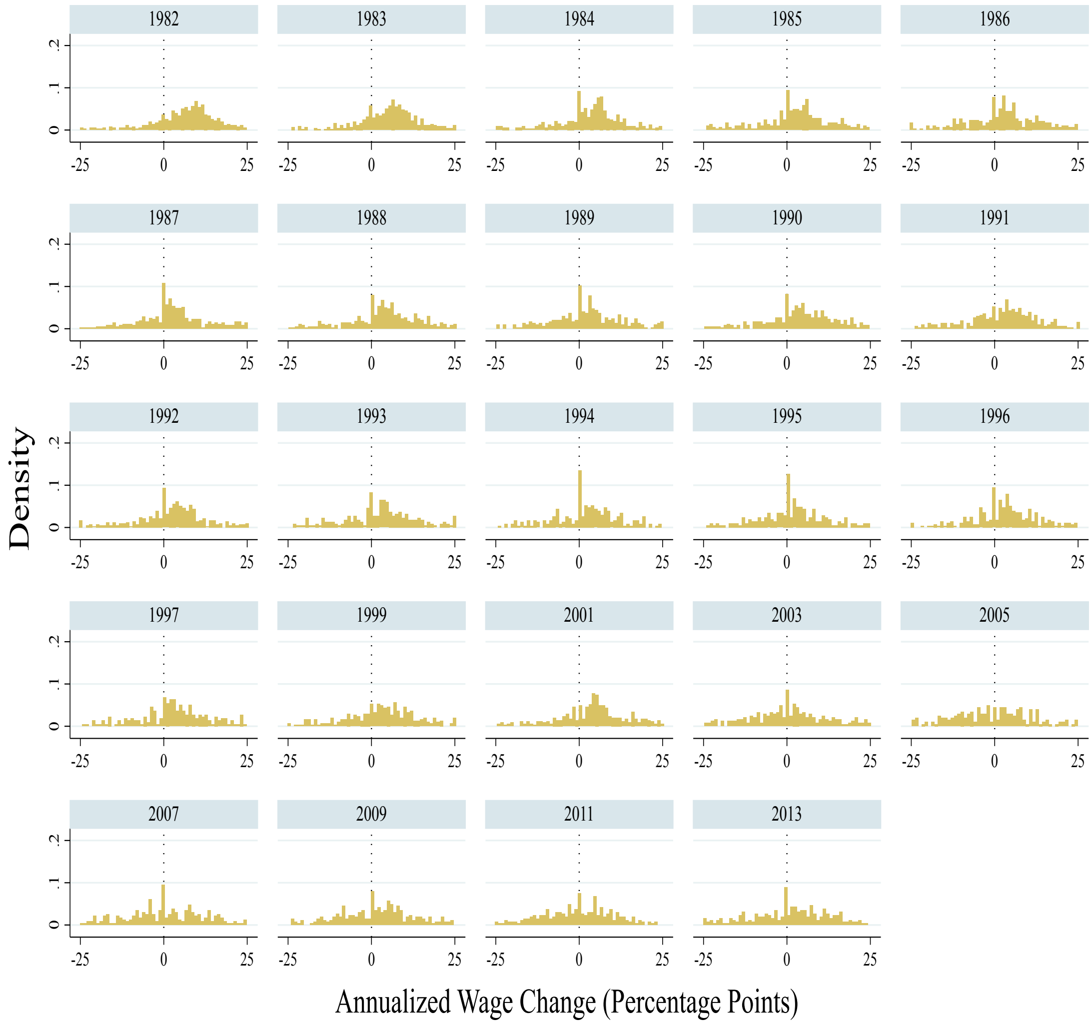

Note: Years show the annualized two-year wage change from two years previously.

CHAPTER II

Fiscal Policy

2.1 Introduction

The macroeconomic effects of fiscal policy are highly contested on both theoretical and empirical grounds. Since the beginning of the Great Recession, and especially once interest rates reached the zero lower bound, there has been a resurgence of interest in the effectiveness of fiscal policy. The lack of consensus on how fiscal levers affect key macro-economic variables extends even to the wisdom of employing countercyclical government spending policies to begin with. In this paper, we use professional forecasts of government spending, revenues, output and prices to estimate both the shocks to fiscal policy and to estimate the effects of changes in government spending on output and prices via local linear projection.

There are two main empirical approaches for obtaining shocks to fiscal policy; either based on Vector Autoregressions (VARs) or via narrative techniques. Standard VAR analyses follow the lead of Blanchard and Perotti (2002) and apply a Cholesky decomposition to back out the shocks to fiscal policy. The main advantage of VARs is their ability to capture complex patterns in the data with very few structural assumptions. The standard narrative approach tends to follow the Ramey and Shapiro (1998) "war dates" analysis where they created a dummy variable to capture government spending shocks from news records of military buildups sourced from *Business Week* magazine. The main advantage of the narrative approach is that the war dummy variable is exogenous to the business cycle.

Both approaches have been criticized recently. The main disadvantage of a VAR is that, due to the small set of included variables, the identified shocks are not truly exogenous innovations since they do not adequately control for future policy actions even though this information might be pub-

This chapter is co-authored with Aditi Thapar.

licly available. Ramey (2011) finds that changes in defense and non-defense spending are anticipated by private agents a few quarters prior to the actual changes in spending. She finds that professional forecasts Granger-cause VAR shocks, implying that VAR shocks are not true innovations. The main disadvantage of the narrative approach is that the results are driven by the military buildup of World War II and the Korean War, which leads to an extremely small sample of innovations.

In this paper, we use the forecast errors of professional forecasters to estimate fiscal shocks and to estimate the dynamic effects of fiscal policy on output, private sector output, and prices. By doing so, we implicitly gain the information set of the forecaster without the need to add more control variables to our estimation. Professional forecasts contain more information than any viable VAR, including information about future expected changes to policy levers. As a result, forecast errors are "clean" of the kind of information that is readily predictable to economic actors but which is not foreseen by, and may even be mis-attributed in, backwards-looking VARs. We obtain a government spending multiplier between 1 and 1.6, with a slightly lower value using defense spending shocks. The multiplier appears to be larger in more recent years, including the Great Recession. This is consistent with the observation that crowding-out of private investment is less likely when interest rates are at or near the zero lower bound.

Simultaneous estimation of output, spending, and (net) taxes requires several assumptions about how to treat contemporaneous effects. As discussed by Perotti (2011) there are four key challenges to any technique for the estimation of fiscal policy. First, typically VARs do not account for information about future government policy that might be publicly available to agents, which he terms the "anticipation problem". Second, government spending is often assumed to be unrelated to contemporary nonfiscal variables in VARs for identification purposes (the "exogeneity problem"). Third, most of the variance in government spending is driven by big war buildups, and these infrequent instances drive the results (the "variance problem"). Fourth, government spending enters into other economic variables in dynamic, variable ways, none of which are captured by simple negative wealth effects a la the neoclassical depiction (the "externality problem").

Empirical analysis of fiscal policy on the macro-economy has evolved over time. Below we discuss some papers that are relevant to our approach to estimating the effects of fiscal policy. Ramey and Shapiro (1998) focus on defense buildups in anticipation of the Korean and Vietnam Wars, as well as the Carter-Reagan buildup in response to the Soviet invasion of Afghanistan, as potentially exogenous shocks to government expenditures. This approach reflects a careful treatment of the exogeneity problem, as well as use of the variance in government spending. It sidesteps the externality problem by focusing on an aspect of government spending that is unlikely to enter into the utility function in complicated ways. They find that total GDP increases but private GDP decreases within two years of the initial buildup, reflecting a decrease in consumption and a spending multiplier less than 1.

Blanchard and Perotti (2002) use a Structural VAR with real GDP, government spending (federal, state, and local consumption and investment) and net taxes (receipts minus transfers to persons and net interest). They identify the shocks as follows: taxes are taken to be affected contemporaneously by output according to sources of taxes and various corresponding elasticities, summed together. Government spending is assumed not to be contemporaneously affected. Use the "cleaned" tax and spending data as instruments to estimate the contemporaneous effects of taxes and spending on output. Then, they order the VAR so either the tax decision comes first or the spending decision comes first (i.e., the spending increase doesn't affect taxes, or the tax changes don't affect spending). They find the order is not very important to the results. They add a time trend for drift, and dummy variables for the 1975q2 tax cut at each lag.

Blanchard and Perotti conclude that spending shocks lead to positive GDP effects, and tax shocks lead to negative GDP effects. The stochastic trend leads to smaller effects on GDP from spending and larger and more persistent negative effects on GDP from tax shocks. Their estimate of the fiscal multiplier is about 1.3, implying higher consumption spending. Interestingly, they find that investment decreases, which contrasts with their otherwise Keynesian results.

Perotti (2007) points out that use of dummies for government defense spending (Ramey-Shapiro dates) subsumes all shocks into the fiscal shock category; using standard SVAR techniques from his 2002 paper, he isolates the effects of fiscal policy and finds that consumption actually rises (more akin to either a Neo-Keynesian interpretation or non-separable utility of leisure and consumption). Noting the timing issues of SVAR techniques and dummy techniques (relevant economic decision-makers can foresee a spending increase or a tax increase), he also does analysis with annual data but excludes pre-1929, which is often interpolated. He finds again that increasing G also increases C. International comparisons yield the same results.

Romer and Romer (2010) look at the effect of tax changes on output growth using a narrative-based approach. They note that the variation in taxes coincides with, and is often explicitly designed in response to, many other factors in the economy which also affect growth. In order to account for the simultaneity issues and omitted variable bias, they focus on a subset of tax changes that are a) legislated, and b) appear to be implemented (based on the available narrative record) for reasons that are independent of these other factors. The motivations they find for legislated tax cuts roughly divide into the following categories: a) offset a change in government spending, b) offset other factors related to growth (e.g. counter-cyclical tax changes), c) address an inherited budget deficit, and d) achieve a long-run goal (smaller government, increased fairness, faster growth). Factors a) and b) are clearly correlated with other growth drivers, so exogenous tax changes are composed of c) and d). The Romers' data range from 1947 to 2007, and estimation includes 12 lags of the tax variable.

Their results are significant in that a 1 percent GDP increase in taxes leads to a 3 percent GDP change in output over the next three years. Tax increases in response to a budget deficit have much smaller effects than tax increases in category d), perhaps because of an interest rate response (explanation ours). Romer and Romer also look at timing issues, finding that the biggest results from implementation rather than announcement (in fact the announcement, when included with implementation, has the opposite sign, but is about one-third the size), and they look at component effects and find that investment has the largest response.

Ramey (2011a) contrasts VAR results with narrative-driven shocks to government spending, in particular military spending. Both types find a positive multiplier of government spending on output, but the VARs provide evidence that consumption and real wages rise, while papers using the narrative-based approach generally find consumption falling. As a result, the multipliers are usually higher for VAR approaches. She then finds that narrative-based shocks Granger-cause VAR shocks, meaning that the VAR shocks are not true innovations. In addition, professional forecasts Granger-cause VAR shocks as well. A more robust measure of expected changes in military spending using news forecasts yields a multiplier of .6 to .8, approaching 1.1 when WWII is included as well. Ramey also runs a VAR accounting for expectations by examining the one-period-ahead forecast error in government spending growth based on the Survey of Professional Forecasters. She finds that temporary rises in government spending do not stimulate the economy with this specification. She calls this approach EVAR, and it attempts to address the concerns with her approach using dummies to measure narrative-driven effects. The key concern with using the dummies is that they absorb all the shocks of the periods in question, regardless of whether other shocks are occurring simultaneously (such as the large tax increases during the Korean War).

Auerbach and Gorodnichenko (2012) create a smooth-transition VAR that measures the effect of log real government purchases (federal, state, and local investment and consumption) and log real government receipts (net of transfers to businesses and individuals) on log real GDP. The model distinguishes between contemporaneous effects (via a covariance matrix) and dynamic effects via lagged polynomials. They focus on the effect of government spending on output because they are skeptical of the SVAR's ability to truly exogenize the tax shocks (which appears to be a reference to Romer and Romer (2010)). The switching regime is smooth so every period estimates both an expansionary and a contractionary parameter matrix. AG incorporate expectations into their model by using the Survey of Professional Forecasters (SPF) estimates of GDP growth, the Research Seminar in Quantitative Economics (RSQE) forecasts of government revenues, and the Greenbook forecasts of government spending. They note that the results of their baseline VAR yield a series of residuals for government spending, and that same process based on forecasts of spending yields a series of residuals that is NOT uncorrelated with the other series. They are positively correlated, implying that some changes in government expenditures are predictable even given the information structure in the VAR.

Auerbach and Gorodnichenko state that ideally they would be able to add the forecasts to the VAR, but it doubles the size and halves the number of data points. They employ two different approaches to deal with this in a tractable way. First, they create forecast errors of output, government spending, and government receipts, and they use these errors to estimate the contemporaneous responses of those variables, which they then allow to propagate in an impulse response function that was previously estimated (under the assumption that all changes to these variables are unanticipated). Second, they include a forecast error for government spending in the VAR directly (as opposed to the full set of forecast errors) and accept the diminished sample size and slightly larger VAR specification. They try this two ways, the first being an addition of the growth forecast, the second being an addition of the forecast error in growth. Both yield higher estimates of the multiplier for both expansions and recessions.

The rest of this paper is organized as follows: section 2.2 details our empirical approach and contrasts it with those of the VAR and standard linear projection techniques. Section 2.3 describes the forecast data from our two sources, the Federal Reserve and the University of Michigan's Research Seminar in Quantitative Economics, as well as the challenges of using forecast data, and assesses the relative performance of professional forecasts versus those generated from VARs. In section 2.4 we highlight the basic empirical results of our methodology in terms of multipliers and impulse response functions. We perform various robustness checks, including 1) using defense shocks as exogenous drivers of government spending shocks, 2) altering our sample period, and 3) comparing our empirical results with those using standard linear projection. Section 2.6 concludes.

2.2 Model

We begin by describing the basic framework of a Vector Autoregression (VAR) when applied to the estimation of the effects of fiscal policy. Consider an $n \times 1$ vector of economic variables,

$X_t = \{G_t, T_t, Y_t\}$, where G_t represents real government spending, T_t represents real tax revenues, and Y_t is real GDP, say. X_t is determined by the history of changes to its components ε according to some functional form:

$$X_t = F(X_0, \varepsilon_t, \varepsilon_{t-1}, ..., \varepsilon_0). \tag{2.1}$$

A structural VAR with p lags represents an approximate linearization of this relationship,

$$AX_t = \Gamma + \beta Z_t + e_t^{VAR}, \tag{2.2}$$

where A is a $n \times n$ matrix, Γ is a $n \times 1$ vector, β is a $n \times p$ matrix, $Z_t \equiv (X_{t-1}, X_{t-2}, ..., X_{t-p})'$, and e_t^{VAR} is a $n \times 1$ vector of structural errors. The corresponding unstructured or reduced-form VAR is defined and estimated as:

$$\begin{aligned} X_t =& A^{-1}\Gamma + A^{-1}\beta Z_t + A^{-1}e_t^{VAR} \\ =& \widetilde{\Gamma} + \widetilde{\beta} Z_t + u_t. \end{aligned} \tag{2.3}$$

We can represent the reduced-form errors (u_t) as the one-step ahead forecast errors generated by the linear VAR model,

$$u_t^{VAR} = X_t - \mathbb{E}_{t-1}[X_t|Z_t]. \tag{2.4}$$

Note that in our approach we use professional forecast errors u_t^f to replace VAR-generated forecast errors u_t^{VAR}. In either case, the vector of residuals u_t is typically used to identify the structural shocks e_t. The structural errors in the VAR are related to the reduced-form errors by:

$$e_t^{VAR} = Au_t^{VAR}.$$

The structural shocks are recovered from the reduced-form errors by making assumptions on the structure of the A matrix. A standard assumption in the VAR literature is to identify structural errors by assuming that A is lower-triangular and that the structural shocks are independent. This recursive identification assumption, also known as a Cholesky decomposition, implies that any contemporaneous covariance between variables is attributed to the variables ordered earlier in the VAR.

To obtain the impulse response functions, note that we can represent the linear VAR with p lags as an $MA(\infty)$ process as given by,

$$X_{t+s} = \sum_{i=0}^{\infty} D_i e^{VAR}_{t+s-i} = D_0 e^{VAR}_{t+s} + D_1 e^{VAR}_{t+s-1} + ... + D_{s-1} e^{VAR}_{t+1} + ... \quad (2.5)$$

The $s-$step ahead forecast error from the VAR can be represented as,

$$X_{t+s} - E_t X_{t+s} = \sum_{i=0}^{s-1} D_i e^{VAR}_{t+s-i} = D_0 e^{VAR}_{t+s} + D_1 e^{VAR}_{t+s-1} + ... + D_{s-1} e^{VAR}_{t+1}. \quad (2.6)$$

Since the e^{VAR}_{t+s} are assumed to be independent, D_{s-1} can be interpreted as the effect of a structural shock next period to the variable of interest s periods later.

We use the forecasts of professional forecasters to obtain empirical estimates of short- and long-run fiscal multipliers via local projection methods, see Jordá (2005) and Thapar (2008) for two applications of these methods.[1] In this paper, we closely follow the basic approach developed in Thapar (2008). Using professional forecasts, we can ignore estimation of equation 2.3 above and calculate residuals directly, based on the insight that $u^f_t = X_t - \mathbb{E}_{t-1}[X_t|I^f_{t-1}]$ is the set of one-step-ahead forecast errors, where I^f_{t-1} is the information set available to forecaster f at time $t-1$.[2] One advantage of our approach is that the information set used to produce our forecasts $\left(I^f_{t-1}\right)$ is presumably much larger than the information set of the VAR (Z_t). By construction, the VAR forecasts can only include the information that is contained in the variables included in the system.

Forecast errors from professional forecasts can identify structural shocks ε_t directly via a Cholesky decomposition. In our baseline case we include only taxes, T, government spending, G, and output, Y, so $X_t = (G_t, T_t, Y_t)'$. The identification process of uncorrelated structural shocks $\varepsilon_t = (\varepsilon^\tau_t, \varepsilon^g_t, \varepsilon^y_t)'$ from the one-step-ahead forecast errors $u_t = (u^\tau_t, u^g_t, u^y_t)'$ requires 3 restrictions. We adopt the identifying method of Blanchard and Perotti (2002), so the (normalized) block:

$$u^\tau_t = \alpha_{11}\varepsilon^g_t + \alpha_{12}\varepsilon^y_t + \varepsilon^\tau_t \quad (2.7)$$

$$u^g_t = \alpha_{21}\varepsilon^\tau_t + \alpha_{22}\varepsilon^y_t + \varepsilon^g_t \quad (2.8)$$

$$u^y_t = \alpha_{31}\varepsilon^\tau_t + \alpha_{32}\varepsilon^g_t + \varepsilon^y_t \quad (2.9)$$

has two order restrictions, namely that $\alpha_{11} = 0$ and $\alpha_{22} = 0$.[3] The third restriction relates u^τ_t to

[1] Ramey (2016) summarizes developments in the literature on the propagation of macroeconomic shocks.

[2] In the VAR above, $I^f_{t-1} = Z_t$ and $\mathbb{E}_{t-1}[X_t|Z_t] = \widetilde{\Gamma} + \widetilde{\beta} Z_t$, or the best linear fit in a mean squared error sense.

[3] The choice of $\alpha_{11} = 0$, which assumes that shocks to government spending do not affect taxes in the contemporaneous quarter, is relaxed and substituted with an alternative restriction $\alpha_{21} = 0$ without any significant change in results for Blanchard and Perotti. We also perform an alternative analysis with this substitute ordering and reach a similar conclusion.

u_t^y based on a weighted set of short-run elasticities between taxes and output. Our estimate of $\eta_{\tau,y}$, the within-quarter elasticity of taxes with respect to output, is 1.18 for gross federal receipts during the sample period for which we have GB forecasts (1978 to 2010) and 1.89 for net federal, state, and local receipts during the sample period for which we have RSQE forecasts (1994 to 2015). A detailed derivation of these values is available in an appendix at the end of the paper, labeled section 2.7. The net effect of this restriction is the creation of a cyclically adjusted tax shock.

By using forecast errors directly, we gain several advantages. First, our errors are not the product of the limited information set of a VAR, $\{X_{t-1}, X_{t-2}, ..., X_{t-p}\}$, but rather include all the information known to the forecasters at time $t-1$. Second, by expanding the system of equations 2.7-2.9 to include other macroeconomic variables of interest (e.g. inflation or the instrument of monetary policy), we can generate additional impulse responses or simply control for more shocks and better isolate the effect of shocks to the existing set of variables.

Increasing the block of equations above is possible in a SVAR format as well; the number of required identifying restrictions r simply increases with the number of variables n according to $r = \frac{n(n-1)}{2}$, just as in our process above. A small advantage might be conferred to the forecast error method in that some identifying restrictions could seem more plausible within the context of a larger information set. The more practical problem in a VAR setting, however, is that the number of degrees of freedom shrinks much more rapidly. For a VAR with 4 lags, adding a 4th variable increases the number of coefficients to be estimated in the first stage alone from $4 \times 3^2 = 36$ to $4 \times 4^2 = 64$, while adding a 5th variable increases the number of coefficients to $4 \times 5^2 = 100$. By dispensing with a linear best fit model to generate one-step-ahead forecast errors, we are potentially able to obtain more precise estimates of the relationship between structural shocks and forecast errors with a smaller data set.

Instead of tracing out an impulse response iteratively through the same linear model, we project our forecast errors at each of the N horizons for which they exist, $t+1$ to $t+N$, onto our structural shocks:

$$X_{t+s} - \mathbb{E}_{t-1}[X_{t+s}|I_{t-1}^f] = c + \beta^s \varepsilon_t + \theta_{t+s} \quad \text{for } s = 1, 2, ..., N, \tag{2.10}$$

The impulse response s periods into the future of the jth element of ε on the ith element of X is the (i,j)th element of β^s. The system of N equations in 2.10 can be estimated individually via least squares using a heteroskedasticity and autocorrelation consistent covariance matrix.

2.2.1 Multipliers

The impulse response functions estimated above provide us the elasticity of output with respect to government spending at various horizons. To convert the elasticities to multipliers, we follow Mountford and Uhlig (2009) and calculate multipliers as the integral or present value of the response of GDP or private GDP divided by the integral response of government spending.[4] The present value multiplier (ϕ) at horizon s is given by

$$\phi^s = \frac{\sum_{i=0}^{s} (1+r)^{-i} \beta^{Y,i}}{\sum_{i=0}^{s} (1+r)^{-i} \beta^{G,i}} \times \frac{Y}{G} \tag{2.11}$$

where $\beta^{j,i}$ refers to the impulse response of variable $Y = \{GDP, private\ GDP\}$ at horizon $i = 1, 2, ..., N$ and r is the real interest rate. The real interest rate is assumed to be the average real interest rate over the relevant sample period.

2.2.2 Comparison of Projection Methods

Our approach differs significantly from a VAR, as described in detail above, but also somewhat from Jordà's popular local projection method, which uses the actual data (as opposed to forecast errors) to estimate impulse response functions. As discussed in Ramey (2016), given a series for the structural shocks one can obtain impulse response functions via Jordà's local projection method by estimating a sequence of regressions given by,

$$Y_{t+s} = b^{Y,s} e_{1t} + control\ variables + \eta_{t+s} \text{ for each } s = 1, 2, ..., H, \tag{2.12}$$

where Y_{t+s} is the variable of interest, say GDP, at time $t+s$, $b^{Y,s}$ is the impulse response of Y at horizon s to a e_{1t} shock in period t. This method can be applied to shocks obtained from any source, such as a Cholesky decomposition, reading news about defense contracts, or through excess returns on the stocks of defense contractors, to name a few. The control variables are added to account for information that is available in period t, when the fiscal policy shock hits the economy. Jordà's local projection method is similar to the VAR projection in equation 2.5. Researchers using this technique typically add various control variables to account for information available to policy makers at the time of the forecast.

One of the main advantages of our approach to calculating impulse response functions, in equation

[4]The impulse response at horizon s, β^s, is the percent response of Y to a one percentage point shock to Y, so $\beta^s = \frac{dY/Y}{dG/G}$.

2.10, is that by using the forecast errors of professional forecasts we have automatically controlled for a larger information set than can be included by adding a small subset of variables into the above estimation equation. Note our projection in equation 2.10 is similar in nature to the projection produced by a VAR in equation 2.5, with one major difference. The expectations in the VAR projection are based only on the information set of the variables included in the VAR system whereas the information set of professional forecasters includes all the information that is available to the forecasters at the time of the forecast. Asymptotically equations 2.12 and equation 2.10 should lead to the same results, however we should expect to see efficiency gains since our approach (equation 2.10) controls for a larger subset of information. As Ramey (2016) summarizes, "*Because the Jordà method for calculating impulse response functions imposes fewer restrictions, the estimates are often less precisely estimated and are sometimes erratic.*" We test this proposition in section 2.5.3 and find that although the point estimates from both approaches are similar, there are strong efficiency gains from using our methodology.

2.3 Data

In this section, we describe the data that are used in this paper. To implement our approach we need forecasts (and forecast errors) of GDP, government spending and tax and transfer information, all on a quarterly basis. We use quarterly forecasts from 2 separate sources, the Federal Reserve Board's Greenbook (GB) forecasts and the University of Michigan's Research Seminar in Quantitative Economics's (RSQE) quarterly forecasts. The GB forecasts, which are made public 5 years after they are created, span from 1966 to 2010 with the relevant data at quarterly intervals. RSQE data, on the other hand, are available from 1983 to 2015.

2.3.1 Greenbook Forecasts

The Greenbook is the colloquial name given to the official report titled "Current Economic and Financial Conditions – Summary and Outlook" that is produced by the research staff at the Board of Governors of the Federal Reserve System. This report is prepared for the Federal Open Markets Committee (FOMC) prior to every FOMC meeting. It includes the forecasts of the US economy and is available to the FOMC members six days prior to every scheduled meeting.

The Greenbooks are publicly available but with a six year lag, which constrains the end of our sample period to the end of 2010. Although the Greenbook forecasts are available from 1966 onwards, it was only beginning in 1975 that the forecast horizon was extended to be a minimum of four quarters ahead in every forecast. The maximum forecast horizon was increased from five quarters ahead to seven quarters ahead in 1979. Eight quarter ahead forecasts are available beginning in 1988. Currently we have data beginning in 1978 but we plan to extend the sample to include the 1975-78 period.[5]

Some of the forecasts of the variables of interest are available in levels, others in growth rates, while a select few are in both levels and growth rates. Still others, like GDP and the GDP deflator, which used to be available in both levels and growth rates began to be published only in growth rate terms beginning in 2005. In this paper, we decided to use the growth rates for all the variables of interest.

The Greenbook forecasts estimate the growth of future real government expenditures on consumption and investment, including federal, state, and local expenditures. We use this set of estimates as our measure of government spending.

Our expected growth rate of government receipts is a gross measure, for want of forecast data on transfers from the government to persons. In addition, state and local receipts are not included in the Greenbook forecasts, so all analyses of Greenbook data that include taxes/receipts are using forecast errors of federal gross receipts, indexed by the GDP deflator.

We are interested in estimating the cumulative effects of government spending on the variables of interest at different horizons. Since we are working with growth rates rather than levels, for the purpose of estimating impulse response functions and multipliers we calculate cumulative growth rates for each variable of interest, $g^x_{t+s} = \frac{X_{t+s}}{X_t}$, and project forecast errors of cumulative growth rates, $g^x_{t+s} - \mathbb{E}_{t-1}[g^x_{t+s}|I^f_{t-1}]$ onto our shock series at each $s = 1, 2, ..., H$.

2.3.2 RSQE Forecasts

RSQE produces quarterly forecasts at regular intervals that have shifted slightly over time. While its current practice is to publish detailed write-ups of its forecasts in March, May, September, and November, in the past its August set of projections was its featured (and sometimes only) 3rd-quarter forecast. In addition, the June forecast was more frequently published in the recent past than the one in May. We draw one forecast from each quarter between 1994 and 2015, according to the following conventions: 1) we use the March, May, and November forecasts throughout our analysis, and 2) we use the August forecast from 1983-2007, and the September forecast from 2008-2015.

[5]Until 1974, the forecasts for the first ten months of the year were only available for the current calendar year, while in the last two months each year the forecast horizon was extended to the next calendar year. As a result of this forecasting convention, we have forecasts of between one and four quarters ahead depending on the quarter. Due to the short forecast horizon of the data in the 1966-74 sample, we focus on the post-1974 sample period for this paper.

The horizon for RSQE forecasts also varies over time. From 1994-1995, forecasts before November included projections through the end of the current year and also the next two years, between 10 and 12 quarters in total (counting the projection for the current quarter, which we do not use). November included an additional year, giving us 13 quarters of projections (and 12-quarters-ahead forecast errors, excluding the current quarter). Beginning in 1996 and extending to 2011, the August forecast included the third full year of projections. From 2012 to the present, RSQE forecasts include between 13 and 16 quarters of projections. The recentness of this change, however, means that we have a 15-step-ahead forecast error for only one quarter, the March 2012 projection for 2015q4.

RSQE forecasts are available in the more recent years, but data from 1952-1982 were destroyed in a fire. Published results, however, are available for 1952-1982 and have never been assembled; they comprise a unique (albeit annual) data set that may be promising if we supplement our existing research agenda with a modified approach to utilize the sparser data (for example, we do not have net government receipts in the published results), shorter horizon, and annually aggregated format.

The RSQE measure of government spending that we employ is the same as for the Greenbook data: total real government expenditures on consumption and investment. For receipts, we construct a measure of real government net receipts, which includes state and local receipts as well as federal and non-federal transfers to persons (excluding interest payments). As with Greenbook forecasts, we work with cumulative forecast errors in growth rates.

2.3.3 Real-Time versus Current-Vintage Data

When using forecast data, the primary issue in obtaining forecast errors is the choice of whether to use the data in its current vintage or to use the value of the variables that were available around the time the forecasts were made (real-time data). Many studies using forecast data tend to use real-time data to construct forecast errors. In this paper we do not follow this convention for two main reasons. First, government spending estimates tend to get revised over time much more than most other variables, and second, since we are working with cumulative forecasts it is difficult to consider real-time data. We describe the issues in detail below.

The main argument for basing forecast errors on real-time data is that calculating forecast errors based on the current vintage might artificially inflate the forecast errors and incorrectly attribute definition changes as shocks. For example, in 2013 the definition of GDP was expanded to include intellectual property. A forecaster in 1980, however, would likely have forecast GDP based on the definition of GDP that was used in 1980. The main argument against using real-time data, and instead using the current vintage data, is that the purpose of forecasting is to obtain the true value of a variable. The most recent versions of data, with all the revisions, are more likely to represent the true value of the variable. Here too, however, there is a problem. Older data, which have been revised more, are likely to have larger errors while newer data that are more closely related to the current definitions will have smaller errors.

The Bureau of Economic Analysis (BEA) publishes three estimates for the NIPA accounts. The advance (first) estimate, the second estimate, and the third estimate are published one, two, and three months respectively after the end of any given quarter. Table 2.1 compares the currently available (current vintage) data with real-time data. We report the correlations between the current vintage data and real-time data as it existed one and two quarters after the end of any given quarter. Most of our real-time data were collected from the Greenbook dataset, which along with forecasts also contain the data available in the past few quarters. Depending on the exact date of the FOMC meeting, the one-quarter-old real-time data is either the advance or second estimate, while the two-quarter-old is the third estimate. Every summer the BEA conducts an annual revision of the data over the past three calendar year and approximately every five years the BEA conducts a comprehensive revision. The third estimate is changed any time the BEA makes major changes.

The last column of table 2.1 contains the correlations of the current-vintage data for all the variables of interest. As expected, the two-quarter-old real-time data is slightly more correlated with the current data than the one-quarter-old real-time data. For most of our variables, the correlation between the two-quarter-old real-time data and the current vintage data is 0.8 or higher. The one notable exception is government spending, the most important variable in our analysis, with a correlation coefficient of only 0.6 with current data.

Figure 2.1 illustrates the quarterly annualized growth rate of government spending and GDP. The first issue with using real-time data as the proxy for the actual data in constructing forecast errors can be observed in the middle panel of the left column of the figure. Government spending tends to be revised even after the third estimate. The period between 1981 and 1988 appears to be responsible for most of the subsequent revisions, and the current vintage data are much less volatile than in the initial telling. GDP growth rates, however, are more similar between the real-time series and the current vintage data. We consider the revisions to government spending in the pre-1990 period to be a serious concern with using real-time data to construct forecast errors.

The second issue with using real-time data is a result of the analysis in this paper being based on growth rates of macroeconomic variables as opposed to levels, due to the restrictions of the Greenbook dataset as discussed above. Our use of cumulative growth rates of our variables of interest creates a dilemma over how to calculate forecast errors when using real-time data. The projected growth rate between today and five quarters ahead can be easily calculated from the Greenbook and RSQE datasets. To obtain the forecast error, however, we need a measure of the actual growth rate between today and five quarters from today. It is not clear at all, even in theory, which real-time data are appropriate to use, especially if an annual or comprehensive revision occurs between the forecast and the realization of the variable of interest.

Due to the issues with revisions to government spending and complications due to using cumulative growth rates, as described above, in this paper we use current data to construct forecast errors.

2.3.4 Comparison of Forecast Errors

In this section we compare the various performance measures of the Greenbook and RSQE forecasts of GDP, government spending, and receipts. We compare their forecast errors with several VAR specifications, all of which include 4 lags of each variable. The VAR specifications are chosen to highlight key points. There are two sets of choices involved. The first choice is whether to use the current vintage data or to use real-time data. The second choice is whether to estimate only one VAR over the entire model, or to estimate a rolling VAR, where the VAR is estimated every period and its predictions $(Y_{t+i}, G_{t+i}, T_{t+i})'$ are based on its simulated values of $(Y_{t+j}, G_{t+j}, T_{t+j})'$ for $j = 1, 2, ..., (i-1)$ rather than the actual values of those variables.

We obtain one- to eight-quarter-ahead forecasts, and forecast errors, from four VARs (see table 2.2). The first VAR specification (*VAR1*) uses current vintage data and is estimated only once over the entire sample period. This specification benefits from consistent definitions over the entire sample and from using the entire sample period to estimate the best linear relationship between the variables. The second VAR specification (*VAR2*) also uses the current vintage data, but it is estimated on a rolling basis and without access to the current period's output, receipts, and government spending (which would not be known at the time a forecast must be made). The third VAR (*VAR3*) is estimated only once over the entire sample period using our real-time data set. The fourth VAR specification (*VAR4*) is a rolling VAR that is estimated using real-time data, which most closely mimics the forecasts that would be made by a professional forecaster. *Ex ante* we expected *VAR1* to perform the best, since it uses current vintage data and is estimated only once over the entire sample, and *VAR4* to perform the worst.

Table 2.3 compares the mean error, root mean-squared error (MSE), and mean absolute error (MAE) of each forecast series over two sample periods, 1983-2010 (the sample common to the available Greenbook and RSQE forecasts) and 1990-2010. The latter sample period is included due to the aforementioned irregularities of government spending growth in the real-time data versus the current-vintage data during the 1980s.

Unsurprisingly, *VAR1*, which uses current vintage data and is estimated only once over the entire sample period, is the best performing of the VARs by most measures. Its MAE and MSE are lower for all three variables over the short run for both sample periods. The interesting exception is *VAR4*, which although it generally does not outperform the first specification, manages to equal and even beat *VAR1* in GDP forecast errors for selected periods.

The two forecast series significantly outperform the VAR specifications using real-time data for both spending and receipts. The one-period-ahead forecast errors for the Greenbook and RSQE series of government spending, respectively, have a MSE of 4.19 and 3.56 between 1983 and 2010, compared with 5.31 and 4.98 for the two VARs. They do even better in receipts; with MSEs of 10.43 and 10.88 respectively, they outperform all 4 VARS, the best of which registers a MSE of 13.68. In GDP, they outperfom all 4 VARs in the short run, with the Greenbook forecasts performing the best at a MSE of 2.19 one period ahead. Over the longer term, this advantage attenuates in GDP and quarterly growth rates do not appear to be more accurate 5 to 8 periods ahead than the best performing VAR specifications.

In the 1990-2010 sample period, the Greenbook and RSQE forecast errors exhibit lower MSE in all three variables than all VAR specifications, and lower MAE in government spending and government receipts.

Table 2.3 also shows the effect of averaging the GB and RSQE forecasts, which improves performance even further. The average of the Greenbook and RSQE forecasts has the lowest MAE and MSE of all specifications for the larger sample period with the exception of government spending compared with the current vintage VARs. In the later sample period, it outperforms everywhere except in later periods for GDP.

While the Greenbook appears to have slightly better performance than RSQE in its GDP forecasts, the further-ahead forecast errors do not perhaps tell a complete story; a forecast's performance over a longer horizon is probably better described by its cumulative error than its errors in each period. Table 2.4 compares the GDP forecasts for the two institutions based on cumulative forecast errors. The performance is much closer than in table 2.3; the Greenbook performs slightly better measured by MSE, while RSQE outperforms based on MAE.

2.4 Results

This section discusses the main results from the estimation of the effects of government spending shocks on key macroeconomic variables in the United States in the past 30 years. All results shown are based on projections onto structural errors derived by ordering (cyclically-adjusted) taxes first, government spending second, and output last. Switching the order of government spending and taxes has minimal impact on the impulses, standard errors, or multipliers.

The top panel of figure 2.2 shows the impulse response of real GDP obtained from a baseline VAR that includes the log levels of government spending, government receipts, real GDP. This small subset of variables is the standard set of variables that are used in the literature. The bottom panel shows the impulse response of the deflator from a VAR that includes the GDP deflator in addition to the variables in the baseline VAR. Figure 2.2 is presented to provide a baseline comparison for our approach. It is the only figure where the confidence bands are the 68% confidence intervals, as is standard in the VAR literature. The impact response is a 0.26 percentage point change in output, which then accumulates to an increase of 1.4 after eight quarters. The results are significant at the 10% level, for output, only for the first two quarters. For the GDP deflator, we find that prices decrease on impact and continue decreasing over our forecast horizon. The results here are significant at the 10% level.

Figures 2.3-2.5 show our baseline results of the effects of a government spending shock to GDP, private spending, and prices, estimated using equations 2.10 above. The results should be interpreted as the cumulative effect to the variable of interest of a 1 percentage-point shock to real government spending, estimated from the first available set of forecasts for each organization through 2010 (1978-2010 for Greenbook, 1983-2010 for RSQE).

In each figure the top panels report the estimated impulse response function using the Greenbook dataset while the bottom panels report the impulse response function for the RSQE dataset. The solid lines are the impulse responses, while the dashed lines are the 90% confidence intervals (solved analytically). Although the convention in the VAR literature has converged to report only one standard deviation bands (68% confidence intervals), our approach allows us to work with conventional measures of significance. As discussed in the data section above, the limitation to the number of horizons for which we can estimate impulse responses is based on data availability. Greenbook results are shown out to 8 quarters while the RSQE results are shown out to 12 quarters, which represents the maximum horizon of each dataset.

Figure 2.3 shows the impulse response function of real GDP. We project GDP cumulative forecast errors on government spending shocks according to the system of equations 2.12, where the first column uses $\varepsilon_t = \varepsilon_t^G$ and the second column uses $\varepsilon_t = (\varepsilon_t^T \ \varepsilon_t^G \ \varepsilon_t^Y)'$. We find that under both specifications GDP increases immediately by about 0.2 percentage points with both Greenbook and RSQE data before gradually rising. The Greenbook estimates are significantly different from zero for most periods, while the RSQE results indicate a smaller response (near 0 in the short term) with a larger confidence interval.

Figure 2.4 shows the effect of a 1 percentage-point government spending shock on private GDP, calculated as GDP minus government spending. Not surprisingly, given the results shown in figure 2.3, the impulse response derived from Greenbook estimates is mostly positive, beginning near zero for the first two periods before picking up, while the RSQE response function is more ambiguous. The IRF from RSQE forecasts remains negative for 7 quarters before becoming positive in the first specification, and negative for the first 10 quarters in the second specification. None of the results for either impulse response functions are statistically significantly different from zero. The lack of a significant response of private spending to a government spending shock, as well as the initial fall in private spending, has been found in other work by Ramey (2013).

Finally, in Figure 2.5 we note an interesting effect of government spending on prices. Increased spending appears to lead to a decrease in the GDP deflator. The qualitative result is the same using both Greenbook and RSQE forecasts (and for both specifications) and is statistically significant for most quarters after a shock. This "fiscal price-puzzle" has also been reported by Canova and Pappa (2007) and Mountford and Uhlig (2009).

Our estimates of the government spending multiplier are presented in tables 2.5 and 2.6. Using Greenbook forecasts imply that, over the 1978-2010 sample period, a government spending shock yields a fiscal multiplier of just under 1 on impact. The multiplier rises to 1.2 one year after the shock and 1.4-1.5 six quarters after the shock. The government spending fiscal multiplier on private GDP is near zero or slightly negative in table 2.5, before beginning to rise in the later quarters.

Table 2.6 presents the same results based on RSQE data for the 1983-2010 sample period. The RSQE multipliers, in contrast to the Greenbook multipliers in table 2.5, are less than 1 for GDP until 10 quarters after the shock. In addition, the private GDP fiscal multiplier is negative through the same timeframe. These results are consistent with the lower impulse responses in the RSQE panels of figures 2.3 and 2.4.

2.5 Robustness

In this section we exploit the flexibility of our methodology to discuss the results of various robustness exercises. We begin by discussing the impulse response functions and multipliers using shocks to federal government spending on defense. Then we examine the importance of sample period choice on our results. Finally, we compare the results in our baseline projections with the Jordà method.

2.5.1 Defense Shocks

Much of the multiplier literature emphasizes the role of defense spending shocks on GDP, as defense spending is in many cases orthogonal to current macroeconomic conditions (see Ramey and Shapiro (1998)). Our approach, in theory, avoids several of the pitfalls that lead researchers to focus on defense shocks; nevertheless, we examine the effect of defense shocks on GDP, private spending, and the GDP deflator here for comparison with our baseline approach.

We project forecast errors of our variables of interest on government spending shocks instrumented by defense shocks, defined according to the same process as government spending shocks.[6] We do not have defense spending projections in Greenbook forecasts until 1981q4, so our sample period spans from 1981q1-2010q4. We do not have defense spending in our RSQE forecasts for 1983-1993, so we do not show the results for RSQE data here.

Figure 2.6 compares our baseline projections with those from instrumenting government spending with defense shocks. In each case, instrumenting for defense shocks mutes the baseline result. The impulse response function for GDP begins at a similar level and rises slightly more slowly than in the baseline, but it is no longer statistically significantly different from 0 at the 90 percent significance level. For private GDP, the response is almost indistinguishable from 0 for most quarters after a shock. Finally, the GDP deflator falls by less than would otherwise be expected, and it borders on statistical significance in a comparison with no effect.

Table 2.7 shows the multipliers of GDP and private GDP based on the projections above. For GDP, the multiplier stays in the range of .7 to .8 for most of the sample before creeping up in later periods. Private GDP holds at a negative value until the very end, in accordance with the impulse responses shown in figure 2.6.

[6]When government spending is ordered first, the forecast error is the structural shock. When government spending is ordered second, which is the result shown, the structural shock is derived from equation 2.8.

2.5.2 Sub-sample

An interesting question is whether or not the efficacy of fiscal policy changed during the great moderation. The advantage of our approach, especially relative to a VAR-based analysis, is that since we treat all expectations and forecast errors as data it is easier for us to answer such questions over small samples. We discuss the effects of a government spending shock on GDP with progressively more recent sample periods and show results excluding the Great Recession.

Figure 2.7 compares the effects of government spending shocks on GDP, private spending, and the GDP deflator over the different samples. Qualitatively, the results are very similar to our baseline. The initial impact on GDP is close to one-for-one, the results are statistically significant at most horizons, and the multiplier increases over the forecast horizon. Quantitatively, the main difference in the results stems from the fact that the impulse responses for both GDP and private GDP rise as we exclude earlier data. The exclusion of the Great Recession from our sample period, shown in the lower right quadrant of each figure, lowers the estimated impulse response, implying that the multiplier during that period was quite strong. This reinforces the mainstream belief that multipliers are most pronounced when interest rates are low, precluding crowding out.

2.5.3 Comparison to Jordà Projections

In section 2.2.2 we compared our method with that of Jordà (2005). The main difference, again, between our approach and Jordà's local projection approach is that we project structural shocks on s-period ahead forecast errors ($Y_{t+s} - E_t Y_{t+s}$), whereas the standard method involves projecting shocks onto the s-period ahead value of the variable, say Y_{t+s}. Here we discuss the effects of the two different estimation techniques.

Figure 2.8 compares the impulse response functions for GDP, private GDP, and the GDP deflator based on the two methodologies. While both methodologies have similar paths for GDP and private GDP, as expected our impulse responses have greater significance. For the GDP deflator, we get a strong, and statistically significant, negative response of prices to a government spending shock whereas projecting onto current data implies a negligible, and statistically insignificant, response of prices.

A well documented problem in the literature on estimating the effects of macroeconomic shocks is the limited information set of a VAR-based analysis. The advantage of our approach is that by using forecasts of professional forecasters, we can control for a variety of information that is available at the time that a shock hits the economy. We do not need to include extra variables in our estimation. The main advantage of Jordà's projections, over our approach, is the possibility of estimating impulses for a longer horizon.

2.6 Conclusion

This paper uses Federal Reserve and a non-profit forecasting organization's forecast errors to estimate the effect of changes in government spending on output. It adapts a local linear projection methodology to estimate IRFs and multipliers that is robust to compounding errors due to misspecification of the data generating process. We obtain multipliers that are stronger than 1 under most specifications, with 1.6 being our best estimate of the multiplier about 2 years after a government spending shock. Multipliers utilizing defense spending are slightly smaller. The government spending multiplier appears to be increasing over time, and the Great Recession was likely an era in which the spending multiplier was quite high.

Forecast errors from the Greenbook and RSQE projections perform better on a variety of metrics than simple VAR residuals. In addition, the projection of cumulative growth rate errors onto our shock series allows us to obtain similar but much more precise estimates to a similar method of projecting the growth rates themselves onto the shocks (with additional controls). This performance improvement allows us to maintain that our estimate of IRFs for output are greater than 0 at the 90 percent confidence interval for two years. The drawback of this methodology is that we are unable to make longer-run estimates of the multiplier or IRF.

Finally, we have identified what appears to be a counterintuitive (but nonetheless previously documented) price response to government spending shocks. The GDP deflator falls after spending increases, contradicting the standard understanding of how demand-side shocks act to raise prices. Further work is necessary to uncover the exact mechanism through which prices fall (or appear to fall).

Figure 2.1: Real-Time versus Current-Vintage Growth Rates
1978-2010

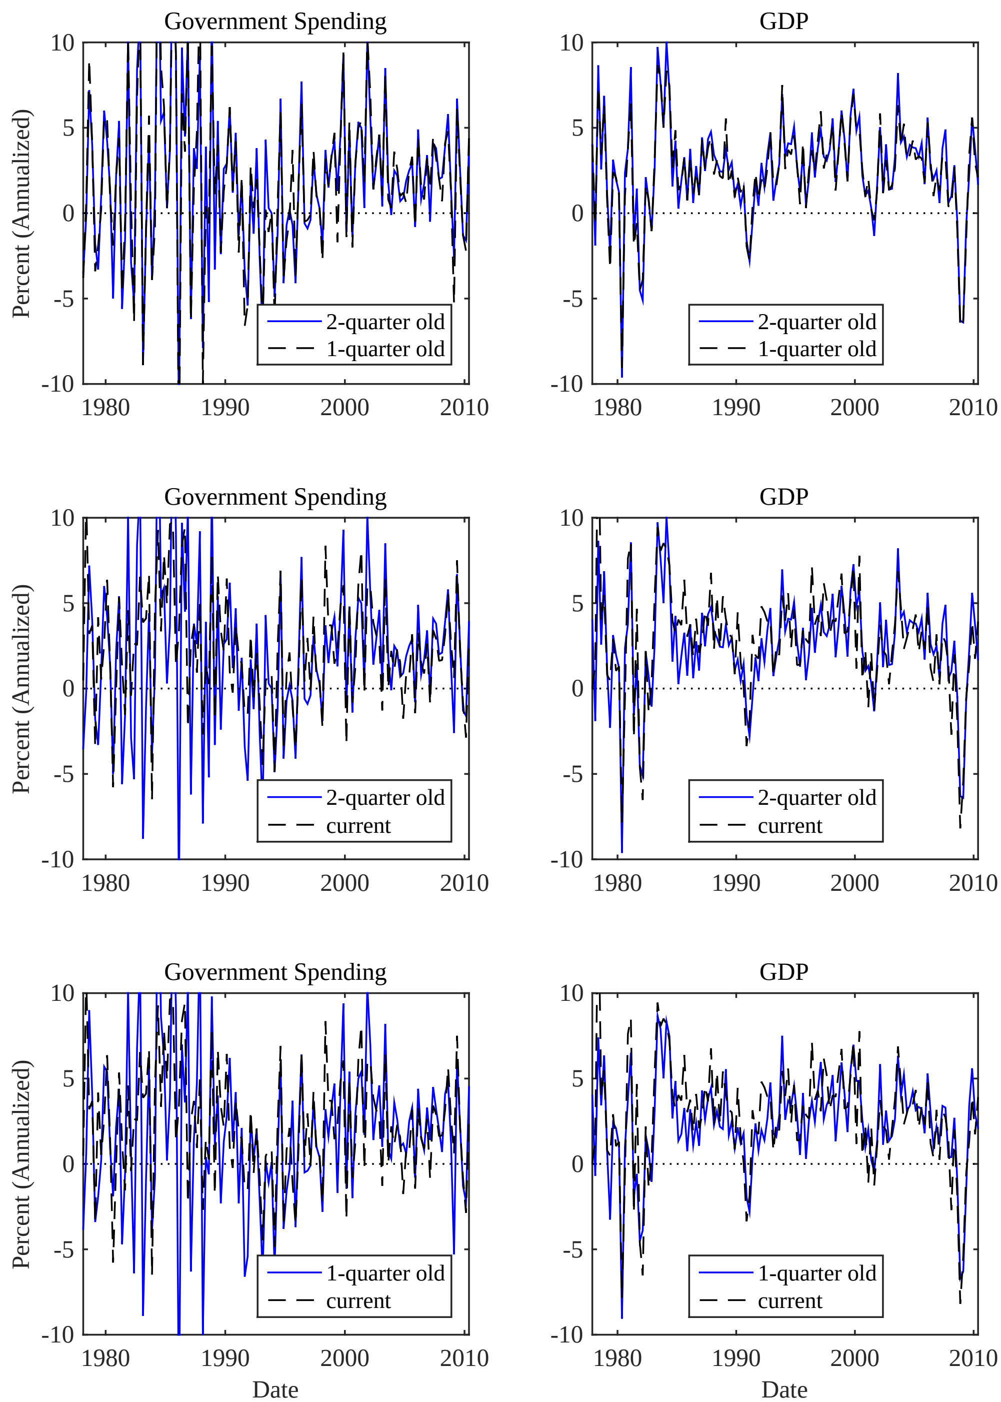

Figure 2.2: IRFs in Response to 1 Percentage Point Increase in G
SVAR Performance, 1978-2010

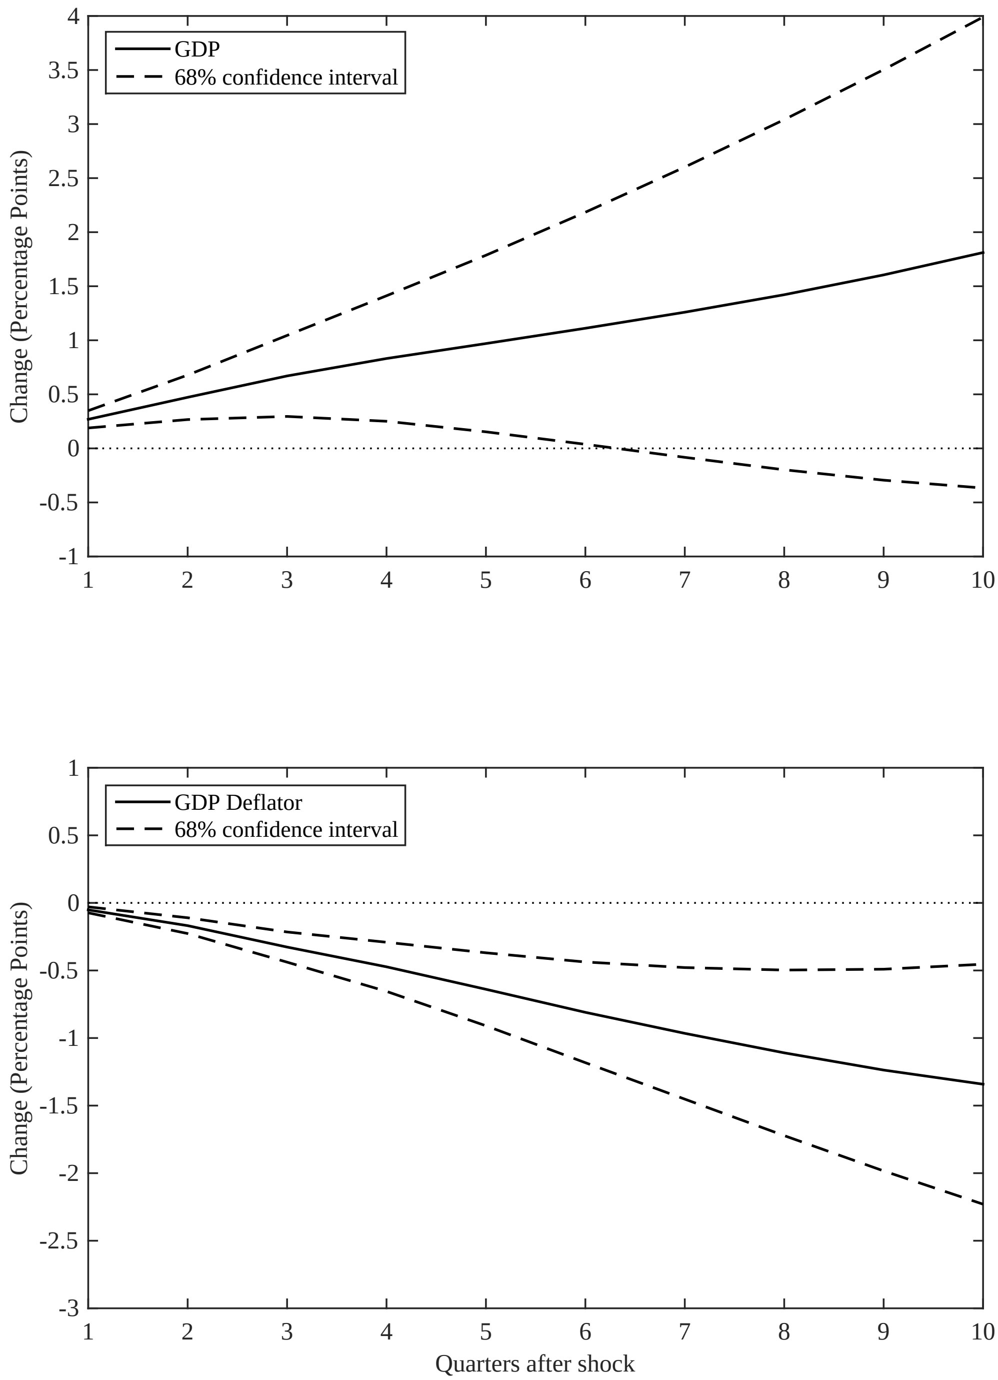

Figure 2.3: IRFs of GDP in Response to 1 Percentage Point Increase in G

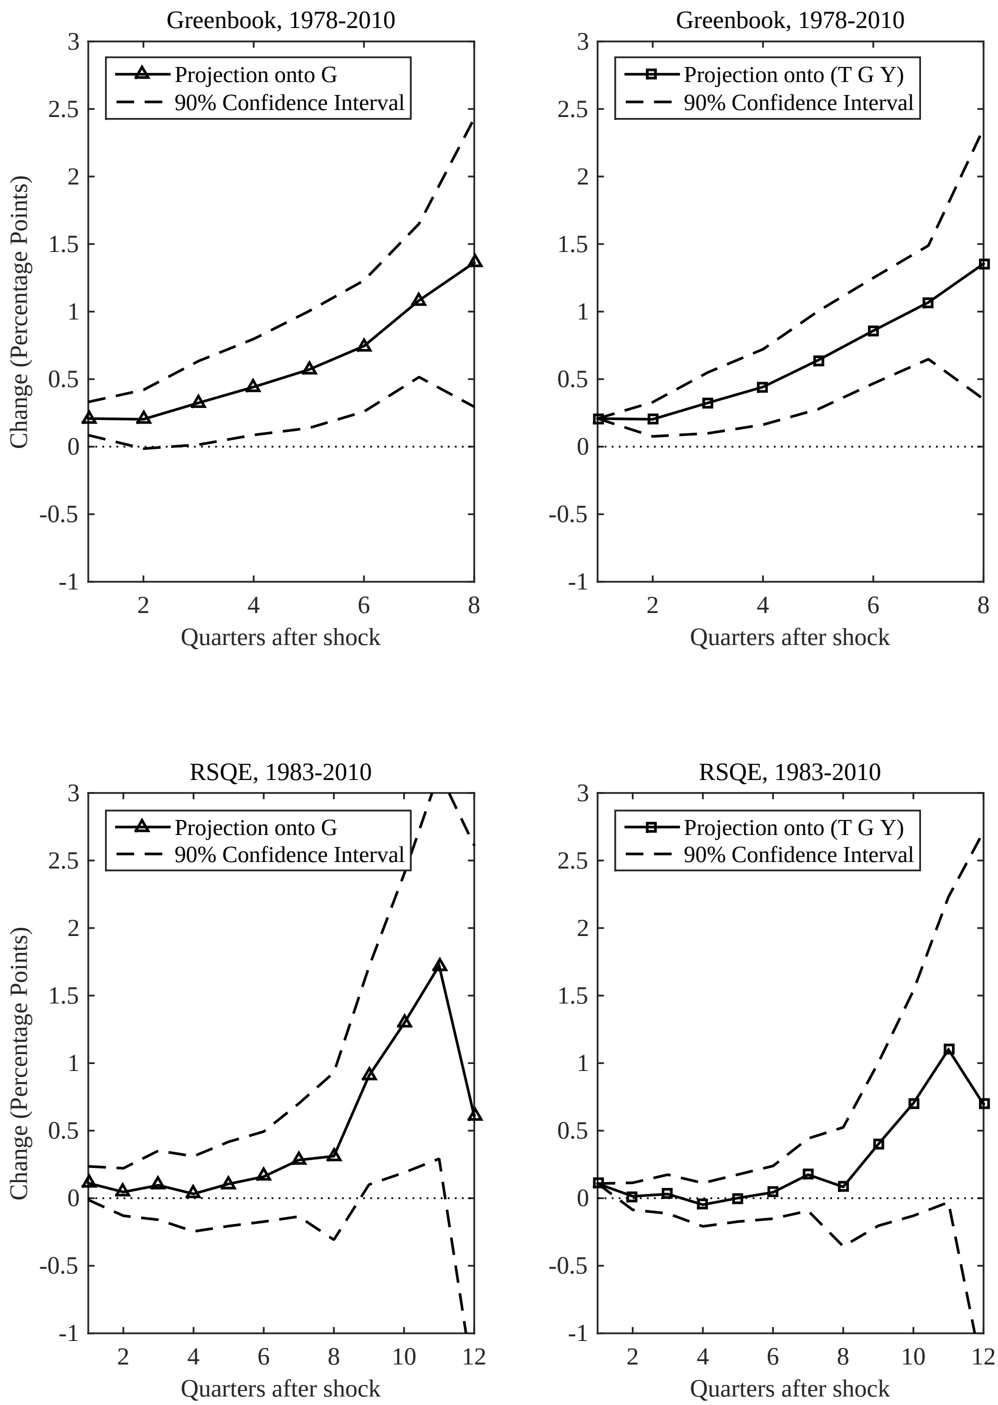

Figure 2.4: IRFs of Private GDP in Response to 1 Percentage Point Increase in G

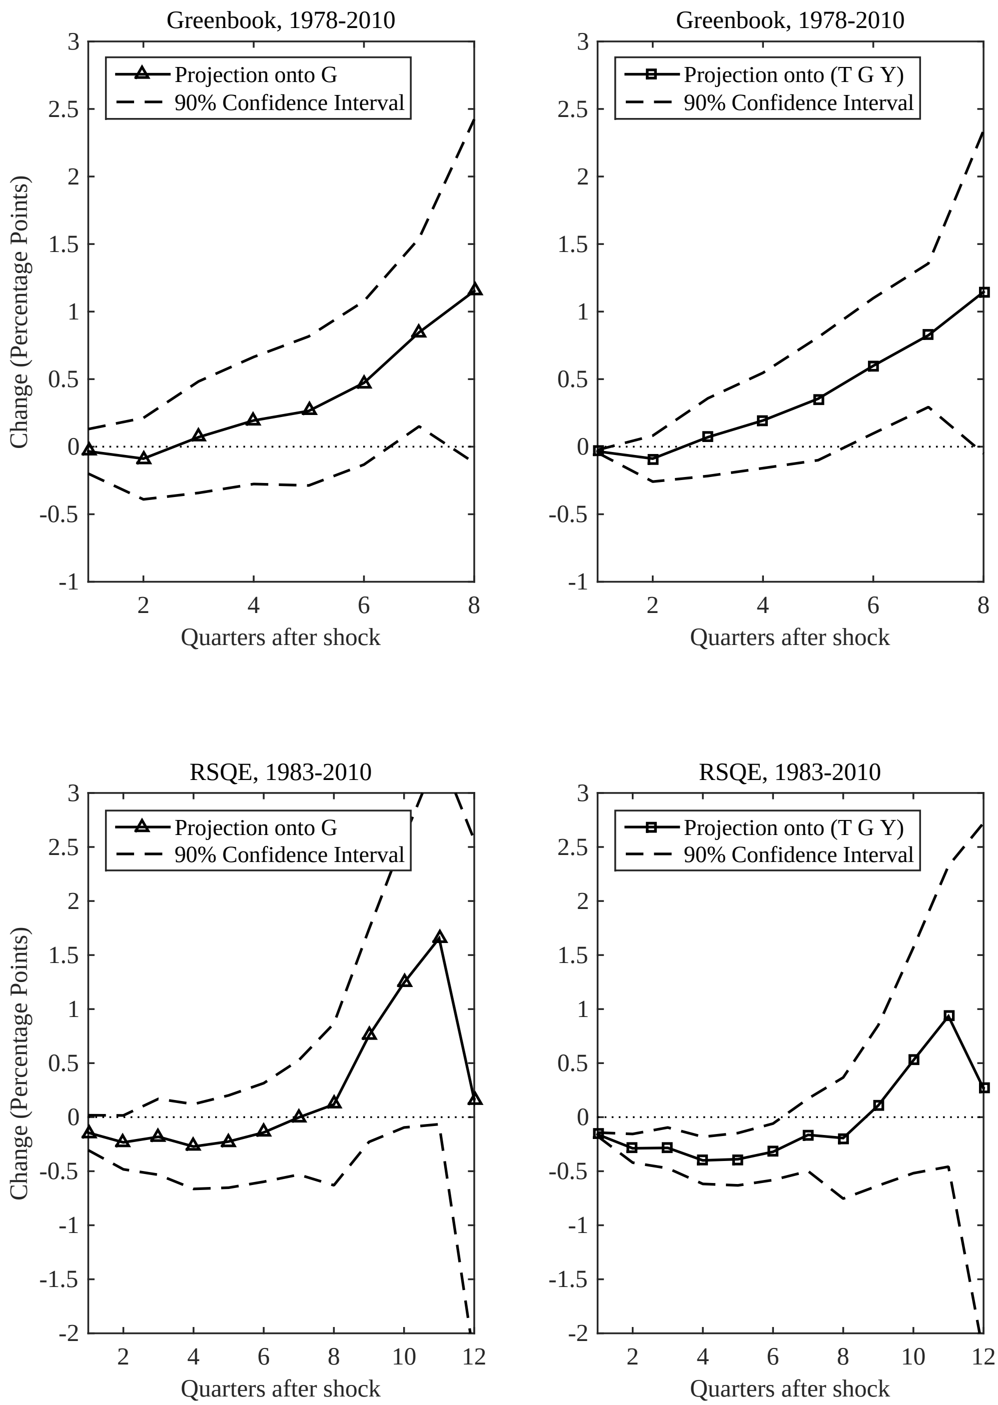

Figure 2.5: IRFs of GDP Deflator in Response to 1 Percentage Point Increase in G

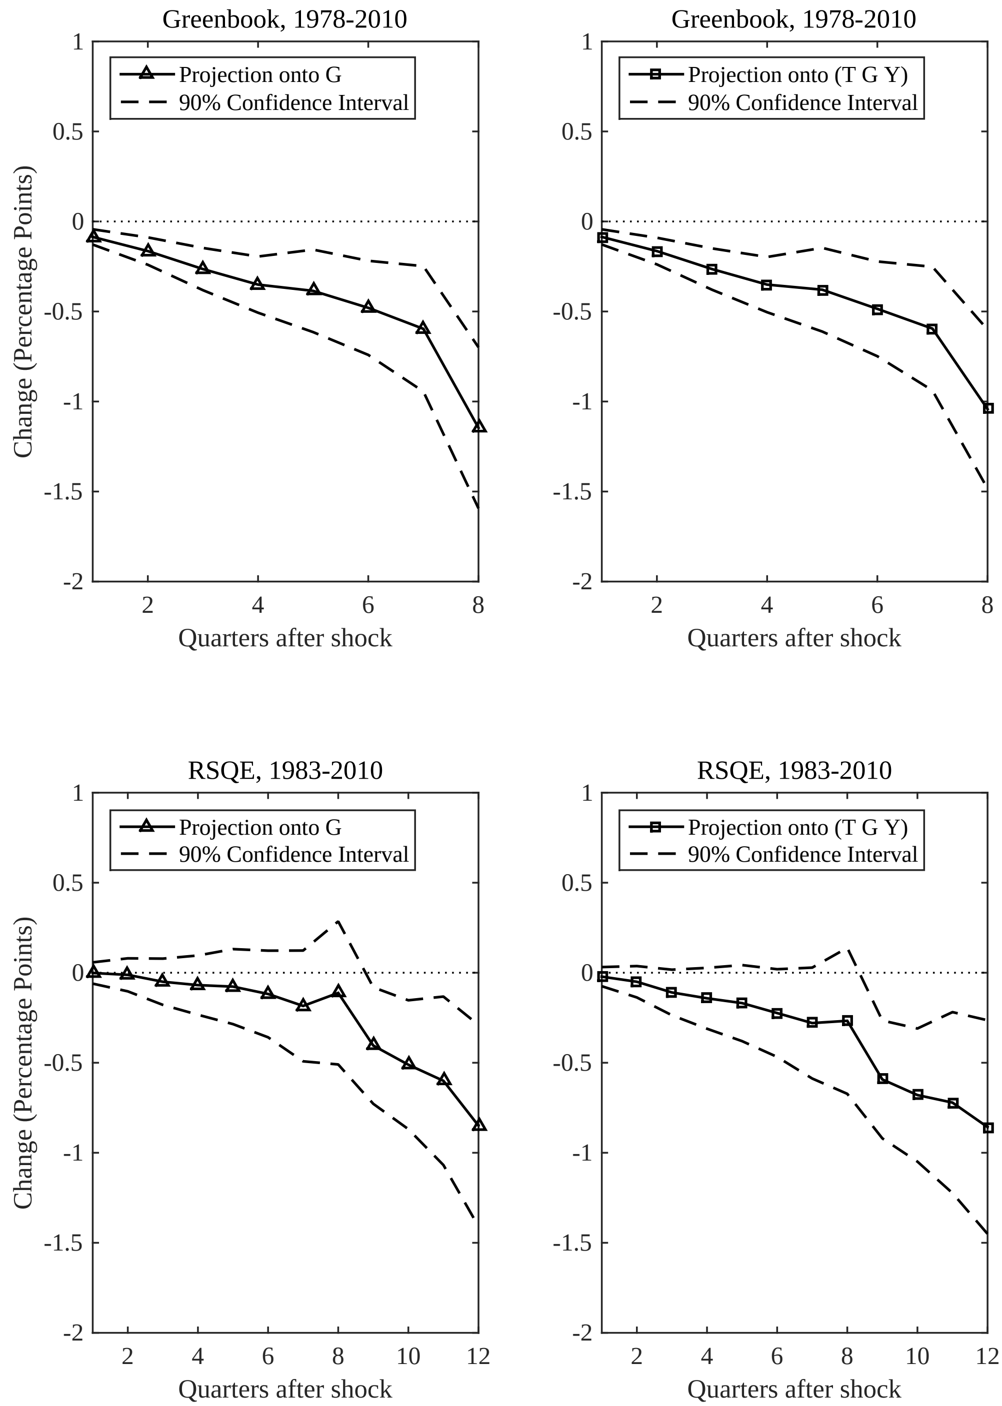

Figure 2.6: IRFs in Response to 1 Percentage Point Increase in G
Greenbook Data, 1982-2010

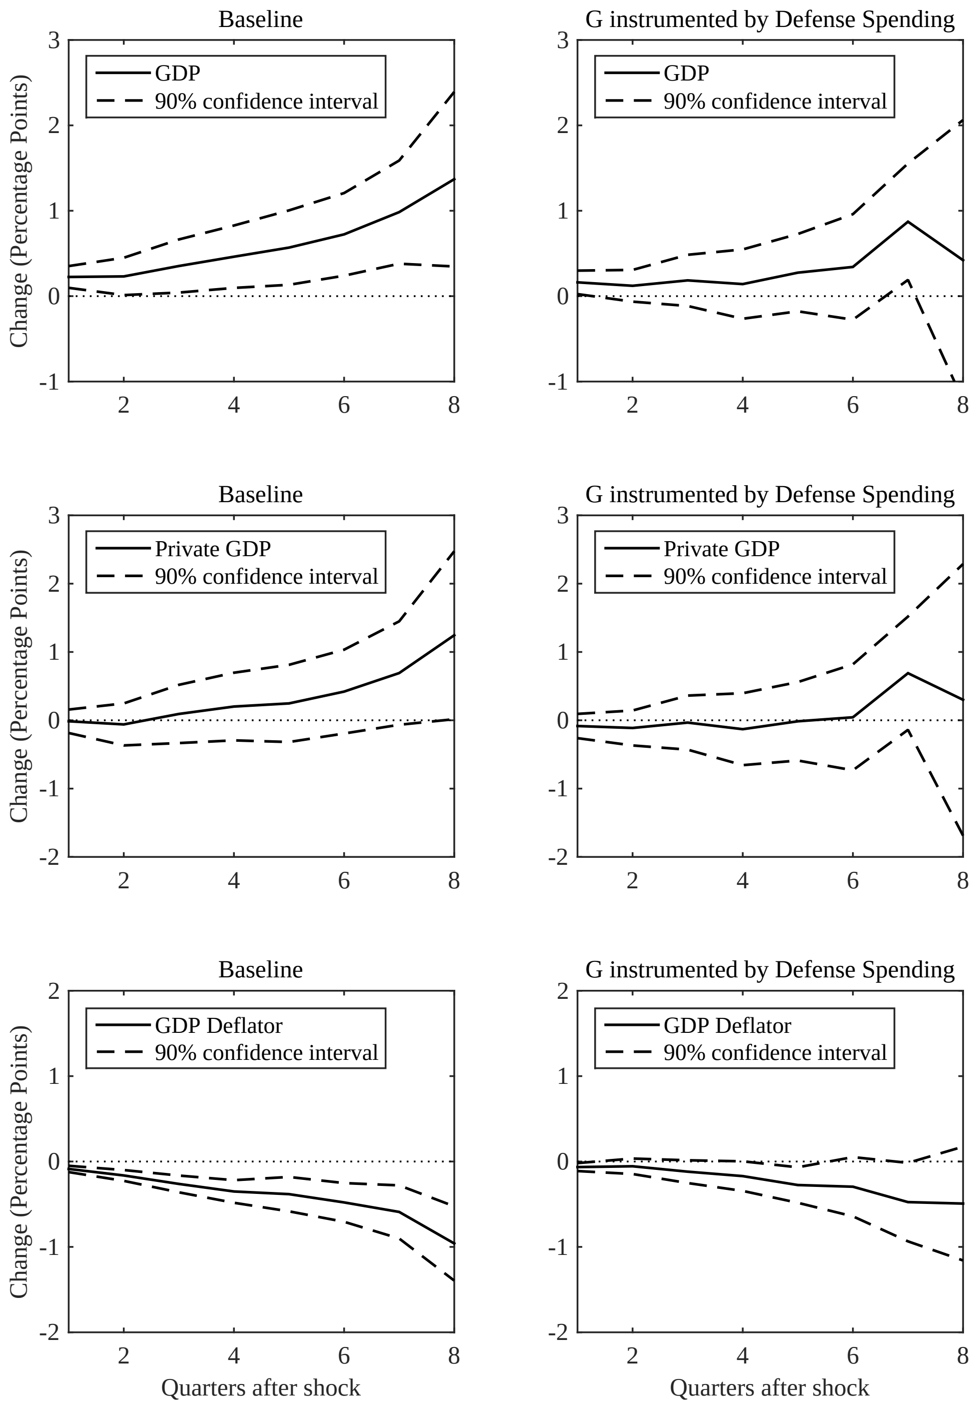

Figure 2.7: IRFs in Response to 1 Percentage Point Increase in G Greenbook, Different Sample Periods

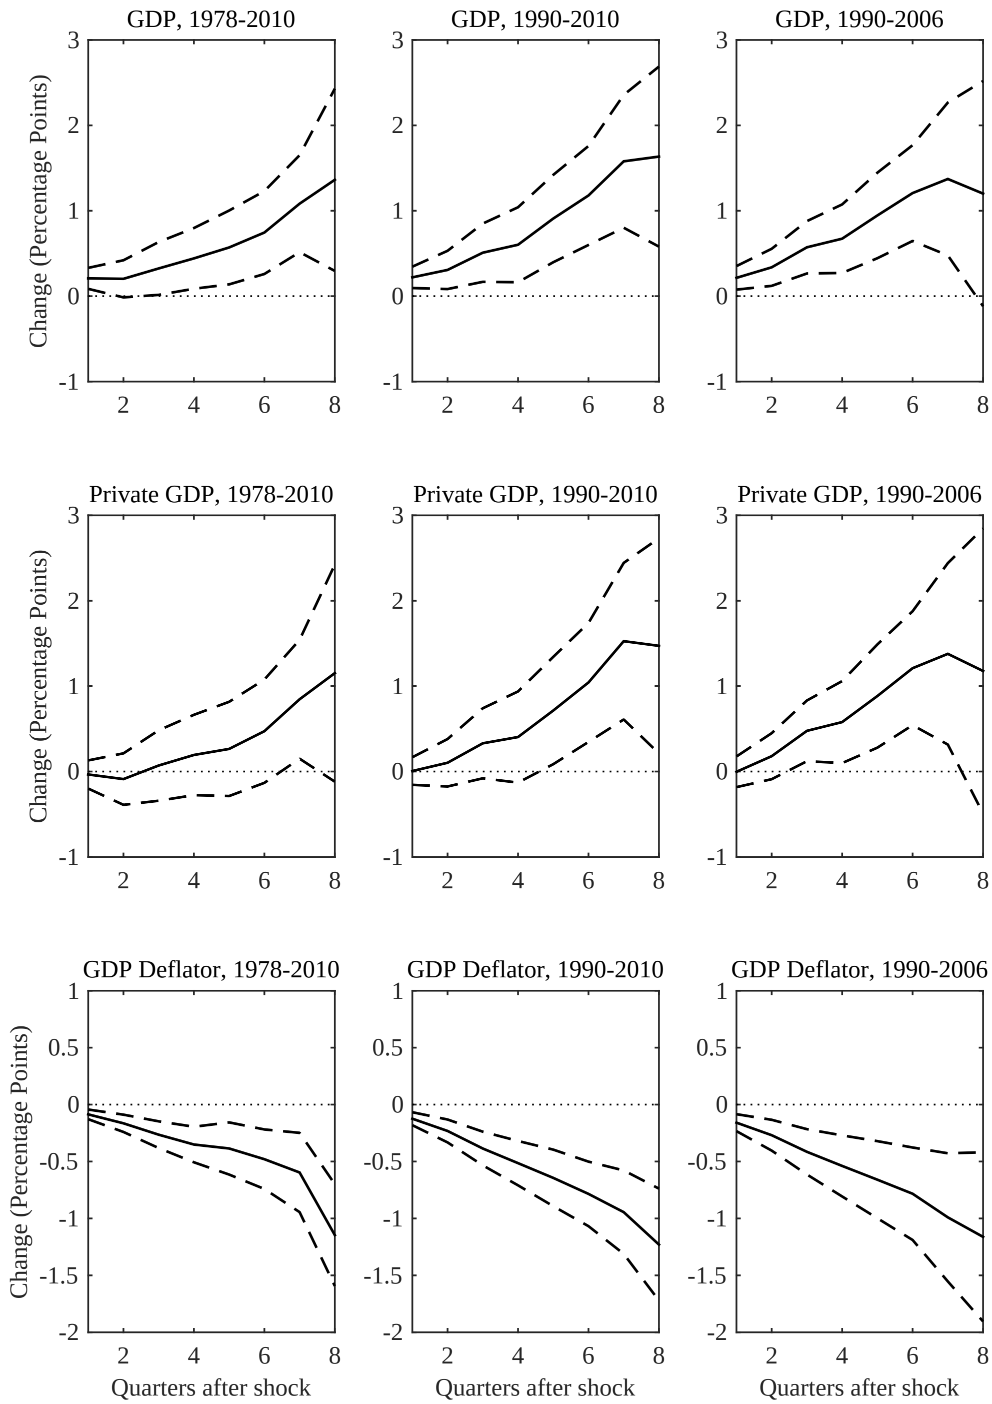

Figure 2.8: IRFs in Response to 1 Percentage Point Increase in G
Greenbook Data, 1982-2010

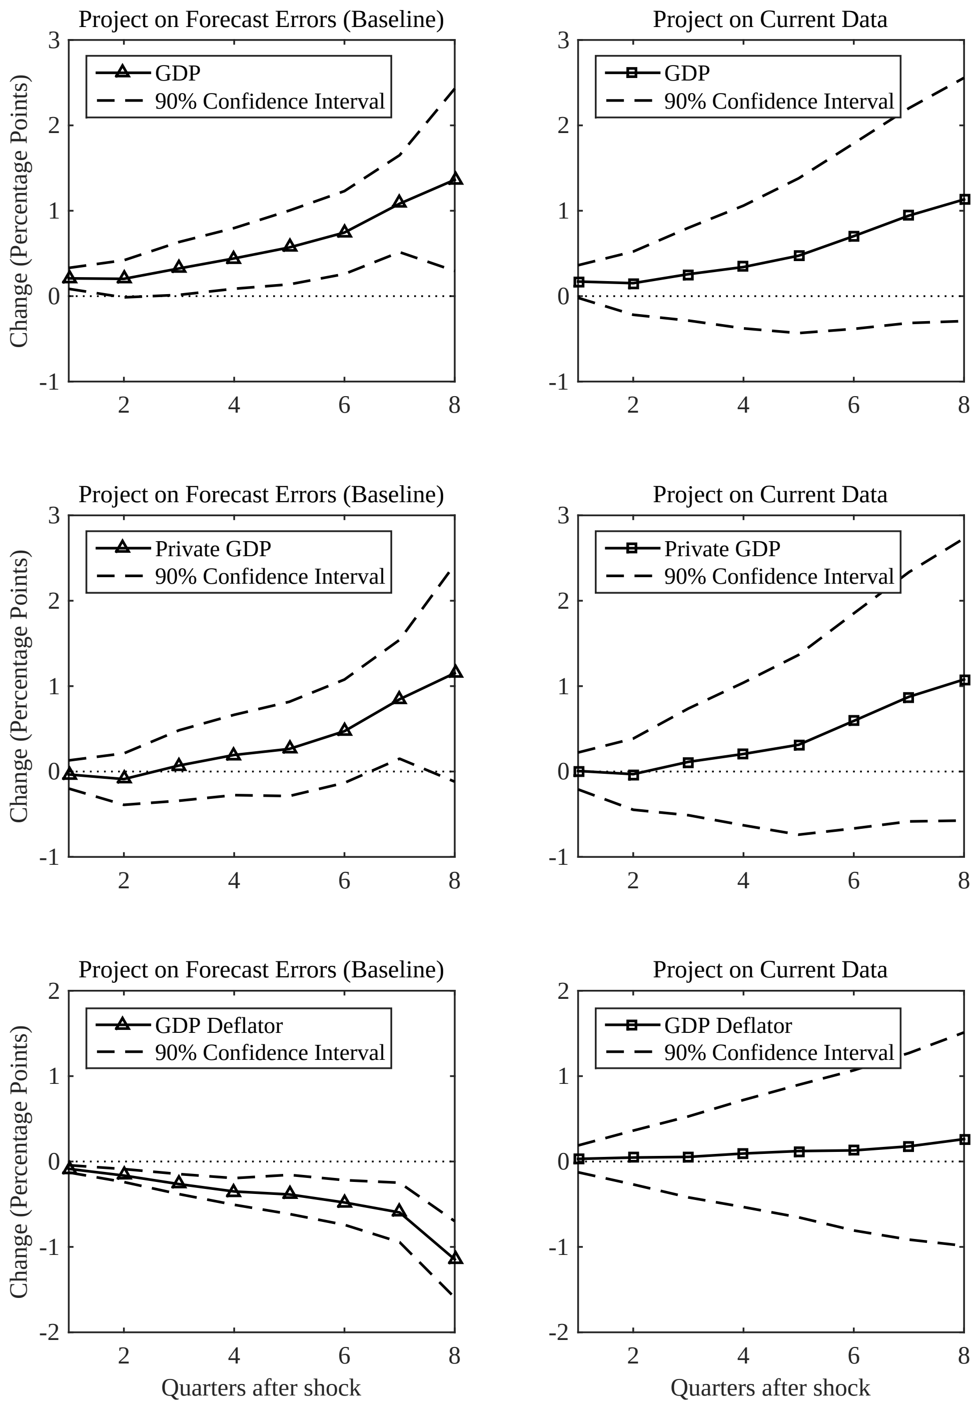

Table 2.1: Comparison of Real-Time Data for Variables of Interest
Greenbook, 1978-2010

	Real-time (2 lags)	Real-time (1 lag)	Current Vintage
GDP			
Real-time (2 lags)	1.00	0.97	0.84
Real-time (1 lag)	--	1.00	0.82
Current Vintage	--	--	1.00
G			
Real-time (2 lags)	1.00	0.96	0.62
Real-time (1 lag)	--	1.00	0.58
Current Vintage	--	--	1.00
Receipts			
Real-time (2 lags)	1.00	0.94	0.81
Real-time (1 lag)	--	1.00	0.74
Current Vintage	--	--	1.00
Private GDP			
Real-time (2 lags)	1.00	0.97	0.81
Real-time (1 lag)	--	1.00	0.79
Current Vintage	--	--	1.00
Defense Spending			
Real-time (2 lags)	1.00	0.94	0.88
Real-time (1 lag)	--	1.00	0.87
Current Vintage	--	--	1.00
Consumer Price Inflation			
Real-time (2 lags)	1.00	1.00	0.99
Real-time (1 lag)	--	1.00	0.98
Current Vintage	--	--	1.00
GDP Deflator			
Real-time (2 lags)	1.00	0.98	0.93
Real-time (1 lag)	--	1.00	0.91
Current Vintage	--	--	1.00

Note: Calculations are coefficients of correlation between series. Defense spending series begins in 1981q3. Consumer inflation series begins in 1979q3.

Table 2.2: VAR Specifications

VAR Specification	Current Vintage or Real-Time Data	Rolling Estimation or Static	Use Current-Period or Forecast Data
VAR1	Current	Static	Curent Period
VAR2	Current	Rolling	Forecast
VAR3	Real-Time	Static	Current Period
VAR4	Real-Time	Rolling	Forecast

Notes: current-period data refers to data corresponding to the period in which the forecast is being made (for future periods). The use of forecast data, then, refers to specifications where the VAR does not have access to current-period data (which would not be available to a forecaster) when estimating future period outcomes.

Table 2.3: Forecast Error Statistics
by Forecaster, 1983-2010 and 1990-2010

	1983 - 2010										1990-2010									
	GDP								G	T	GDP								G	T
	fe^Y_{t+1}	fe^Y_{t+2}	fe^Y_{t+3}	fe^Y_{t+4}	fe^Y_{t+5}	fe^Y_{t+6}	fe^Y_{t+7}	fe^Y_{t+8}	fe^G_{t+1}	fe^T_{t+1}	fe^Y_{t+1}	fe^Y_{t+2}	fe^Y_{t+3}	fe^Y_{t+4}	fe^Y_{t+5}	fe^Y_{t+6}	fe^Y_{t+7}	fe^Y_{t+8}	fe^G_{t+1}	fe^T_{t+1}
Mean Error (percentage points)																				
GB	0.13	-0.04	-0.12	-0.17	-0.31	-0.40	-0.29	-0.31	0.59	0.46	0.07	-0.12	-0.19	-0.24	-0.25	-0.33	-0.18	-0.25	0.08	-0.09
RSQE	0.29	-0.04	-0.11	-0.09	-0.11	-0.16	-0.22	-0.38	0.95	0.06	0.06	-0.15	-0.23	-0.08	-0.03	-0.15	-0.29	-0.41	0.55	0.17
Average*	0.16	-0.06	-0.16	-0.19	-0.21	-0.15	-0.20	-0.26	0.76	0.20	0.19	0.00	-0.08	-0.08	-0.14	-0.13	-0.14	-0.36	0.25	0.07
VAR1	-0.11	-0.19	-0.24	-0.25	-0.27	-0.28	-0.30	-0.30	-0.04	-0.75	-0.42	-0.41	-0.39	-0.28	-0.23	-0.27	-0.32	-0.34	-0.33	-1.62
VAR2	-0.42	-0.60	-0.75	-0.80	-0.86	-0.91	-0.95	-0.97	-0.19	-0.87	-0.74	-0.91	-1.04	-0.96	-0.94	-0.95	-0.98	-0.98	-0.46	-1.56
VAR3	-0.07	-0.16	-0.22	-0.17	-0.18	-0.21	-0.19	-0.22	-0.19	-1.41	-0.28	-0.31	-0.33	-0.21	-0.08	-0.02	0.00	0.00	-0.36	-2.35
VAR4	0.01	-0.08	-0.11	-0.09	-0.14	-0.17	-0.12	-0.13	0.02	-1.14	-0.15	-0.15	-0.13	-0.07	0.00	0.00	0.00	-0.01	-0.21	-2.06
Mean Squared Error (percentage points)																				
GB	2.19	2.30	2.37	2.38	2.33	2.22	2.43	2.90	4.19	10.43	2.17	2.40	2.53	2.55	2.46	2.33	2.53	2.93	3.06	10.52
RSQE	2.41	2.54	2.65	2.66	2.60	2.62	2.68	2.77	3.56	10.88	2.50	2.79	2.90	2.85	2.77	2.72	2.79	2.83	3.05	11.43
Average*	2.05	2.24	2.33	2.34	2.36	2.29	2.49	2.85	3.51	9.32	2.15	2.37	2.50	2.55	2.53	2.42	2.59	2.97	2.79	9.76
VAR1	2.58	2.43	2.40	2.39	2.38	2.38	2.39	2.39	3.41	13.68	2.61	2.61	2.59	2.50	2.47	2.46	2.47	2.49	3.16	13.92
VAR2	2.71	2.77	2.82	2.66	2.61	2.59	2.60	2.63	3.50	14.35	2.82	2.98	3.02	2.79	2.66	2.68	2.69	2.71	3.26	14.86
VAR3	2.81	2.69	2.72	2.45	2.29	2.25	2.25	2.25	5.31	16.34	2.74	2.78	2.80	2.53	2.34	2.33	2.33	2.32	3.70	16.05
VAR4	2.58	2.44	2.42	2.42	2.36	2.34	2.30	2.30	4.98	15.05	2.52	2.52	2.52	2.50	2.44	2.43	2.43	2.43	3.55	14.66
Mean Absolute Error (percentage points)																				
GB	1.73	1.78	1.83	1.78	1.76	1.67	1.84	2.22	3.03	6.85	1.74	1.89	2.00	1.93	1.89	1.77	1.91	2.23	2.36	6.78
RSQE	1.91	1.87	1.93	1.97	1.84	1.86	1.95	2.02	2.77	7.94	2.00	2.08	2.11	2.09	1.96	1.92	2.02	2.03	2.44	8.37
Average*	1.64	1.69	1.75	1.73	1.71	1.67	1.80	2.01	2.64	6.66	1.72	1.80	1.88	1.89	1.85	1.78	1.85	2.06	2.15	7.04
VAR1	1.89	1.81	1.78	1.77	1.76	1.76	1.77	1.77	2.69	9.90	1.94	1.93	1.91	1.85	1.82	1.80	1.82	1.84	2.45	9.95
VAR2	1.99	2.00	2.03	1.86	1.83	1.81	1.83	1.86	2.75	10.59	2.12	2.14	2.16	1.96	1.87	1.89	1.90	1.91	2.54	10.89
VAR3	2.11	2.01	2.09	1.83	1.67	1.65	1.66	1.65	3.77	11.99	2.05	2.07	2.13	1.89	1.71	1.70	1.70	1.69	2.93	11.38
VAR4	1.88	1.81	1.78	1.79	1.74	1.72	1.69	1.70	3.64	11.36	1.87	1.87	1.87	1.84	1.79	1.78	1.78	1.78	2.79	10.89

Notes: Average forecast error is the unweighted arithmetic mean of Greenbook and RSQE forecast errors. VAR1 is a 4 period lag on current vintage data, estimated over entire sample period. VAR2 is a 4 period lag on current vintage data, estimated on a rolling basis without access to the current period's data. VAR3 is a 4 period lag on real-time data, estimated over the entire sample period. VAR4 is a 4 period lag on real-time data, estimated on a rolling basis without access to the current period's real-time data.

Table 2.4: Cumulative Forecast Error Statistics
by Forecaster, 1983-2010 and 1990-2010

	1983 - 2010									
	GDP								G	T
	fe^Y_{t+1}	fe^Y_{t+2}	fe^Y_{t+3}	fe^Y_{t+4}	fe^Y_{t+5}	fe^Y_{t+6}	fe^Y_{t+7}	fe^Y_{t+8}	fe^G_{t+1}	fe^T_{t+1}
Mean Error (percentage points)										
GB	0.07	0.11	0.12	0.12	0.03	-0.02	-0.06	-0.46	0.17	-0.51
RSQE	0.07	0.05	0.02	0.00	-0.03	-0.07	-0.19	-0.31	0.23	-0.08
Average*	0.06	0.06	0.04	0.01	-0.06	-0.11	-0.18	-0.44	0.19	-0.30
VAR1	-0.03	-0.09	-0.15	-0.22	-0.29	-0.37	-0.45	-0.53	-0.02	-0.38
VAR2	-0.11	-0.27	-0.46	-0.66	-0.87	-1.10	-1.34	-1.58	-0.06	-0.43
VAR3	0.00	-0.03	-0.06	-0.09	-0.13	-0.18	-0.21	-0.25	-0.02	-0.51
VAR4	-0.03	-0.07	-0.13	-0.17	-0.22	-0.27	-0.32	-0.38	-0.07	-0.63
Mean Squared Error (percentage points)										
GB	0.57	1.01	1.46	1.86	2.28	2.63	2.99	3.52	0.88	2.83
RSQE	0.59	1.03	1.49	1.90	2.27	2.62	3.07	3.59	0.88	2.79
Average*	0.56	0.99	1.43	1.82	2.25	2.59	2.92	3.47	0.82	2.49
VAR1	0.65	1.01	1.29	1.48	1.63	1.74	1.87	2.00	0.86	3.63
VAR2	0.68	1.18	1.68	2.10	2.47	2.79	3.09	3.40	0.88	3.84
VAR3	0.65	1.01	1.21	1.41	1.54	1.61	1.71	1.78	1.24	3.94
VAR4	0.70	1.21	1.64	2.01	2.29	2.50	2.69	2.84	1.33	4.37
Mean Absolute Error (percentage points)										
GB	0.45	0.79	1.16	1.49	1.86	2.16	2.48	2.81	0.69	1.71
RSQE	0.47	0.80	1.15	1.44	1.72	1.96	2.35	2.81	0.68	1.98
Average*	0.44	0.76	1.11	1.41	1.78	2.05	2.34	2.75	0.63	1.67
VAR1	0.47	0.71	0.93	1.06	1.17	1.24	1.34	1.45	0.67	2.53
VAR2	0.50	0.83	1.21	1.53	1.79	2.00	2.23	2.48	0.69	2.73
VAR3	0.47	0.72	0.90	1.01	1.08	1.15	1.28	1.35	0.91	2.91
VAR4	0.53	0.90	1.25	1.50	1.72	1.89	2.05	2.23	0.94	3.10

	1990-2010									
	GDP								G	T
	fe^Y_{t+1}	fe^Y_{t+2}	fe^Y_{t+3}	fe^Y_{t+4}	fe^Y_{t+5}	fe^Y_{t+6}	fe^Y_{t+7}	fe^Y_{t+8}	fe^G_{t+1}	fe^T_{t+1}
Mean Error (percentage points)										
GB	0.00	-0.06	-0.13	-0.21	-0.29	-0.34	-0.31	-0.40	0.00	-0.53
RSQE	0.01	-0.03	-0.09	-0.11	-0.12	-0.16	-0.26	-0.35	0.13	-0.06
Average*	0.03	0.02	-0.02	-0.06	-0.11	-0.13	-0.15	-0.36	0.06	-0.27
VAR1	-0.11	-0.22	-0.32	-0.40	-0.46	-0.53	-0.62	-0.71	-0.09	-0.61
VAR2	-0.19	-0.43	-0.69	-0.93	-1.17	-1.40	-1.65	-1.89	-0.13	-0.63
VAR3	-0.04	-0.08	-0.12	-0.14	-0.15	-0.15	-0.16	-0.17	-0.06	-0.74
VAR4	-0.08	-0.16	-0.25	-0.30	-0.32	-0.33	-0.33	-0.33	-0.10	-0.87
Mean Squared Error (percentage points)										
GB	0.58	1.04	1.50	1.94	2.37	2.75	3.11	3.57	0.75	2.89
RSQE	0.62	1.12	1.63	2.09	2.51	2.90	3.34	3.79	0.76	2.97
Average*	0.59	1.04	1.51	1.97	2.41	2.79	3.14	3.62	0.71	2.63
VAR1	0.66	1.05	1.35	1.54	1.68	1.81	1.94	2.07	0.79	3.79
VAR2	0.71	1.25	1.78	2.22	2.59	2.92	3.23	3.53	0.82	4.08
VAR3	0.64	1.04	1.26	1.43	1.54	1.63	1.71	1.79	0.89	3.95
VAR4	0.69	1.23	1.70	2.06	2.31	2.52	2.68	2.83	0.93	4.44
Mean Absolute Error (percentage points)										
GB	0.46	0.82	1.21	1.56	1.93	2.26	2.60	2.86	0.59	1.68
RSQE	0.49	0.88	1.29	1.62	1.96	2.23	2.60	2.99	0.60	2.12
Average*	0.47	0.81	1.19	1.54	1.91	2.23	2.57	2.91	0.56	1.74
VAR1	0.49	0.75	0.98	1.12	1.24	1.31	1.41	1.49	0.61	2.59
VAR2	0.54	0.90	1.29	1.64	1.90	2.14	2.36	2.63	0.64	2.84
VAR3	0.47	0.74	0.93	1.03	1.08	1.15	1.26	1.33	0.70	2.83
VAR4	0.51	0.91	1.27	1.52	1.72	1.87	2.00	2.17	0.73	3.00

Notes: Average forecast error is the unweighted arithmetic mean of Greenbook and RSQE forecast errors.

Table 2.5: Fiscal Multiplier of G on GDP
Greenbook Data, 1978-2010

Number of Quarters	1	2	3	4	5	6	7	8
Project onto G shock								
GDP	0.924	0.856	0.998	1.158	1.281	1.445	1.673	1.896
Private GDP	-0.119	-0.198	-0.057	0.105	0.227	0.389	0.618	0.847
Project onto (T, G, Y)' shock								
GDP	0.924	0.856	0.998	1.158	1.332	1.539	1.739	1.949
Private GDP	-0.119	-0.198	-0.057	0.105	0.277	0.482	0.683	0.899
sample size	132	132	132	132	120	89	57	25

Note: Estimates are obtained using net present value calculation procedure described in Mountford and Uhlig (2009) using r=.04. Cyclically-adjusted real federal gross receipts are ordered first.

Table 2.6: Fiscal Multiplier of G on GDP
RSQE Data, 1983-2010

Number of Quarters	1	2	3	4	5	6	7	8	9	10	11	12
Project onto G shock												
GDP	0.503	0.352	0.367	0.304	0.317	0.363	0.456	0.550	0.888	1.278	1.679	1.797
Private GDP	-0.504	-0.653	-0.627	-0.682	-0.664	-0.614	-0.517	-0.415	-0.132	0.251	0.648	0.681
Project onto (T, G, Y)' shock												
GDP	0.468	0.261	0.211	0.107	0.082	0.093	0.165	0.183	0.342	0.594	0.928	1.110
Private GDP	-0.537	-0.740	-0.778	-0.872	-0.892	-0.876	-0.800	-0.773	-0.729	-0.530	-0.221	-0.143
sample size	131	130	129	128	127	126	114	102	87	76	57	38

Note: Estimates are obtained using net present value calculation procedure described in Mountford and Uhlig (2009) using r=.04. Cyclically-adjusted real government net receipts are ordered first.

Table 2.7: Fiscal Multiplier of G on GDP, Instrument G with Defense Spending Greenbook Data 1982-2010

Number of Quarters	1	2	3	4	5	6	7	8
Project onto G shock								
GDP	0.727	0.668	0.748	0.714	0.786	0.866	1.206	1.317
Private GDP	-0.293	-0.360	-0.287	-0.328	-0.261	-0.185	0.150	0.257
Project onto (T, G, Y)' shock								
GDP	0.752	0.708	0.798	0.763	0.831	0.910	1.144	1.238
Private GDP	-0.269	-0.323	-0.239	-0.281	-0.218	-0.142	0.088	0.178
sample size	117	117	117	117	109	83	54	25

Note: Estimates are obtained using net present value calculation procedure described in Mountford and Uhlig (2009) using r=.04. Cyclically-adjusted real federal gross receipts are ordered first.

Appendix

2.7 Tax Elasticity of Output

The within-period elasticity of taxes with respect to output $\eta_{\tau,y}$ is estimated in a similar manner to Blanchard and Perotti (2002). $\eta_{\tau,y}$ is an average of various tax elasticities weighted by their respected shares of total taxes, $\sigma_{\tau_i,\tau}$, as follows:

$$\eta_{\tau,y} = \sum_i \eta_{\tau_i,B_i}\eta_{B_i,Y}\sigma_{\tau_i,\tau}. \tag{2.13}$$

In equation 2.13 above, the elasticity of each tax is broken into two separate components: 1) the elasticity of tax i to its tax base B_i, and 2) the elasticity of the tax base to output. With respect to relatively proportional taxes, such as payroll taxes, η_{τ_i,B_i} is closer to unity, while for progressive taxes, such as the federal income tax, η_{τ_i,B_i} is a larger number.

In constructing these various elasticities, Blanchard and Perotti (2002) rely on estimates provided by Giorno et al (1995) and supply their own estimates based on a series of regressions of log changes of the tax base on leads and lags of log changes to output. The contemporaneous effect is treated as the within-period elasticity. We create a similar series of estimates below for each measure of taxation used in our analysis.

2.7.1 Greenbook

Using BEA data of personal income, output, and various tax receipts, supplemented by CPS measures of employment and average weekly earnings among the employed, we directly estimate elasticity measures for three federal taxes that make up more than 90 percent of gross federal receipts: the personal income tax (TPF), the corporate income tax (TCF), and social insurance taxes (TSIF).

2.7.1.1 Social Insurance Taxes

The relationship between social insurance taxes and output is characterized by Giorno et al (1995) as $T = t(w)w(E)E(Y)$, where $t(w)$ is the tax rate on wages, $w(E)$ is earnings, and $E(Y)$ is the employment level. With a little bit of rearrangement, the elasticity of federal social insurance taxes $\eta_{TSIF,y}$ can then be separated into three parts: 1) the elasticity of employment with respect to output $\eta_{E,y}$, 2) the elasticity of weekly average wages among the employed with respect to employment $\eta_{w,E}$, and 3) the elasticity of social insurance taxes with respect to wages $\eta_{wt,w}$.

$$\eta_{TSIF,y} = \eta_{E,y}\left[\eta_{w,E}\eta_{wt,w} + 1\right] \tag{2.14}$$

We use the quarterly average of seasonally-adjusted privately employed production and nonsupervisory employees as our measure of E. For wages w, we use the average weekly earnings of privately employed production and nonsupervisory employees. Regressing the log change of each variable on one lead and four lags of the log change in output and employment, respectively, we obtain within-period elasticity estimates of $\eta_{E,y} = .33$ and $\eta_{w,E} = .66$ when we restrict the sample to 1978-2010 (the period for which we have GB forecasts of federal receipts). These estimates are close to those of Blanchard and Perotti (2002), who find $\eta_{E,y} = .42$ and $\eta_{w,E} = .62$. When we change our sample to 1964-1997, the last year included in Blanchard and Perotti's analysis, our estimates are .36 and .68, respectively. Using different measures of earnings and wages slightly alters our results, but the product of the two elasticities is surprisingly robust to our choice of wage and employment pairing.

The elasticity of taxes with respect to earnings, $\eta_{wt,w}$, is taken to be 1.0 in Giorno et al and Blanchard and Perotti. We adopt instead the convention that the elasticity is equal to the share of covered earnings, which changes slowly over time. Thus, we use $\eta_{wt,w} = 0.85$. Combining the three measures, we get $\eta_{TSIF,y}$ slightly greater than 0.5.

2.7.1.2 Federal Personal Income Taxes

The elasticity of the federal personal income tax $\eta_{TPF,y}$ is calculated in a departure from Blanchard and Perotti. They utilize the elasticities calculated above and combine them with an estimate of $\eta_{wt,w}$ from Giorno et al, who find that in the United States the gross-earnings elasticity of income tax rises from 2.5 in 1978 to 3.9 in 1992. As this is an annual number, and the earnings base for the federal income tax is substantially different from earned income, we instead repeat the analysis above, but with personal income directly. We define $TPF(B(Y))$ to be the federal income tax TPF as a function of the income base B, which is a function of output Y. The elasticity $\eta_{TPF,y}$, then, is

$$\eta_{TPF,y} = \eta_{TPF,B}\eta_{B,y}. \tag{2.15}$$

We estimate $\eta_{TPF,B}$ using the same lead and lag formulation as for social insurance taxes, where B is aggregate personal income, and obtain an estimate of 2.34. The elasticity $\eta_{B,y}$ is 0.48, yielding $\eta_{TPF,y} = 1.12$.

2.7.1.3 Federal Corporate Income Taxes

Our corporate tax analysis is performed identically to the personal income tax. Given

$$\eta_{TCF,y} = \eta_{TCF,B}\eta_{B,y}, . \tag{2.16}$$

we estimate $\eta_{TCF,B} = 1.0$ and $\eta_{B,y} = 3.69$, yielding $\eta_{TCF,y} = 3.7$.

2.7.1.4 Other Federal Taxes and Aggregate Elasticity

Our measure of gross federal receipts in Federal Reserve Greenbook forecasts includes other federal taxes that are not included above, but which constitute on average almost 10 percent of revenues. We assume a within-period elasticity of 1.0 for these other taxes with respect to output.

The weighted average within-period elasticity of federal taxes with respect to GDP, $\eta_{\tau,y}$, is estimated to be about 1.18.

2.7.2 RSQE

The primary differences between our measure of receipts in RSQE forecasts and our Greenbook measure are 1) the presence of state and local revenues in RSQE forecasts, and 2) RSQE receipts are net receipts, so they include transfers such as social security payments. We define net receipts as follows:

$$NR_t = GFR_t + GSLR_t - GTRF_t - GTRSL_t, \tag{2.17}$$

where GFR_t = gross federal receipts (personal income taxes, corporate income taxes, indirect business taxes, social insurance taxes, household transfers to the federal government, and business transfers to the federal government);[7] $GSLR_t$ = gross state receipts (personal income taxes, corporate income taxes, indirect business taxes, social insurance taxes, and federal aid to states); $GTRF_t$ = federal transfers (government transfers to persons and federal aid to states); and $GTRSL_t$ = state and local transfers.

In addition, our RSQE forecasts range from 1983-2015 rather than 1978-2015, so some minor differences arise from this change of sample. The three federal taxes above, as well as state and local personal income taxes (TPSL), are calculated in the same manner as for the Greenbooks, and they

[7] Before 2004, RSQE did not itemize forecasts of household and business transfers to the federal government.

yield the following estimated within-period elasticities as follows:

$$\begin{aligned}\eta_{TPF,y} =& 0.45\\ \eta_{TSIF,y} =& 0.40\\ \eta_{TCF,y} =& 3.63\\ \eta_{TPSL,y} =& 0.21\end{aligned}$$

2.7.2.1 Indirect State and Local Business Taxes

Indirect business taxes mostly consist of sales taxes collected from retailers, but paid to businesses by consumers as a proportional tax on some subset of final goods. Blanchard and Perotti use an elasticity of 1.0 for this category, while noting that some goods are exempt from sales taxes. We estimate the elasticity using the BEA measure of indirect business taxes collected by state and local governments (TIBSL), and find $\eta_{TIBSL,y} = .61$.

2.7.2.2 Transfers and Aggregate Elasticity

Transfers enter into the average elasticity equation as a negative share, and their respective within-period elasticities span a large range. Social security benefits, which are a large portion of transfers, likely have little-to-no relationship with contemporaneous changes to output, while unemployment benefits react quite strongly (and inversely) to GDP. We use -0.2 as our elasticity of total transfers to GDP, following Blanchard and Perotti in using OECD estimates. Given net receipts (NR) is equal to gross receipts (GR) less transfers (TR),

$$\eta_{\tau,y} = \eta_{NR,y} = \eta_{GR,y}\sigma_{GR,NR} - \eta_{TR,y}\sigma_{TR,NR}, \tag{2.18}$$

where $\sigma_{GR,NR}$ is the share of gross receipts to net receipts and $\sigma_{TR,NR}$ is the share of transfers to net receipts.

The weighted average within-period elasticity of government net receipts with respect to GDP, $\eta_{\tau,y}$, is estimated to be about 1.75.

2.7.3 The Tax Elasticity Restriction

From section 2.2, the system of equations 2.7-2.9 combined with identifying restrictions $\alpha_{11} = 0$ and $\alpha_{22} = 0$ yields the following:

$$u_t^\tau = \alpha_{12}\varepsilon_t^y + \varepsilon_t^\tau \tag{2.19}$$

$$u_t^g = \alpha_{21}\varepsilon_t^\tau + \varepsilon_t^g \tag{2.20}$$

$$u_t^y = \alpha_{31}\varepsilon_t^\tau + \alpha_{32}\varepsilon_t^g + \varepsilon_t^y \tag{2.21}$$

The additional restriction $\eta_{\tau,y} = \eta$ allows us to estimate this system of equations uniquely. If forecast errors u_t are transformed to represent growth rate errors, then $\eta_{\tau,y} = \frac{du_t^\tau}{du_t^y} = \frac{du_t^\tau}{d\varepsilon_t^y} = \alpha_{12} = \eta$.

CHAPTER III

Tax Salience and Charitable Giving

3.1 Introduction

Economists have often assumed in their models what they decry in their public statements, which is that consumers internalize the full costs of transactions, inclusive of taxes. Effective tax incidence is fully segregated from statutory tax incidence in standard economic theory, but we know better and have often said so. Milton Friedman argued against the income tax withholding system (in peace time) on primarily salience grounds.[1] John Stuart Mill (1848) also hypothesized that hidden tax systems lead to bigger government. Gregory Mankiw has on several occasions mentioned this possibility as a potential argument against the VAT, which he otherwise lauds as efficient.[2] Each of these notable economists has treated salience as an important component in ensuring that consumers completely account for the effect that taxes and tax rates will have on their choice sets.

Chetty et al. (2009) determine that sales taxes assessed at the checkout counter lead to qualitatively different consumer behavior than equivalent taxes embedded in the sticker price of a wide range of goods. The higher relative salience of embedded taxes leads to a stronger disincentive to consume goods than the less salient regime of adding taxes to the sticker price at the point of purchase, even when the consumer knows the sales tax rate. In a companion paper, Chetty et al. (2007) develop a model of bounded rationality to explain the importance of tax salience on consumer behavior, arguing that the first-order cognitive cost of determining the full price of a good is weighed against the second-order effect of re-optimizing consumption allocation in response to a small tax.

Finkelstein (2009) also finds evidence that the salience of a price regime, in this case highway tolls, distorts the consumer response to pricing changes. She determines that after the adoption of E-ZPass and similar electronic toll collection methods, the demand for toll roads became significantly

[1] Interview with Reason Magazine available at http://reason.com/archives/1995/06/01/best-of-both-worlds.
[2] See http://gregmankiw.blogspot.com/2009/10/value-added-tax.html for one example.

less price elastic.

This paper examines a component of the U.S. tax system long known for its opacity: the federal income tax code. In particular, I look at one particularly complicated part of the federal income tax: a shadow tax with its own rules that often controvert well-understood heuristics regarding deductibility and progressive tax brackets, known as the Alternative Minimum Tax.

The AMT was originally created as a back-up tax regime in response to a national outcry over several hundred high-net-worth individuals who, due to their judicious use of deductions, owed no federal income tax. Because 1) the AMT was not indexed to inflation until 2013, 2) AMT preferences such as state and local taxes have grown in scope relative to other deductions in the tax code, and 3) regular tax bracket rates have fallen significantly in the past 40 years, this tax now affects four million filers a year.

There are several reasons to believe that the AMT has lower salience than the rest of the federal income tax code. It primarily affects the tax return of the prior year as opposed to being withheld during the year in which income is earned and charitable donations are bestowed. Second, it is a relatively uncommon provision of the tax code in the sense that few taxpayers carry an AMT liability. Third, the AMT operates in the background of the primary tax bracket, rate, and deduction structure. Fourth, the universe of AMT preferences is quite complex (Burman et al. (2003) gives a comprehensive treatment of this issue). Even though the existence of the AMT is well-known and AMT liability is readily calculable, these timing and salience issues may have a substantial effect on consumer behavior.

The AMT has distinct tax brackets and marginal rates relative to the regular federal income tax. Typically taxpayers who end up with an AMT liability find themselves with a higher total tax bill but lower marginal rate than if they were governed solely by the primary tax regime. This variation in income and marginal rates presents an opportunity to test whether AMT status is fully anticipated or fully incorporated into the decision-making of the taxpayer through the vehicle of charitable giving, a consumption good that has a strong and well-established relationship with marginal tax rates and after-tax income.

Beginning with Taussig (1967), hundreds of papers been devoted to characterizing the influence that the tax code has on the level of charitable contributions. The richness of this body of research demonstrates that the relationship between marginal tax rates, after-tax income, and charitable giving can be readily estimated, which is helpful in any endeavor to test the salience of a particular aspect of the tax code.

The early cross-sectional studies of tax filer data showed that charitable contributions are both normal goods and relatively price elastic, spanning a rather narrow band of elasticities (Clotfelter 1985). Subsequent work, focused on eliminating omitted variable bias and endogeneity issues stemming from the timing of charitable contributions, has utilized panel data of both the tax file and survey varieties. The balance of this body of work is well-characterized by the meta-analysis of Peloza and Steel (2005) and literature survey by List (2011), which conclude that the (absolute value of the) permanent price elasticity of demand for charitable giving is somewhat greater than 1.

The rest of the paper goes as follows: in section 3.2 I describe the interaction between the AMT and charitable giving. In section 3.3 I discuss the data I will use in the paper. In section 3.4 I lay out the model(s) for charitable giving and modify them to test the salience of the AMT. In section 3.5 I discuss estimation results. In section 3.6 I conclude.

3.2 The Alternative Minimum Tax and Charitable Giving

3.2.1 AMT Structure

The federal Alternative Minimum Tax has its origins in 1969 when Congress first passed a backstop minimum tax to address public concerns that high-income filers were paying little or no tax as a result of their use of deductions to reduce taxable income.[3] At the time, no tax brackets were indexed for inflation, but the 1981 tax cut indexed most the regular income tax while omitting the AMT brackets and exemptions. Inflation was the largest culprit in making the AMT a relevant tax for a significant minority of taxpayers by the 2000s, at which point the president and Congress began issuing a series of one-time "fixes" that raised the exemption for a year or two to prevent millions more AMT taxpayers. In the course of resolving the "fiscal cliff" at the end of 2012, the AMT exemptions and tax brackets were permanently indexed at higher levels to prevent further bracket creep.

The federal AMT consists of a set of supplemental calculations that are conducted to determine an alternative measure of federal income taxes owed. This measure is then compared with the traditional personal income tax apparatus, and the filer owes the greater of the two. The term "minimum" in the AMT refers, then, to the fact that a taxpayer owes at least as much as their AMT liability.

The AMT takes gross income and subtracts qualifying deductions and exemptions, then applies its rate structure to determine AMT liability. Some nonrefundable and refundable credits are then

[3]See Burman et al (2003) for much more detail. Ironically, the taxpayers who escaped income taxes completely were mostly widows and retirees deriving their income from tax-exempt municipal bonds, which is still allowed under current AMT rules as long as the bonds are not "private activity bonds."

netted against this liability to determine a filer's final liability, just as in the federal income tax. The biggest differences between the regular tax and the AMT are that 1) the AMT is much less generous in the type of deductions that it allows to be counted against gross income, and 2) the rate structure of the AMT is much flatter, with a higher exemption. As a result, the type of filer who is most likely to owe an AMT liability is one with a high-enough income to exhaust the relatively large exemption and a significant number of deductions that qualify under the rules of the regular income tax but not the AMT.

The types of deductions that cannot be taken against the AMT are called AMT preferences and adjustments, and they include state and local income taxes, personal exemptions, the standard deduction, property taxes, and most miscellaneous itemized deductions, as well as a host of less common deductions. In addition, unexercised stock options must be recognized for AMT purposes in the year in which they are earned, valued at the difference between their market price and option price. The deductions which can be counted against the AMT include medical expenses (in excess of AGI), charitable contributions, and mortgage interest (so long as the mortgage was used to purchase or improve your home).

The AMT exemption, which substitutes for a standard deduction and personal exemptions, was \$53,600 in 2015 for single taxpayers and \$83,400 for married taxpayers filing jointly. The two statutory rates are 26 percent and 28 percent, with the higher rate applicable for taxable income (for AMT purposes) above \$185,400 for both single and married filers. In addition, beginning at \$119,200 (\$158,900 for married filing jointly) in taxable income the AMT exemption begins to phase out at a rate of \$1 for each \$4 in additional income. The true marginal (federal) rate, then, of an AMT filer progresses from 26 percent to $26 * 1.25 = 32.5$ percent, reaching a high of $28 * 1.25 = 35$ percent before the exemption is exhausted, at which point the rate returns to 28 percent.

3.2.2 Interaction with Charitable Giving

As a non-AMT preference item, charitable donations can be deducted against a filer's AMT taxable income, just like the regular federal income tax. Therefore, the after-tax price of giving, P_{it}, is either $1 - \tau_{it}^{regular}$ or $1 - \tau_{it}^{AMT}$ depending on whether the filer has an AMT liability, where $\tau_{it}^{regular}$ is the marginal tax rate for a filer under the traditional income tax structure for filer i at time t and τ_{it}^{AMT} is the corresponding marginal tax rate under the AMT regime.

For filers who are at risk of incurring an AMT liability, their statutory federal marginal tax rate can vary widely; if they have many personal exemptions or live in a high-tax state, they might be in the 15 percent bracket and still end up owing the AMT. More commonly, filers are in the upper middle class, as the relatively large exemption zeros out any AMT tax liability at modest income levels, and the more progressive tax rates under the regular rate structure create a higher tax liability for very-high income filers than the flat rate structure of the AMT.

As a result, an AMT filer might have a smaller, equal, or larger marginal tax rate than they would under the regular income tax, and therefore a larger, equal, or smaller tax price of charitable donations. In practice, AMT filers are about twice as likely to have lower marginal tax rates than higher ones (see Burman et al (2003)), so the AMT tends to raise the price of charitable giving.

3.3 Creating Panel Tax Data

My analysis uses longitudinal data from the Panel Survey on Income Dynamics (PSID). The PSID contains data on wages and salary, self-employment income, business income, transfers, and passive income such as rent, pensions, and dividends, as well as rich demographic information about each family unit. I combine the 1999-2013 biannual surveys to construct a panel dataset for respondents that spans the tax years 1998 to 2012. For each year I construct key income variables following Kim et al. (2014) and input these variables into the TAXSIM program maintained by Daniel Feenberg.[4] Since these are survey data rather than tax data, the inputs into TAXSIM are less precise than the actual tax data available via the Statistics of Income (SOI) or other confidential panel data from the Department of the Treasury. This source of measurement error may cause some amount of attenuation bias, investigated more fully below. As Peloza and Steel (2005) note, however, estimates of the price elasticity of demand for charitable giving are nevertheless frequently higher for survey data. Biased reporting of contribution levels, such as in Slemrod (1989), may inflate elasticities if misreporting increases with income levels (and also tax rates).

The key benefit of using the PSID is that it is a publicly available panel dataset with demographic, income, and charitable giving information for those family units who itemize their federal income taxes. In addition, the income data in the PSID is quite detailed and may represent total income more accurately than income reported on a tax form. The key downside is that the sample size of the PSID is moderate (about 8,000 family units responded to each survey between 1999 and 2013) and the sample is skewed towards moderate-income families who are neither as likely to itemize their deductions nor to find themselves with an AMT liability in a given year. Nevertheless, the sheer magnitude of the increase in the number of AMT filers over the past 15 years has meant that both

[4]The TAXSIM program is able to identify only a narrow definition of an AMT taxpayer, specifically one who owes an AMT liability after the regular federal income tax and credits have been applied. A broader definition of AMT taxpayer, such as any taxpayer whose net liabilities increase as a result of the AMT (due to the limitation that the AMT puts on credits like the general business credit), outstrips the capabilities of TAXSIM.

moderate-income and upper-middle-class families have found themselves facing an AMT liability with greater frequency than in the past. About 2 percent of family units in 2012 owed tax as a result of the AMT in the PSID, which over-samples poorer families, who are unlikely to have an AMT liability.

In addition to income, tax, and demographic detail, I also constructed wealth variables for each tax unit. The presence of wealth in any charitable giving regression can help eliminate confusion of temporary with permanent price and income elasticities, and it is especially important when using survey data that are lacking important information on realized capital gains, which correlate with contributions. Since changes in wealth are the primary drivers of this correlation, inclusion of wealth data is a necessary step in preventing bias in the estimates.

The summary statistics of the tax-augmented PSID dataset are presented in table 3.1. Roughly 31% of tax units make itemizable charitable contributions, and those who itemize and make charitable contributions have significantly higher median income, average wealth, and median wealth. The average charitable contribution among itemizers with a non-zero contribution is $3,441, with significant variance between and within income groupings (see figure 3.1). The sample skews slightly older and is much more likely to be married than the sample as a whole. AMT filers, at 1.1 percent, represent a small fraction of the total sample, and they are quite wealthy relative to the PSID sample. A key point arises from the summary statistics in table 1, which is that one challenge of any estimation technique that distinguishes AMT filers from their fellow taxpayers is to allow sufficient flexibility of income and wealth effects on charitable contributions so that AMT filers are compared with similarly situated families rather than imposing a constricting functional form that implicitly contrasts behavior of high income families with those whose incomes are much more modest.

Table 3.2 records key summary statistics on itemizing deductions, charitable donations, and AMT liability by income group in the sample. The share of itemizers, donors, and AMT filers all rise (unsurprisingly) with income. In addition, the number of AMT filers is overrepresented in the sample as a share of high-income taxpayers, while it is underrepresented in the middle-income range. One potential cause of this underrepresentation is the limited information provided on itemized deductions; another is the simplification of the tax code when in TAXSIM, which ignores some credits that might be disqualified by the AMT.

The number of AMT taxpayers, shown as a whole in table 3.1 and by income group in table 3.2, is quite small in my sample. Only 733 records have a positive AMT liability. Of these records, over 2/3 itemize and made a charitable donation in the year in which they incurred this liability, averaging $4,291 in contributions. Although this number is higher than the average for itemizers with a contribution, table 3.2 shows that the comparison is misleading; within each income class, mean charitable donations are lower for AMT taxpayers than for their non-AMT counterparts, whether the sample is restricted to donors or not.

Lower giving rates and levels could mean 1) that higher tax bills and higher tax prices on donations are at least somewhat salient to AMT filers; 2) that AMT filers are situated at or near the bottom of each income category; 3) that AMT filers are disproportionately likely to be earning high temporary income (which presumably is less likely to be spent on charity and more likely to be saved relative to permanent income); 4) that AMT filers have other attributes that correlate with lower giving; or 5) that the small sample size is leading to over-interpretation of sample means. I investigate these hypotheses in more detail below.

3.4 Model

3.4.1 Cross-Sectional Charitable-Giving Models

The canonical model of charitable contributions, first specified by Taussig (1967),[5] is given by:

$$\ln g_{it} = \alpha_1 \ln Y_{it}^d + \alpha_2 \ln P_{it} + X_{it}\beta + \varepsilon_{it}, \tag{3.1}$$

where g_{it} is the level of giving that appears on an itemized deduction, Y_{it}^d is discretionary (or after-tax) income, and P_{it} is the price of charitable contributions, which depends on the marginal federal and state income tax rates as well as the extent to which contributions are deductible on both the federal and state returns. The additional controls X_{it} are traditionally the kinds of demographic information that can be found on a tax return, such as indicators for married or single, the number of dependent exemptions taken, and an indicator for whether a tax unit has one or more members over the age of 65. With the additional demographic information available from survey data, controls can be richer, including age, number of children, educational attainment, race, and wealth. The coefficient α_1 in this specification represents the income elasticity of charitable giving, while α_2 represents the tax price elasticity of demand.

Taussig (1967) used the marginal tax rate, net of contributions, as his regressor in determining the price effect on giving. This rate constituted, in his mind, the true price at the margin for incremental decision-making. Feldstein and Taylor (1976), in contrast, used the first-dollar marginal

[5] As is mentioned below, Taussig used the marginal tax rate as opposed to the price of charitable donations, which leads to a slightly different interpretation of the "price" coefficient.

rate (or the marginal rate ex contributions) to avoid the bias that stems from higher contributions driving down the net-of-contribution marginal rate. This could induce a positive correlation between the after-tax price and the amount of giving, in contrast to the assumed negative relation between price and quantity demanded. A solution to this dilemma, now widely adopted in the literature, is to use the net-of-contribution price for charitable giving, instrumented by the first-dollar price.

An additional problem with the canonical model of contributions is that it does not measure "true" income (in the Haig-Simons sense); if the types of income that are unreported on tax forms influence giving, which seems quite likely, then their omission from the model can potentially bias the other results. Wealth data are a valuable correction for this potential problem, and in the PSID I am able to create a measure of net wealth (assets minus debts, including home mortgage) for each tax unit.

A final concern with the estimation of equation 3.1 is the enforcement of uniform price and income elasticities across a range of personal financial situations. Taussig himself estimated the elasticities separately for different adjusted gross income (AGI) categories, as well as Feldstein and Taylor and many subsequent analyses.

The results of estimating the canonical model on PSID data, which will be amended later in order to test several hypotheses regarding AMT status and giving behavior, are shown in table 3.3. In column (a), I show the estimates from a simple least squares regression of the specification in equation 3.1. Column (b) in table 3.3 shows the effects of instrumenting the last-dollar tax rate (the tax rate after deducting charitable donations) with the first-dollar tax rate (the tax rate if no charitable donation had been made). Column (c) includes wealth and gift data as additional regressors. Column (d) estimates income elasticities separately by gross income class, as well as allowing non-price regressors (not counting wealth) to have distinct effects in each income class. Columns (c) and (d) represent the baseline charitable giving models on which I test AMT salience. The specification in Column (d) adds additional flexibility to the model, which is otherwise characterized by a single elasticity across all income levels, in a manner that is consistent with the literature (see Bakija and Heim (2011)). On the other hand, it imposes a cost in power given the small sample size of AMT taxpayers. Even in the baseline scenario, at the 5 percent significance level I cannot reject the null that any two of the separate estimates by class are equal to each other.

The price elasticity of demand is -0.6 and -0.68 in specifications (c) and (d), respectively, meaning that an increase of 10 percent in the after-tax price of giving (e.g. from a 35 percent marginal tax rate to a 28.5 percent marginal tax rate) lowers donations by 6 percent, holding after-tax income constant.[6] This result is on the weaker end of measured price responsiveness in the literature, as the central tendency of the price elasticity of demand is around -1 or more negative (Peloza and Steel (20005), List (2011)). The income elasticity of demand is 0.46 in the constant income elasticity specification, implying that a 10 percent increase in after-tax income should raise donations by 4.6 percent. Allowing the income elasticity to vary by gross income category generates similar effects except at the lowest income class (less than $20,000), where the elasticity is 0.25. This change is not statistically significant, however, at the 95 percent confidence level.

As a first step in determining the effect of AMT status on charitable giving, I amend the formulation discussed above with an additional indicator variable that takes the value of 1 if the filer is an AMT taxpayer and 0 if the filer is not:

$$\ln g_{it} = \gamma D_{it}^{AMT} + \alpha_1 \ln Y_{it}^d + \alpha_2 \ln P_{it} + X_{it}\beta + \varepsilon_{it}. \tag{3.2}$$

In table 3.4 I present the results of this preliminary multivariate analysis. The effect of including the indicator variable is minimal with respect to the other regressors, but suggests that AMT status is correlated with lower charitable contributions. The simplest interpretation of the coefficients on D_{it}^{AMT} in columns (a) and (b) is that AMT taxpayers donate 9 percent or 19 percent fewer dollars to charity than equivalent non-AMT taxpayers; only the estimate in (b) (where income elasticity is allowed to vary by gross income class) is statistically significant. In general, the small sample size of AMT filers (733), of which only 2/3 have recorded donations, makes it difficult to determine whether the apparent effect is real.

3.4.2 Estimation Challenges

To the researcher, estimation of an equation such as 3.2 is full of pitfalls. One potential source of endogeneity is that a filer might realize capital gains in a year in which they intend to donate, which makes the error term positively correlated with realized income Y_{it}^d. Another is that temporary boosts to income might not reflect underlying permanent income, so Y_{it}^d and the error term could be negatively correlated. A third arises from the fact that larger contributions g_{it} change the level of after-tax income and potentially the marginal tax rate (in a progressive tax code). And to the extent that omitted variables affect charitable giving and correlate with observables like income or AMT status, their existence could also bias estimation.

[6] In this example, the after-tax price rises from 65 percent to 71.5 percent, a 10 percent increase, while the marginal tax falls by a greater share.

Measurement error is a source of endogeneity that is always a source of concern in non-tax data. Classical measurement error tends to bias results towards zero, but in the case of measurement error in income, the price variable can be affected as well. Even measurement error of the dependent variable, log charitable giving, can in this case bias results; an overstated level of giving will lower the perceived tax burden to the econometrician, creating a corresponding positive error in presumed Y_{it}^d, but according to the dynamics of the tax code rather than the underlying income elasticity of demand.

I am able to address several, but not all, of these potential sources of bias in the cross-sectional framework. As mentioned above, I instrument the actual values of Y_{it}^d and P_{it} with their first-dollar counterparts. This approach eliminates the endogeneity that stems from simultaneous increased contributions altering the regressors and biasing α_1 and α_2. In addition, by using wealth data as a supplemental regressor to be combined with income data, I can mitigate the bias of temporary and realized income on estimates of the price and income elasticity of demand for charitable donations, although the timing of realization of income to coincide with charitable contributions is not directly addressable.

In the context of the results in table 3.4, the persistently negative effect of AMT status on charitable contributions, even after accounting for the change in after-tax price of donating and the change in income due to owing an AMT liability, suggests that either 1) AMT filers are overreacting to the change in their taxes, 2) the specification above is missing systematic unobserved differences between AMT filers that drive charitable donations, or 3) my sample size is inadequate to address the question satisfactorily.

The first explanation, while it can not be dismissed out of hand, is at best simply question begging and at worst completely backwards. I find that the negative effect on charitable contributions of owing AMT liability amplifies over time, i.e. that new filers react less quickly than filers who had faced a similar situation in the prior survey (not shown). The second explanation, that the simple cross-sectional specification has omitted variable bias, leads naturally towards a repeat of the foregoing investigation, this time with a fixed-effects model in order to make fuller use of the panel dataset.

3.4.3 Longitudinal Charitable-Giving Models

The basic log-linearized charitable demand function, estimated in a full panel setting, is represented as:

$$\ln g_{it} = \delta_i + \alpha_1 \ln Y_{it}^d + +\alpha_2 \ln P_{it} + X_{it}\beta + \varepsilon_{it}. \tag{3.3}$$

The interpretations of the coefficients α_1, α_2, and β remain the same if the model is correctly estimated. The presence of the individual fixed effect term, δ_i, improves estimation of the other coefficients if, in the cross-sectional estimation procedure in equation 3.1, omitted variables systematically correlated with one or more of the RHS variables. To the extent that those omitted variables, such as temporarily distorted income, change over time, the estimates in equation 3.3 will remain biased.

Although Peloza and Steel (2005) have concluded that the bias from omitted covariates and temporary changes to income and tax liabilities in charitable giving models appears to be small, the difference between permanent and temporary effects might be magnified for those individuals who find themselves on the AMT.

An error-correction model for price and income effects, as shown by Bakija and Heim (2011)[7], can address the potentially confounding effect of temporary income (and temporarily distorted marginal tax rates):

$$\ln g_{it} = \delta_i + \alpha_1 \Delta \ln Y_{it}^d + \alpha_2 \ln Y_{it}^d + \alpha_3 \Delta \ln P_{it} + \alpha_4 \ln P_{it} + X_{it}\beta + \varepsilon_{it} \tag{3.4}$$

The interpretation of the above model is that the permanent price elasticity of demand for charitable giving is α_4, while the temporary elasticity is $\alpha_3 + \alpha_4$. In order to estimate these equations, I instrument $\ln Y_{it}^d$ and $\ln P_{it}$ with their first-dollar equivalent, as well as $\Delta \ln P_{it}$ and $\Delta \ln Y_{it}^d$, as in the cross-sectional models in section 3.4.1.

Adding an AMT dummy to equations 3.3 and 3.4, I estimate the following specifications,

$$\ln g_{it} = \delta_i + \alpha_1 \ln Y_{it}^d + \alpha_2 \ln P_{it} + \gamma D_{it}^{AMT} + X_{it}\beta + \varepsilon_{it} \tag{3.5}$$

$$\ln g_{it} = \delta_i + \alpha_1 \Delta \ln Y_{it}^d + \alpha_2 \ln Y_{it}^d + \alpha_3 \Delta \ln P_{it} + \alpha_4 \ln P_{it} + \gamma D_{it}^{AMT} + X_{it}\beta + \varepsilon_{it}, \tag{3.6}$$

[7] Bakija and Heim go a step further and estimate future expected income, a procedure that has a modest effect on their results and that I do not follow here.

and show the results in table 3.5. The results are strikingly different from those in table 3.4. To begin with, the estimated permanent price elasticity, α_2 in equation 3.5 and α_4 in equation 3.6, is significantly attenuated from the cross-sectional estimate. In all the specifications provided in table 3.5 the price elasticity is not statistically different from 0 at any standard significant level. In addition, the sign on the D_{it}^{AMT} term is now positive, albeit not statistically different from 0 either. That the point estimate of γ becomes positive when person-specific fixed effects are included (and remains positive when temporary income and price effects are controlled for) should lead to skepticism at the results in table 3.4.

Unlike in the cross-sectional procedure, an individual fixed-effects coefficient soaks up systematic differences in unobserved characteristics between AMT filers and non-AMT taxpayers, potentially reducing the bias such differences may have imparted on the results in table 3.4. The D_{it}^{AMT} term, then, is identified off individual variation in tax status. Unfortunately, the low sample size of AMT filers is somewhat exacerbated by that fact that only a subset of them move on and off the AMT, reducing the number of records that can identify γ from 733 to 611.

Leaving the discussion of significance aside for the moment, the results from Table 3.5 (if accurate) indicate that AMT taxpayers made slightly *higher* charitable contributions than their marginal tax rates and after-tax income implied. But, in a pattern that will be repeated throughout the discussion of the results of the panel models, these coefficients are only slightly positive and much smaller than their respective standard errors. This observation holds true whether the income elasticity is uniform across all gross income classes or allowed to vary, and whether multiple tax years are included as regressors.

3.4.4 Salience Tests in Cross-Sectional Charitable-Giving Models

The estimation of a demand function for charitable donations has a graphical representation that motivates the tests for tax salience below. In figure 3.2 I illustrate an individual's hypothetical demand curve for giving as a function of the after-tax price (for a given level of after-tax income). At a 35 percent marginal tax rate, the after-tax price is 0.65 and the quantity demanded is $1,000 in contributions. The triggering of an AMT liability, with a 28 percent marginal rate, should raise the price to 0.72, which is an increase of about 11 percent. Consequently, the taxpayer's demand for charitable contributions should fall to around $900 (in red), since in this example the price elasticity is -1. In such a scenario, the econometrician would not find any particular AMT effect distinct from the standard predicted level of charitable contributions. If the lower marginal rate (and thus higher after-tax price) is not incorporated into the filer's information set, however, they might contribute the same amount to charity as that implied by the tax price corresponding to their traditional marginal rate (again, \$1,000 in the example), which is off the demand curve in green. If either some fraction of filers incorporate the new tax price into their decision-making process or all filers incompletely incorporate the price into their decision, then the result is something in between, in blue.

To investigate tax salience empirically in this setting I divide the after-tax price and income terms in the charitable giving model into pre-AMT and post-AMT components. In the absence of the AMT, a taxpayer perceives a marginal tax rate applicable to charitable giving P^*_{it} and owes tax liability T^*_{it}, leaving disposable income $Y^{d^*}_{it} = Y_{it} - T^*_{it}$. In most cases $P^*_{it} = P_{it}$ and $Y^{d^*}_{it} = Y^d_{it}$, but for those hit by an AMT liability $Y^{d^*}_{it} < Y^d_{it}$ and P^*_{it} is less than, equal to, or greater than P_{it} depending on the tax bracket of the individual involved. We can express equation 3.1 above decomposed to separate out the "AMT effect" like this:

$$\ln g_{it} = \alpha_1 \ln Y^{d^*}_{it} \frac{Y^d_{it}}{Y^{d^*}_{it}} + \alpha_2 \ln P^*_{it} \frac{P_{it}}{P^*_{it}} + X_{it}\beta + \varepsilon_{it}, \tag{3.7}$$

which I estimate as:

$$\ln g_{it} = \alpha_1 \ln Y^{d^*}_{it} + \gamma_1 \ln \frac{Y^d_{it}}{Y^{d^*}_{it}} + \alpha_2 \ln P^*_{it} + \gamma_2 \ln \frac{P_{it}}{P^*_{it}} + X_{it}\beta + \varepsilon_{it}. \tag{3.8}$$

The test for salience, then, involves comparison between the terms α_1 and γ_1, as well as α_2 and γ_2. In figure 3.3 I illustrate the additional income effect of triggering an AMT liability; the demand curve for charitable contributions is lowered by the extent to which after-tax income is decreased, scaled by the elasticity of demand for charitable giving with respect to income (0.4 in this example). In this illustration, full- and zero- incorporation of information into the decision-making process are both depicted, as well as an intermediate response that reflects incomplete adjustment to tax information on both the income and tax price dimensions.

If the filer responds in a predictable way to the increase in price and decrease in after-tax income, shown in red, then $\alpha_1 = \gamma_1$ and $\alpha_2 = \gamma_2$. If the filer is oblivious to the change in after-tax income and marginal tax rate due to triggering an AMT liability, shown in green, then they continue to give \$1,000 and $\gamma_1 = \gamma_2 = 0$. In the intermediate response, shown in blue, a variety of estimate combinations could generate the same change in giving. Because the econometrician only views the final giving amount, disentangling the additional income and price effects of an AMT trigger is just as challenging as separating the original income and price effects, with the added difficulties of 1) a smaller sample size of AMT filers than the tax base as a whole, and 2) needing enough controls or proxies to ensure that AMT filers are not systematically different in unobservables from the families to which they are being compared in a counterfactual sense. Nevertheless, figure 3.3 inspires a range of possible tests to determine whether taxpayers on the AMT incorporate their status into their decision-making.

The relevant tests on equation 3.8 are:

1. Full Income Salience: $\alpha_1 = \gamma_1$
2. No Income Salience: $\gamma_1 = 0$
3. Full Price Salience: $\alpha_2 = \gamma_2$
4. No Price Salience: $\gamma_2 = 0$
5. Full Joint Salience: $\alpha_1 = \gamma_1$ & $\alpha_2 = \gamma_2$
6. No Joint Salience: $\gamma_1 = 0$ & $\gamma_2 = 0$

A risk in attributing all the differences, if they exist, between α_i and γ_i to misperceptions of price and after-tax income is the potential for composition bias, a consequence of assuming one single price- and income- elasticity for all income levels. AMT payers differ systematically from non-AMT payers in that they have more income on average. I follow the approach taken in Bakija and Heim (2011) and interact all non-price controls, as well as the income variables, with a set of dummy variables that indicate different income levels.

3.4.5 Salience Tests in Longitudinal Charitable-Giving Models

Again, cross-sectional results from direct estimation of equation 3.8 suffer from several shortcomings that can bias the estimates of α_1, α_2, γ_1, and γ_2.

I extend the identification process in equation 3.8 to a fixed-effect model and an error-correction model in order to disentangle the income and price effects of AMT incidence. The first is given as:

$$\ln g_{it} = \delta_i + \alpha_1 \ln Y_{it}^{d^*} + \gamma_1 \ln \frac{Y_{it}^d}{Y_{it}^{d^*}} + \alpha_2 \ln P_{it}^* + \gamma_2 \ln \frac{P_{it}}{P_{it}^*} + X_{it}\beta + \varepsilon_{it}. \tag{3.9}$$

The relevant tests are the same, namely that $\alpha_1 = \gamma_1$ and $\alpha_2 = \gamma_2$ if taxpayers have fully incorporated their AMT status into their decision process. The panel specification of equation 3.9 ought to absorb individual fixed effects and eliminate bias from stationary omitted variables, though not ones that evolve over time.

Performing the same decomposition as in equations 3.8 and 3.9, I embed a test for tax salience in the equation above and instrument each tax variable with its first-dollar counterpart to estimate the following model:

$$\ln g_{it} = \delta_i + \alpha_1 \Delta \ln Y_{it}^{d^*} + \alpha_2 \ln Y_{it}^{d^*} + \gamma_1 \ln \frac{Y_{it}^d}{Y_{it}^{d^*}} + \alpha_3 \Delta \ln P_{it}^* + \alpha_4 \ln P_{it}^* + \gamma_2 \ln \frac{P_{it}}{P_{it}^*} + X_{it}\beta + \varepsilon_{it} \qquad (3.10)$$

where $\Delta \ln Y_{it}^{d^*}$ is defined as $\ln \frac{Y_{it}^{d^*}}{Y_{it-1}^{d}}$. In other words, taxpayers always understand their last year's income tax liability, and the only question is whether they anticipate AMT liability in the year in which it is being incurred. This specification for tax salience has different tests that correspond to different interpretations of why a filer becomes an AMT filer. If AMT liability results from aberrations from a normal year, then $\gamma_1 = \alpha_1 + \alpha_2$ and $\gamma_2 = \alpha_3 + \alpha_4$ are the relevant tests, the derivation of which are straightforward and consigned to Appendix 3.7.2. This is my preferred set of tests. The incremental deviation from the (regular-income-tax-derived) demand curve resulting from AMT liability ought to match the temporary elasticity.

If, instead, an AMT liability is a normal state of affairs for a filer, then the test for salience is $\gamma_1 = \alpha_2$ and $\gamma_2 = \alpha_4$, equivalent to the permanent income and price elasticity measures. This latter interpretation, however, assumes the answer to the hypothesis it is testing: if filers are regular AMT taxpayers, then they are undoubtedly aware of this state of affairs, so failing the test indicates something other than what is being tested.

Thus the appropriate salience tests on equation 3.10 are as follows:

1. Full Income Salience: $\gamma_1 = \alpha_1 + \alpha_2$
2. No Income Salience: $\gamma_1 = 0$
3. Full Price Salience: $\gamma_2 = \alpha_3 + \alpha_4$
4. No Price Salience: $\gamma_2 = 0$
5. Full Joint Salience: $\gamma_1 = \alpha_1 + \alpha_2$ & $\gamma_2 = \alpha_3 + \alpha_4$
6. No Joint Salience: $\gamma_1 = 0$ & $\gamma_2 = 0$

These tests provide the proper "salience" interpretation of the coefficients and standard errors of the charitable giving demand function estimation above.

3.5 Results

3.5.1 Cross-Sectional Results

The results of the estimation of the cross-sectional specifications in section 3.4.4 are shown in table 3.6. In all specifications, the tax price and the after-tax income variables are instrumented by their first-dollar counterparts. In the first grouping, the price elasticity of demand and income elasticity of charitable giving are both assumed to be constant across income levels. The base case, which yields a pair of elasticities that are readily compared with those in the considerable aforementioned literature, appears to heighten fears of at least some attenuation bias stemming from measurement error. The tax price elasticity of demand for charitable giving is relatively inelastic, coming in at -0.6, while the income elasticity of demand is estimated at 0.46 (column a). When discretionary income, as well as other non-price regressors, is allowed to vary by gross income categories (column c), the price elasticity rises to -0.68, while the income elasticity ranges between 0.25 and 0.41. While these results are well within the bounds of the various estimates of the past 50 years, more recent data and survey-based data have tended to find relatively high price elasticities.

The base cases reflect attempts to remove or otherwise minimize endogeneity from the cross-sectional specification. A simple OLS estimation of equation 1 above yields very similar elasticities to those reported in table 3.6, column a (see table 3.1). Instrumenting the tax price of charitable giving with the first-dollar equivalent and doing the same for discretionary income gives almost an identical result for the price elasticity of demand, albeit raising the income elasticity somewhat. This similarity reflects the extent that the marginal tax rate is rather unlikely to change as a result of the magnitude of charitable donations observed in the data.[8] Adding net wealth and gifts lessens the income effect somewhat; wealth and income are positively correlated, so the extent to which realized income just represents changes in net wealth (i.e. capital gains), it is captured by the latter term.

Inclusion of the nested terms for testing AMT salience in table 3.6 yields results with conflicting interpretations. The incremental income effect appears to be much weaker than the baseline model would indicate, but the price effect implies higher-than-normal responsiveness to the corresponding marginal rates. In the language of equation 3.8, both the income-salience tests $\alpha_1 = \gamma_1$ and $\gamma_1 = 0$ fail to reject the null hypothesis because of the large standard errors on the coefficient of the $\ln \frac{Y_{it}^d}{Y_{it}^{d*}}$ term, while both the price-salience tests $\alpha_2 = \gamma_2$ and $\gamma_2 = 0$ can be rejected at the 1 percent significance level because of the high degree of responsiveness of the $\ln \frac{P_{it}}{P_{it}^*}$ term.

[8]Table 3.9 illustrates the degree to which instruments co-vary with their instrumented variables; the first stage of each IV regression is quite strong.

Allowing the non-price variables, including demographic, wealth, and income measures, to vary by gross income class yields slightly different results, but with the same upshot. The standard error on the AMT income effect is too large to reject the possibility of full salience, while the incremental AMT price effect is strangely powerful and significantly different from the estimate for the rest of the sample.

The constant-income-elasticity specification is readily tested against restrictions that imply either full salience or no salience, namely the joint test that $\alpha_1 = \gamma_1$ and $\alpha_2 = \gamma_2$ (full salience) and $\gamma_1 = 0$ and $\gamma_2 = 0$ (no salience). Both of these tests are rejected with a high degree of confidence (at the 1 percent level). The loosened income-elasticity specification is harder to test; I show the test on the highest income class in table 3.6, but a comparison of the AMT after-tax adjustment term with every single class's income elasticity in a joint test yields the same result - rejection of the nulls of both full and no salience about AMT tax variables.

These joint rejections of both hypotheses are interesting, but ultimately all the rejections are on the back of the outlier coefficients on the price adjustment terms. The combination of low sample sizes and difficulty in selectively apportioning the combined effect of lower income and, on average, a higher tax price of giving appears to create a significant barrier to fully identifying the salience of the AMT in augmented PSID data using cross-sectional approaches.

3.5.2 Panel Results

The results from estimating the panel specifications in section 3.4.3 are given in tables 3.7 and 3.8. In light of the suggestive but statistically insignificant preliminary analysis in table 3.6, we turn to results for the more detailed decomposition of the fixed-effects charitable giving demand function. Table 3.7 restricts the model to a uniform (permanent) income elasticity. In column (b) I report the results of estimation of equation 3.9 (column (a) shows the corresponding base case IV model with fixed effects but no AMT decomposition). Column (d) shows the results from estimating equation 3.10 (with column (c) again playing the role of the base case). The income elasticity of the base case in both the standard fixed-effects specification and the error-correction specification is similar to that of the cross-sectional results. The tax price effects, however, are quite minor and in some cases are the wrong sign (i.e., inconsistent with a downward-sloping demand curve). In contrast with the cross-sectional results, which appear to suffer from significant mis-specification, the panel data results are distinguished by a lack of variation and thus low-power tests.[9] Both the standard fixed-

[9]Measurement error, as I alluded to above, is a concern in all of these specifications, but particularly in panel data. If, for example, individual income is characterized by a high degree of auto-correlation, but measurement error has low auto-correlation, then the attenuation bias is magnified beyond that for cross-sectional results, which may be the case

effects estimation and the error-correction alternative fail to reject the null of either full salience or no salience on both the income effect and the tax price effect of the AMT on charitable giving.

Table 3.8 shows results of estimation of the same set of demand functions, with the income elasticity and other non-price variables free to differ by gross income class. Again, high standard errors in the terms that compare 1) the tax price of charitable giving under the regular income tax regime versus the AMT regime and 2) the discretionary income available to spend on charity under the regular income tax regime versus the AMT regime, lead to inconclusive tests for salience in both the fixed-effect and error-correction specifications.

3.5.3 Measurement Error

To determine whether measurement error was a factor in the spate of inconclusive tests above, and for lack of an instrument to remove measurement error from survey data, I added iid, normally distributed, mean zero measurement error to income in order to determine whether it resulted in attenuation bias. Attenuation bias resulting from this exercise could imply the existence of existing bias in the "unbiased" estimates in tables 3.3-3.8. In table 3.10, the simple cross-sectional demand model is compared with the results from the addition of error in income. Also, I compare both sets of results with a second exercise in which classical mean zero, iid measurement error is added to reported charity (for those who already report non-zero contributions). Surprisingly, income error biases the price term towards zero but not the income term. Measurement error in charitable contributions, however, attenuates both variables. These unexpected results are probably driven by some nonlinear relationships among the three variables that have not been fully explored.

3.6 Conclusion

This paper develops a method for testing the salience of the Alternative Minimum Tax via a nested demand model in cross-sectional and panel formats. The ambiguity of empirical results highlights sample-size limitations, as well as the limitations of computing tax variables from survey data, in which small errors can aggregate and potentially bias important relationships.

I find a negative effect of AMT liability on charity in cross-tabs and in cross-sectional regressions, even after accounting for the expected changes to demand for contributions given lower after-tax income and the higher after-tax price of donating. In contrast, the difference between actual and

here. Compounding the problem is that measurement error in income variables is not obviously classical measurement error, since the error passes through a tax simulator and affects the price regressor. Classical measurement error is addressable through IV, as the error is uncorrelated with the instrument used on it, but non-random measurement error is more intractable. I analyze the effects of measurement error in income and charitable donations below.

expected contributions for AMT taxpayers swings to become weakly positive but statistically insignificant when person fixed effects or temporary income effects are accounted for.

I find similar results when I measure income and price effects in order to determine tax salience - the wrong sign in cross-sectional specification implies an overreaction to AMT incidence, while panel specifications suffer from large standard errors throughout that make interpretation of their results difficult.

The weak effects in both types of analysis point to some related problems (small sample size, measurement error) and some unique ones. Cross-sectional analysis appears to be insufficient to capture the attributes of AMT taxpayers, which systematically differ from others in their income cohort (likely due to higher than average temporary income), while there is good reason to suspect that panel data magnifies the problems of measurement error and relies too heavily on changes in AMT status from one year to the next to generate exogenous variation.

Nevertheless, the inability to distinguish between such extremely different behavioral models as 1) full incorporation of tax status into the taxpayer's information set, and 2) full ignorance about tax status, is intriguing. Further investigation with more accurate tax data and larger sample sizes, especially in the income region where AMT liabilities are more prevalent, is warranted in order to determine with greater confidence whether the AMT and other non-salient tax features (such as the PEASE limitation) are being treated rationally by taxpayers, or are simply being ignored.

3.7 Appendix

3.7.1 Technical Appendix on Data

This section of the appendix details the methods employed to build a panel data set with income, demographic, tax, and wealth data using the PSID and TAXSIM.

3.7.1.1 Income Aggregates

I extracted relevant income information from each of the 1999-2013 family-level surveys in order to create aggregate income variables that could be fed into a tax simulator (in this case TAXSIM9). In order to do this I built on and amended code that was generously provided by M. Marit Rehavi. The key variables include tax year, state of residence, marital status, number of dependent exemptions, number of age exemptions, earned income of the household head, earned income of the spouse (if relevant), dividend income, property income, pension income, Social Security income, non-taxable transfer income, property taxes, rental payments, itemized deductions, childcare expenditures, unemployment compensation, number of dependent children, and mortgage interest. TAXSIM requires short-term and long-term capital gains as well, but this information is not available in the PSID.

To create an earned income measure, I combined wages, bonuses, tips, overtime, commissions, professional income, garden income, and other labor income. Property income is the combination of interest income, trust income, rental income, net alimony income, and other income. Beginning in 2003, property income is amended to be inclusive of dividends, as the TAXSIM calculator asks for only "qualified" dividends. Pension income is the combination of VA pensions, non-VA pensions, annuities, and other pensions (not inclusive of Social Security). Prior to the 2013 survey, the PSID aggregated pension, annuity, and "retirement income" for the wife of the household head, but beginning in 2013 these categories are separated. Social Security income is provided on a total family basis prior to 2005, but afterwards head and spouse benefits are provided separately. Non-taxable transfer income includes child support, TANF/ADC income, supplemental security income, and other welfare income. Unemployment compensation is created from the combination of unemployment insurance and workers' compensation.

Mortgage interest is calculated in two ways. The simpler method is to use the outstanding mortgage on the first two properties listed multiplied by the mortgage interest rate provided. The second way, used in the analysis of this paper, follows the method in Kimberlin, Kim, and Schaefer (2014). Property taxes and insurance premia are subtracted from monthly mortgage payments, which are compared with the principal and the mortgage interest rate to infer the share of each payment that constitutes interest.

Although no AMT preference items are available in the PSID, non-AMT preference items in addition to mortgage interest include miscellaneous medical expenses and charitable contributions.

3.7.1.2 Wealth Data

The PSID has several asset categories that allow for estimation of net wealth of a household. I follow the procedure in Bosworth and Anders (2008) and aggregate housing wealth net of remaining principal, other real estate assets, vehicles, business/farm wealth, stocks, retirement assets, other accounts, and other wealth. Counted against this gross measure is an undifferentiated "debt" category until 2011, at which point credit card debt, student loans, medical bills, legal bills, loans from relatives, and other debt are separated into individual categories.

3.7.1.3 Tax Data

I feed the income aggregates into TAXSIM9, which returns a host of variables, including taxes owed, AMT tax, and statutory marginal tax rates at the federal and state level. I increment earned income and construct empirical marginal tax rates that compare well with these statutory rates, so the analysis above uses the empirical rates.

The tax price of charitable donations is calculated by augmenting the non-AMT preference item category and determining the empirical change to total taxes owed. I do this from the existing level of donations for each filer as well as from a counter-factual tax simulation in which the filer has no charitable donations. The two distinct empirical tax prices reflect the last-dollar price and first-dollar price, respectively, of giving. In addition, the total tax owed in the counter-factual simulation in which no charitable donations are counted towards non-AMT preference items doubles as the tax from which a first-dollar measure of discretionary income can be obtained.

3.7.2 Test for Tax Salience in Panel Data Specification

As in the case of the cross-sectional specification, I define disposable income Y^d as $Y^d_{it} = Y_{it} - T_{it}$ and ex-AMT disposable income Y^{d^*} as $Y^{d^*}_{it} = Y_{it} - T^*_{it}$, where $T_{it} = T^*_{it} + AMT_{it}$. The terms $\Delta \ln Y^d_{it}$ and $\ln Y^d_{it}$ in Equation 4 can be expressed as

$$\alpha_1 \Delta \ln Y^d_{it} = \alpha_1 \ln \frac{Y^d_{it}}{Y^d_{it-1}} = \alpha_1 \ln \frac{Y^{d^*}_{it}}{Y^d_{it-1}} \frac{Y^d_{it}}{Y^{d^*}_{it}} = \alpha_1 (\Delta \ln Y^{d^*}_{it} + \ln \frac{Y^d_{it}}{Y^{d^*}_{it}}) \tag{3.11}$$

and

$$\alpha_2 \ln Y_{it}^{d} = \alpha_2 \ln Y_{it}^{d^*} \frac{Y_{it}^{d}}{Y_{itt}^{d^*}} = \alpha_2 (\ln Y_{it}^{d^*} + \ln \frac{Y_{it}^{d}}{Y_{it}^{d^*}}). \tag{3.12}$$

Performing the same decomposition on P_{it} and combining terms, I estimate the following equation:

$$\begin{aligned} \ln g_{it} = \delta_i + \alpha_1 \Delta \ln Y_{it}^{d^*} + \alpha_2 \ln Y_{it}^{d^*} + (\alpha_1 + \alpha_2) \ln \frac{Y_{it}^{d}}{Y_{it}^{d^*}} \\ + \alpha_3 \Delta \ln P_{it}^{*} + \alpha_4 \ln P_{it}^{*} + (\alpha_3 + \alpha_4) \ln \frac{P_{it}}{P_{it}^{*}} + X_{it}\beta + \varepsilon_{it} \end{aligned} \tag{3.13}$$

From Equation 3.10, then, the tests for tax and marginal rate salience are $\gamma_1 = \alpha_1 + \alpha_2$ and $\gamma_2 = \alpha_3 + \alpha_4$, as stated above.

Figure 3.1: Frequency of Charitable Contributions by Income Class
PSID, 1999 to 2013

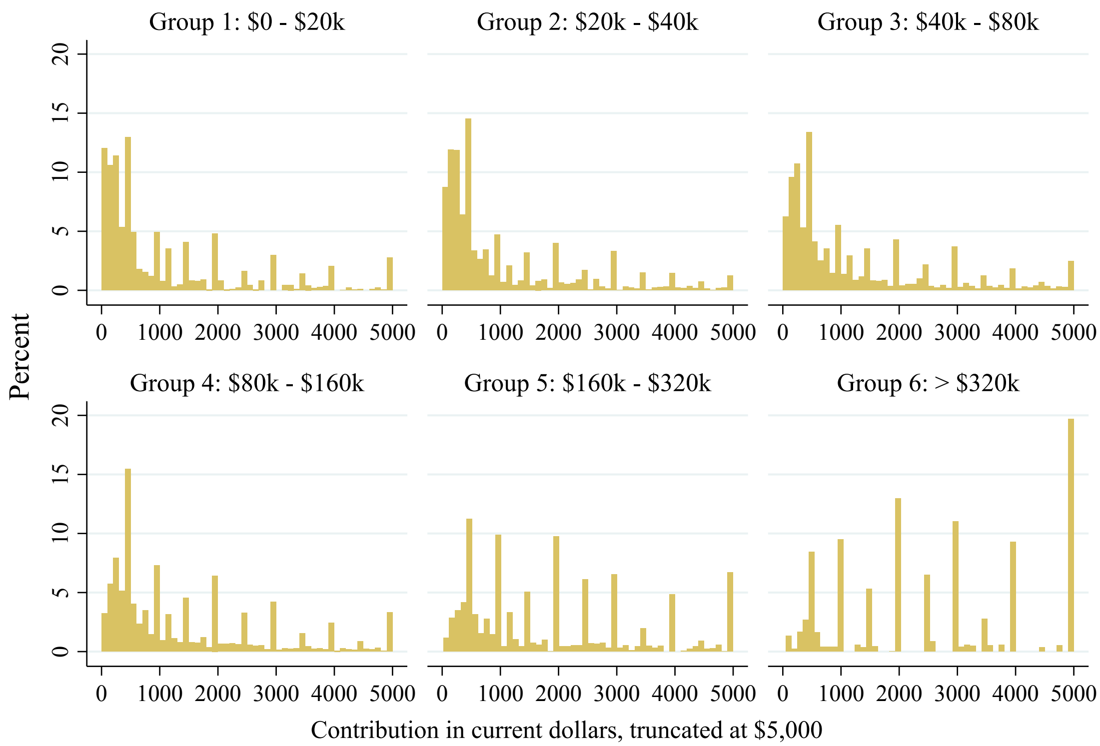

Note: Non-contributors are excluded. Bin size is $100.

Figure 3.2: Rate Effect of the AMT on Charitable Contributions

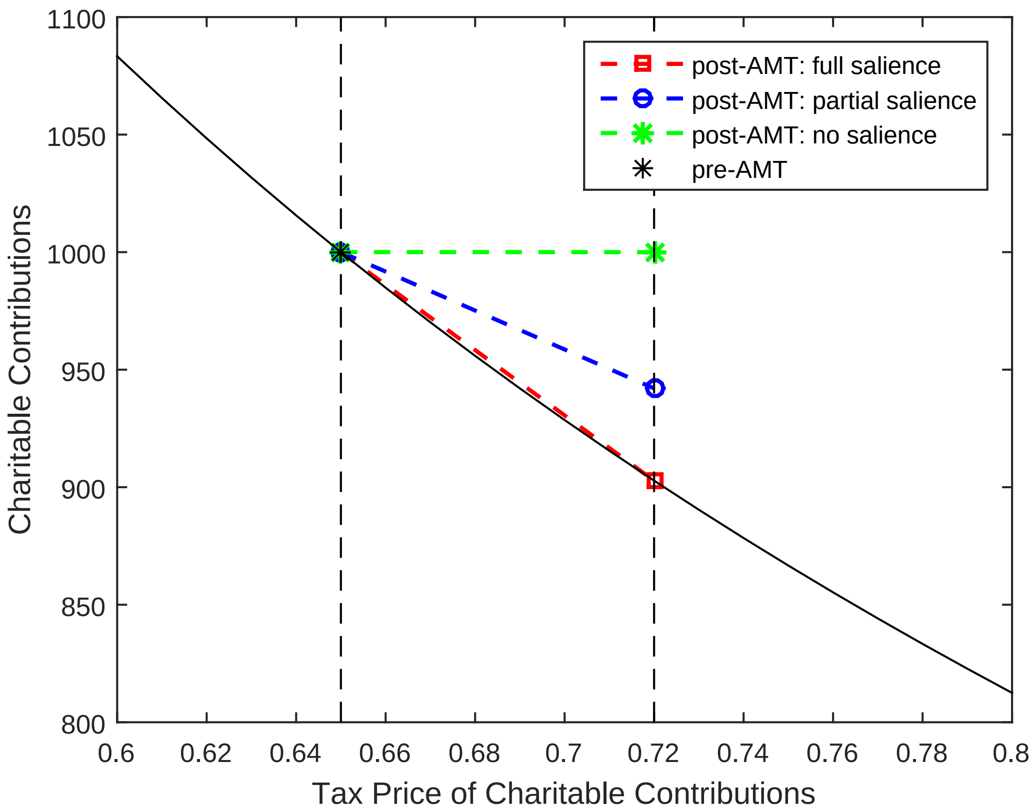

Figure 3.3: Income and Rate Effects of the AMT on Charitable Contributions

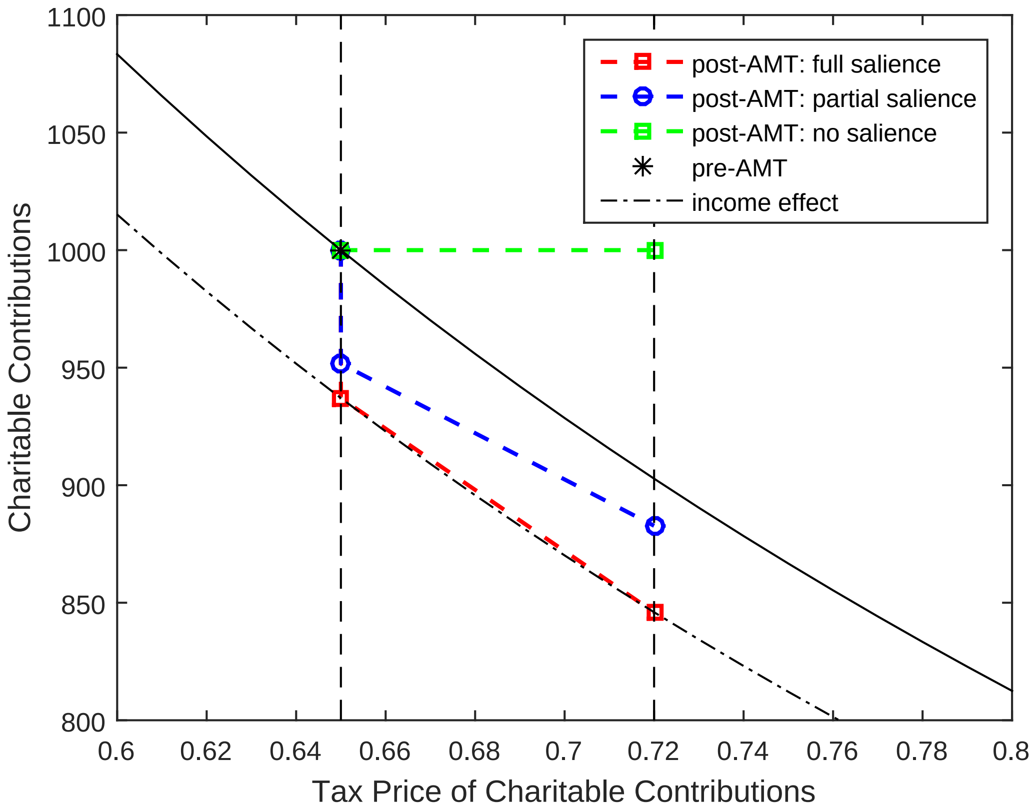

Table 3.1: Descriptive Statistics
PSID, 1999 to 2013

Variable	All Tax Units	Itemizing Tax Units with Charitable Contribution	AMT Filers
Number of Units	65,151	11,637	733
Average Income	57,370	117,078	271,867
Median Income	38,770	90,581	255,000
Average Wealth	281,623	637,686	1,869,790
Median Wealth	48,500	231,000	768,000
Average Tax Liability Owed	14,776	35,270	94,672
Median Tax Liability Owed	6,943	24,932	88,694
Marginal Tax Rate - Federal	12.5	23.0	33.0
Marginal Tax Rate - State	3.2	5.0	6.1
Marginal Tax Rate - FICA	14.1	12.1	5.2
Marginal Tax Rate - All	27.4	35.8	40.6
Average Charitable Contribution, Itemizers	2,797	3,441	4,294
Proportion Itemizing and Making Charitable Contribution	30.9%	--	69.4%
Share of Units who are Married	48.2%	72.2%	86.2%
Average Age of Head	50.23	50.73	49.44
Average Number of Dependent Children	0.57	0.69	1.29

Notes: All numbers are weighted by PSID longitudinal family weights except number of units. Marginal tax rates are for earned income.

Table 3.2: Charitable Contribution, Itemizing Status, and AMT Status by Income Group PSID, 1999 to 2013

Gross Income Category	<$20,000	$20,000-$40,000	$40,000-$80,000	$80,000-$160,000	$160,000-$320,000	>$320,000
Number of Units	20,620	14,973	16,960	9,988	2,161	435
Itemize Taxes	2,172	3,260	7,637	7,233	1,846	378
Have Charitable Contribution	1,120	2,047	5,735	6,050	1,612	327
Share with Charitable Contribution	7.6%	17.4%	38.8%	63.7%	75.2%	75.2%
Have AMT Liability	0	0	16	73	470	174
Share with AMT Liability	0.0%	0.0%	0.1%	0.7%	22.5%	38.8%
Average Liability Among AMT Filers	--	--	1,215	1,242	1,991	4,962
Among All Filers:						
Average Charitable Contribution	163	317	821	1,936	3,602	9,108
Average Among Contributors	2,151	1,822	2,117	3,039	4,791	12,109
Median Among Contributors	600	520	900	1,200	2,100	5,000
Among AMT Filers:						
Average Charitable Contribution	--	--	94	623	2,879	7,522
Average Among Contributors	--	--	736	1,084	4,071	10,280
Median Among Contributors	--	--	820	500	2,500	5,000

Notes: Results weighted by PSID longitudinal family weights except for number of units, number of units who itemize taxes, number of units with a charitable contribution, and number of units with an AMT liability. Shares are weighted. Charitable contributions are recorded only for filers who itemize deductions.

Table 3.3: Building a Cross-Sectional Model of Charitable Giving
PSID, 1999 to 2013

	OLS	Instrument for first dollar giving price	Wealth and Gift	Variables Vary by Income Class
	(a)	(b)	(c)	(d)
$ln(P_{it})$	-0.61	-0.59	-0.60	-0.68
	(0.17)	(0.18)	(0.20)	(0.21)
$ln(Y^d_{it})$	0.79	0.57	0.46	--
	(0.03)	(0.04)	(0.04)	--
$ln(W_{it})$			0.13	0.13
			(0.01)	(0.01)
$ln(Gift_{it})$			0.01	0.01
			(0.01)	(0.01)
$D_1 ln(Y^d_{it})$				0.25
				(0.18)
$D_2 ln(Y^d_{it})$				0.41
				(0.10)
$D_3 ln(Y^d_{it})$				0.36
				(0.08)
$D_4 ln(Y^d_{it})$				0.33
				(0.08)
$D_5 ln(Y^d_{it})$				0.32
				(0.08)
$D_6 ln(Y^d_{it})$				0.34
				(0.10)

Notes: classes for gross income are 0-20k, 20-40k, 40-80k, 80-160k, 160k-320k, and >320k. Other controls in regressions include marriage dummy; number of children; age of head of household; years of education of head of household; dummy variable indicating that household head is black; and dummy variable indicating the household head is any other nonwhite, nonhispanic ethnic or racial group. In the regression specifications with separate income classes, these variables are interacted with dummies corresponding to which income class the tax unit belongs.

Table 3.4: Cross-Sectional Model of Charitable Giving, Including AMT Indicator Variable PSID, 1999 to 2013

	Constant Elasticity for Giving by Income	Elasticity for Giving Varies by Income Class
	(a)	(b)
D^{AMT}_{it}	-0.09	-0.19
	(0.07)	(0.08)
$ln(P_{it})$	-0.64	-0.76
	(0.20)	(0.22)
$ln(Y^d_{it})$	0.47	--
	(0.04)	--
$D_1\, ln(Y^d_{it})$		0.26
		(0.18)
$D_2\, ln(Y^d_{it})$		0.42
		(0.10)
$D_3\, ln(Y^d_{it})$		0.37
		(0.08)
$D_4\, ln(Y^d_{it})$		0.34
		(0.08)
$D_5\, ln(Y^d_{it})$		0.33
		(0.08)
$D_6\, ln(Y^d_{it})$		0.33
		(0.10)

Notes: Sample size equal to 7,634 records. Classes for gross income are 0-20k, 20-40k, 40-80k, 80-160k, 160k-320k, and >320k. Other controls in regressions include marriage dummy; number of children; age of head of household; years of education of head of household; dummy variable indicating that household head is black; dummy variable indicating the household head is any other nonwhite, nonhispanic ethnic or racial group; the log of wealth; and the log of gifts (including inheritances). In the regression specification with separate income classes, these variables are interacted with dummies corresponding to which income class the tax unit belongs.

Table 3.5: Panel Model of Charitable Giving with Fixed Effects, Including AMT Indicator Variable
PSID, 1999 to 2013

	Use Only Current-Period Tax Data		Incorporate Year-over-Year Changes in Tax Data	
	Constant Elasticity for Giving by Income	Elasticity for Giving Varies by Income Class	Constant Elasticity for Giving by Income	Elasticity for Giving Varies by Income Class
	(a)	(b)	(c)	(d)
D^{AMT}_{it}	0.03	0.03	0.01	0.02
	(0.06)	(0.06)	(0.06)	(0.06)
$\Delta ln(P_{it})$	--	--	0.20	0.18
	--	--	(0.13)	(0.13)
$ln(P_{it})$	-0.04	-0.13	-0.44	-0.49
	(0.25)	(0.27)	(0.31)	(0.33)
$\Delta ln(Y^d_{it})$	--	--	-0.02	-0.02
	--	--	(0.03)	(0.03)
$ln(Y^d_{it})$	0.32	--	0.30	--
	(0.06)	--	(0.07)	--
$D_1 ln(Y^d_{it})$		-0.05		0.08
		(0.19)		(0.23)
$D_2 ln(Y^d_{it})$		0.21		0.16
		(0.10)		(0.11)
$D_3 ln(Y^d_{it})$		0.25		0.22
		(0.08)		(0.10)
$D_4 ln(Y^d_{it})$		0.26		0.24
		(0.08)		(0.10)
$D_5 ln(Y^d_{it})$		0.29		0.25
		(0.08)		(0.09)
$D_6 ln(Y^d_{it})$		0.33		0.34
		(0.09)		(0.10)

Notes: sample size of 7,634 using only current-period tax data, 6,429 incorporating year-over-year changes in tax data. Other controls in regressions include marriage dummy; number of children; age of head of household; years of education of head of household; dummy variable indicating that household head is black; dummy variable indicating the household head is any other nonwhite, nonhispanic ethnic or racial group; the log of wealth; and the log of gifts (including inheritances). 395 of 733 AMT filer year records are first-time filers, while 216 filer year records correspond to non-AMT taxpayers who owed AMT liability in the prior survey.

Table 3.6: Cross-Sectional Model of Charitable Giving, Including AMT Salience PSID, 1999 to 2013

	Coefficient are Uniform Across Income Levels		Non-Price Variables Vary by Income Class	
	IV Base Case	AMT Decomposition	IV Base Case	AMT Decomposition
	(a)	(b)	(c)	(d)
$ln(P^*_{it})$	-0.60	-0.82	-0.68	-0.85
	(0.20)	(0.21)	(0.21)	(0.22)
$ln(P_{it}/P^*_{it})$	--	-3.86	--	-3.83
	--	(1.08)	--	(1.16)
$ln(Y^{d*}_{it})$	0.46	0.45	--	--
	(0.04)	(0.04)	--	--
$ln(Y^d_{it}/Y^{d*}_{it})$	--	-3.95	--	1.08
	--	(4.56)	--	(4.94)
$D_1\, ln(Y^{d*}_{it})$			0.25	0.26
			(0.18)	(0.18)
$D_2\, ln(Y^{d*}_{it})$			0.41	0.42
			(0.10)	(0.10)
$D_3\, ln(Y^{d*}_{it})$			0.36	0.37
			(0.08)	(0.08)
$D_4\, ln(Y^{d*}_{it})$			0.33	0.34
			(0.08)	(0.08)
$D_5\, ln(Y^{d*}_{it})$			0.32	0.33
			(0.08)	(0.08)
$D_6\, ln(Y^{d*}_{it})$			0.34	0.32
			(0.10)	(0.10)
Chi-SquaredTests				
Full Price Salience		0.00		0.01
No Price Salience		0.00		0.00
Full Income Salience		0.33		0.88
No Income Salience		0.39		0.83
Full Joint Salience		0.01		0.00
No Joint Salience		0.00		0.00

Notes: Sample size equal to 7,634 records. Classes for gross income are 0-20k, 20-40k, 40-80k, 80-160k, 160k-320k, and >320k. Other controls in regressions include marriage dummy; number of children; age of head of household; years of education of head of household; dummy variable indicating that household head is black; dummy variable indicating the household head is any other nonwhite, nonhispanic ethnic or racial group; the log of wealth; and the log of gifts (including inheritances). In the regression specifications with separate income classes, these variables are interacted with dummies corresponding to which income class the tax unit belongs.

Table 3.7: Panel Model of Charitable Giving with Fixed Effects, Including AMT Salience
Coefficients are Uniform Across Income Classes
PSID, 1999 to 2013

	Use Only Current-Period Tax Data		Incorporate Year-over-Year Changes in Tax Data	
	IV Fixed Effects	AMT Decomposition	IV Fixed Effects	AMT Decomposition
	(a)	(b)	(c)	(d)
$\Delta ln(P^{*}_{it})$	--	--	0.20	0.20
	--	--	(0.13)	(0.13)
$ln(P^{*}_{it})$	-0.07	0.00	-0.45	-0.36
	(0.24)	(0.27)	(0.29)	(0.33)
$ln(P_{it}/P^{*}_{it})$	--	0.36	--	0.53
	--	(0.86)	--	(0.93)
$\Delta ln(Y^{d*}_{it})$	--	--	-0.02	-0.02
	--	--	(0.03)	(0.03)
$ln(Y^{d*}_{it})$	0.32	0.33	0.30	0.31
	(0.06)	(0.06)	(0.07)	(0.07)
$ln(Y^{d}_{it}/Y^{d*}_{it})$	--	0.76	--	2.57
	--	(3.27)	--	(3.39)
Chi-Squared Tests				
Full Price Salience		0.62		0.37
No Price Salience		0.68		0.57
Full Income Salience		0.89		0.50
No Income Salience		0.82		0.45
Full Joint Salience		0.88		0.62
No Joint Salience		0.91		0.70

Notes: sample size of 7,634 using only current-period tax data, 6,429 incorporating year-over-year changes in tax data. Other controls in regressions include marriage dummy; number of children; age of head of household; years of education of head of household; dummy variable indicating that household head is black; dummy variable indicating the household head is any other nonwhite, nonhispanic ethnic or racial group; the log of wealth; and the log of gifts (including inheritances).

Table 3.8: Panel Model of Charitable Giving with Fixed Effects, Including AMT Salience
Non-Price Variables Vary Across Income Classes
PSID, 1999 to 2013

	Use Only Current-Period Tax Data		Incorporate Year-over-Year Changes in Tax Data	
	IV Fixed Effects	AMT Decomposition	IV Base Case	AMT Decomposition
	(a)	(b)	(c)	(d)
$\Delta ln(P^{*}_{it})$	--	--	0.18	0.18
	--	--	(0.13)	(0.13)
$ln(P^{*}_{it})$	-0.17	-0.15	-0.52	-0.46
	(0.26)	(0.28)	(0.31)	(0.33)
$ln(P_{it}/P^{*}_{it})$	--	-0.03	--	0.22
	--	(0.96)	--	(1.02)
$\Delta ln(Y^{d*}_{it})$	--	--	-0.03	-0.03
	--	--	(0.03)	(0.03)
$ln(Y^{d}_{it}/Y^{d*}_{it})$	--	-0.27	--	0.84
	--	(3.61)	--	(3.74)
$D_1 ln(Y^{d*}_{it})$	-0.05	-0.05	0.08	0.09
	(0.19)	(0.19)	(0.23)	(0.23)
$D_2 ln(Y^{d*}_{it})$	0.20	0.21	0.16	0.16
	(0.10)	(0.10)	(0.11)	(0.11)
$D_3 ln(Y^{d*}_{it})$	0.25	0.25	0.22	0.22
	(0.08)	(0.09)	(0.10)	(0.10)
$D_4 ln(Y^{d*}_{it})$	0.26	0.26	0.24	0.24
	(0.08)	(0.08)	(0.10)	(0.10)
$D_5 ln(Y^{d*}_{it})$	0.28	0.29	0.25	0.25
	(0.08)	(0.08)	(0.09)	(0.09)
$D_6 ln(Y^{d*}_{it})$	0.33	0.33	0.34	0.34
	(0.09)	(0.09)	(0.10)	(0.11)
Chi-Squared Tests				
Full Price Salience		0.89		0.57
No Price Salience		0.98		0.83
Full Income Salience		0.87		0.89
No Income Salience		0.94		0.82
Full Joint Salience		0.96		0.84
No Joint Salience		1.00		0.96

Notes: sample size of 7,634 using only current-period tax data, 6,429 incorporating year-over-year changes in tax data. Classes for gross income are 0-20k, 20-40k, 40-80k, 80-160k, 160k-320k, and >320k. Other controls in regressions include marriage dummy; number of children; age of head of household; years of education of head of household; dummy variable indicating that household head is black; dummy variable indicating the household head is any other nonwhite, nonhispanic ethnic or racial group; the log of wealth; and the log of gifts (including inheritances). In all regression specifications these variables are interacted with dummies corresponding to which income class the tax unit belongs.

Table 3.9: First-Stage Regression Results for Instruments
PSID, 1999-2013

	Coefficient of Correlation	Coefficient	T-Statistic	R-Squared of First Stage
	(a)	(b)	(c)	(d)
Among all itemizers:				
$ln(P_{it})$	0.971	0.995	447.8	0.944
		(0.002)		
$ln(Y^{d}_{it})$	1.000	1.002	5112.7	1.000
		(0.000)		
$ln(P^{*}_{it})$	0.972	0.994	451.2	0.945
		(0.002)		
$ln(Y^{d*}_{it})$	1.000	1.002	5152.2	1.000
		(0.000)		
$ln(P_{it}/P^{*}_{it})$	0.958	0.968	367.1	0.919
		(0.003)		
$ln(Y^{d}_{it}/Y^{d*}_{it})$	0.994	0.963	1015.9	0.989
		(0.001)		
Among AMT filers:				
$ln(P_{it})$	0.959	0.964	78.8	0.920
		(0.012)		
$ln(Y^{d}_{it})$	0.999	1.007	519.3	0.998
		(0.002)		
$ln(P^{*}_{it})$	0.949	0.950	69.8	0.900
		(0.014)		
$ln(Y^{d*}_{it})$	0.999	1.008	570.6	0.998
		(0.002)		
$ln(P_{it}/P^{*}_{it})$	0.943	0.948	65.8	0.889
		(0.014)		
$ln(Y^{d}_{it}/Y^{d*}_{it})$	0.988	0.986	147.9	0.976
		(0.007)		

Notes: Results are given for primary instrument of each endogenous variable. Sample restricted to itemizers. The first two variables correspond to equation 1, while the latter four correspond to equation 3.

Table 3.10: Measurement Error in a Cross-Sectional Model of Charitable Giving
PSID, 1999-2013

Data in PSID:	
$ln(P_{it})$	-0.60
	(0.20)
$ln(Y^d_{it})$	0.46
	(0.04)
Measurement error in income (modest):	
$ln(P_{it})$	-0.55
	(0.20)
$ln(Y^d_{it})$	0.47
	(0.04)
Measurement error in income (significant):	
$ln(P_{it})$	-0.45
	(0.20)
$ln(Y^d_{it})$	0.47
	(0.04)
Measurement error in giving:	
$ln(P_{it})$	-0.71
	(0.20)
$ln(Y^d_{it})$	0.43
	(0.04)

Notes: controls in regressions include marriage dummy; number of children; age of head of household; years of education of head of household; dummy variable indicating that household head is black; and dummy variable indicating the household head is any other nonwhite, nonhispanic ethnic or racial group. In the panel for modest measurement error in income, random error with mean of 0 and standard deviation of $1,000 was added to wages. In the panel for significant measurement error in income, random error with mean of 0 and standard deviation of $10,000 was added to wages. In the panel for measurement error in charitable giving, random error with mean of 0 and standard deviation of $100 was added to donations, conditional on the tax filer having a positive charitable donation.

BIBLIOGRAPHY

BIBLIOGRAPHY

Akerlof, G., W. Dickens, and G. Perry (1996): "The Macroeconomics of Low Inflation," *Brookings Papers on Economic Activity*, 1, 1–76.

Altonji, J. and P. Devereux (2000): "The Extent and Consequences of Downward Nominal Wage Rigidity," *Research in Labor Economics*, 19, 383–431.

Altonji, J. and N. Williams (1997): "Do Wages Rise with Job Seniority? A Reassessment," *NBER Working Paper*, 6279.

Auerbach, A. J. and Y. Gorodnichenko (2012): "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 4, 1–27.

Bakija, J. and B. Heim (2011): "How Does Charitable Giving Respond to Incentives and Income? New Estimates from Panel Data," *National Tax Journal*, 64, 615–650.

Barattieri, A., S. Basu, and P. Gottschalk (2010): "Some Evidence on the Importance of Sticky Wages," *NBER Working Paper*, 16130.

Barrett, K. S., A. M. McGuirk, and R. Steinberg (1997): "Further Evidence on the Dynamic Impact of Taxes on Charitable Giving," *National Tax Journal*, 50, 321–334.

Barro, R. J. (1977): "Long-term Contracting, Sticky Prices, and Monetary Policy," *Journal of Monetary Economics*, 3, 305–316.

Bewley, T. (1999): *Why Wages Don't Fall During Recessions*, Harvard University Press.

Bils, M., Y. Chang, and S.-B. Kim (2014): "How Sticky Wages in Existing Jobs Can Affect Hiring," *NBER Working Paper*, 19821.

Blanchard, O. and R. Perotti (2002): "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output," *Quarterly Journal of Economics*, 117, 1329–1368.

Blatter, M., S. Muehlemann, and S. Schenker (2012): "The Costs of Hiring Skilled Workers," *European Economic Review*, 56, 20–35.

Bosworth, B. P. and S. Anders (2008): "Saving and Wealth Accumulation in the PSID, 1984-2005," *Center for Retirement Research Working Paper 2008-2*.

Burman, L. E., W. G. Gale, J. Rohaly, and B. H. Harris (2003): "The Individual AMT: Problems and Potential Solutions," *Tax Policy Center Discussion Paper No. 5*.

Burnside, C., M. Eichenbaum, and J. D. Fisher (2004): "Fiscal shocks and their consequences," *Journal of Economic Theory*, 115, 89–117.

Calvo, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383–398.

Card, D. and D. Hyslop (1997): "Does Inflation "Grease the Wheels of the Labor Market"?" in *Reducing Inflation: Motivation and Strategy*, ed. by C. Romer and D. Romer, University of Chicago Press.

Chetty, R., A. Looney, and K. Kroft (2007): "Salience and Taxation: Theory and Evidence," *NBER Working Paper*, 13330.

——— (2009): "Salience and Taxation: Theory and Evidence," *American Economic Review*, 99, 1145–1177.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005): "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of Political Economy*, 113, 1–45.

Christiano, L. J., M. Eichenbaum, and M. Trabandt (2014): "Understanding the Great Recession," *NBER Working Paper*, 20040.

Clotfelter, C. T. (1985): *Federal Tax Policy and Charitable Giving*, University of Chicago Press.

Daly, M. C. and B. Hobijn (2013): "Downward nominal wage rigidities bend the phillips curve," *Federal Reserve Bank of San Francisco Working Paper*.

Dickens, W. T., L. Goette, E. L. Groshen, S. Holden, J. Messina, M. E. Schweitzer, J. Turunen, and M. Ward (2007): "How Wages Change: Micro Evidence from the International Wage Flexibility Project," *Journal of Economic Perspectives*, 21, 195–214.

DiNardo, J., N. M. Fortin, and T. Lemieux (1996): "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica*, 64, 1001–1044.

Edelberg, W., M. Eichenbaum, and J. D. Fisher (1999): "Understanding the Effects of a Shock to Government Purchases," *Review of Economic Dynamics*, 2, 166–206.

Ehrlich, G. and J. Montes (2015): "Wage Rigidity and Employment Outcomes: Evidence from Administrative Data," *Unpublished Working Paper.*

Elsby, M. W. L. (2009): "Evaluating the Economic Significance of Downward Nominal Wage Rigidity," *Journal of Monetary Economics*, 56, 154–189.

Elsby, M. W. L., D. Shin, and G. Solon (2013): "Wage Adjustment in the Great Recession," *NBER Working Paper*, 19478.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000): "Optimal monetary policy with staggered wage and price contracts," *Journal of Monetary Economics*, 46, 281–313.

et al., C. G. (1995): "Estimating Potential Output, Output Gaps and Structural Budget Balances," *OECD Economics Department Working Papers*, 152.

Fatas, A. and I. Mihov (2001): "The effects of Fiscal Policy on Consumption and Employment: Theory and Evidence," *CEPR Discussion Papers 2760.*

Feenberg, D. (1987): "Are Tax Price Models Really Identified: the Case of Charitable Giving," *National Tax Journal*, 40, 629–633.

Fehr, E. and L. Götte (2005): "Robustness and Real Consequences of Nominal Wage Rigidity," *Journal of Monetary Economics*, 52, 779–804.

Feldstein, M. and A. Taylor (1976): "The Income Tax and Charitable Contributions," *Econometrica*, 44, 1201–1222.

Finkelstein, A. (2009): "E-ZTax: Tax Salience and Tax Rates," *Quarterly Journal of Economics*, 124, 969–1010.

Groshen, E. and M. Schweitzer (1999): "Identifying Inflation's Grease and Sand Effects in the Labor Market," in *The costs and benefits of price stability*, ed. by M. Feldstein, University of Chicago Press, 273–308.

Haefke, C., M. Sonntag, and T. V. Rens (2013): "Wage rigidity and job creation," *Journal of Monetary Economics*, 60, 887–899.

Hosios, A. (1990): "On the Efficiency of Matching and Related Models of Search and Unemployment," *The Review of Economic Studies*, 57, 279–298.

Jones, J. B. (2002): "Has fiscal policy helped stabilize the postwar U.S. economy?" *Journal of Monetary Economics*, 49, 709–746.

Jordá, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161–182.

Kahn, S. (1997): "Evidence of Nominal Wage Stickiness from Microdata," *American Economic Review*, 993–1008.

Karlan, D. and J. A. List (2007): "Does Price Matter in Charitable Giving? Evidence from a Large-Scale Natural Field Experiment," *American Economic Review*, 97, 1774–1793.

Katz, L. and A. Krueger (1999): "The High-Pressure U.S. Labor Market of the 1990s," *Brookings Papers on Economic Activity*, 1, 1–65.

Kaur, S. (2012): "Nominal Wage Rigidity in Village Labor Markets," *Mimeo, Harvard University.*

Kimberline, S., J. Kim, and L. Shaefer (2014): "An updated method for calculating income and payroll taxes from PSID data using the NBER's TAXSIM, for PSID survey years 1999 through 2011," *Mimeo.*

Knoppik, C. and T. Beissinger (2009): "Downward nominal wage rigidity in Europe: an analysis of European micro data from the ECHP 1994?2001," *Empirical Economics*, 36, 321–338.

Lebow, D., R. E. Saks, and B. A. Wilson (1999): "Downward Nominal Wage Rigidity: Evidence from the Employment Cost Index," *Finance and Economics Discussion Series 99/31.*

Lebow, D., D. Stockton, and W. Waschter (1995): "Inflation, Nominal Wage Rigidity, and the Efficiency of Labor Markets," *Finance and Economics Discussion Series 94/95.*

List, J. A. (2011): "The Market for Charitable Giving," *Journal of Economic Perspectives*, 25, 157–180.

Mankiw, N. G. and M. Weinzierl (2011): "An Exploration of Optimal Stabilization Policy," *Brookings Papers on Economic Activity*, 209–249.

McFadden, D. (1989): "A method of simulated moments for estimation of discrete response models without numerical integration," *Econometrica*, 57, 995–1026.

McLaughlin, K. (1994): "Rigid Wages?" *Journal of Monetary Economics*, 34, 383–414.

——— (1998): "Are Wage Changes Skewed Away from Nominal Wage Cuts?" *Mimeo.*

Mertens, K. and M. O. Ravn (2011): "Understanding the aggregate effects of anticipated and unanticipated tax policy shocks," *Review of Economic Dynamics*, 14, 27–54.

Monacelli, T., R. Perotti, and A. Trigari (2010): "Unemployment fiscal multipliers," *Journal of Monetary Economics*, 57, 531–553.

Mountford, A. and H. Uhlig (2009): "What are the Effects of Fiscal Policy Shocks?" *Journal of Applied Econometrics*, 24, 960–992.

Nickell, S. and G. Quintini (2003): "Nominal Wage Rigidity and the Rate of Inflation," *Economic Journal*, 113, 762–781.

Peloza, J. and P. Steel (2005): "The Price Elasticities of Charitable Contributions: A Meta-Analysis," *Journal of Public Policy & Marketing*, 24, 260–272.

Perotti, R. (1999): "Fiscal Policy in Good Times and Bad," *Quarterly Journal of Economics*, 114, 1399–1436.

——— (2007): "In Search of the Transmission Mechanism of Fiscal Policy," *NBER Macroeconomics Annual*, 22, 169–249.

——— (2011): "Expectations and Fiscal Policy: An Empirical Investigation," *IGIER Working Paper*, 429.

Pissarides, C. A. (2009): "The Unemployment Volatility Puzzle: is Wage Stickiness the Answer?" *Econometrica*, 77, 1339–1369.

Pissarides, C. A. and D. T. Mortensen (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *The Review of Economic Studies*, 61, 397–415.

Ramey, V. A. (2011a): "Can Government Purchases Stimulate the Economy?" *Journal of Economic Literature*, 49, 673–685.

——— (2011b): "Identifying Government Spending Shocks: It's All in the Timing," *Quarterly Journal of Economics*, 126, 1–50.

——— (2011c): "A Reply to Roberto Perotti's 'Expectations and Fiscal Policy: An Empirical Investigation'," *mimeo.*

——— (2016): "Macroeconomic Shocks and their Propagation," *NBER Working Paper*, 21978.

Ramey, V. A. and M. D. Shapiro (1998): "Costly capital reallocation and the effects of government spending," *Carnegie-Rochester Conference Series on Public Policy*, 48, 145–194.

Ricardo, D. (1887): *Letters of David Ricardo to Thomas Robert Malthus, 1810-1823*, Oxford University Press.

Romer, C. D. and D. H. Romer (2010): "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, 100, 763–801.

Shafir, E., P. Diamond, and A. Tversky (1997): "Money Illusion," *Quarterly Journal of Economics*, 112, 341–374.

Shapiro, M. D. (1986): "The Dynamic Demand for Capital and Labor," *Quarterly Journal of Economics*, 101, 513–542.

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95, 25–49.

——— (2012): "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 15, 127–148.

Shin, D. (1994): "Cyclicality of real wages among young men," *Economics Letters*, 46, 132–142.

Shin, D. and G. Solon (2006): "New Evidence on Real Wage Cyclicality within Employer-Employee Matches," *NBER Working Paper*, 12262.

Sims, C. A. (1980): "Macroeconomics and Reality," *Econometrica*, 48, 1–48.

Slemrod, J. (1989): "Are Estimated Tax Elasticities Really Just Tax Evasion Elasticities? The Case of Charitable Contributions," *Review of Economics and Statistics*, 71, 517–552.

Smets, F. and R. Wouters (2003): "An estimated dynamic stochastic general equilibrium model of the euro area," *Journal of the European Economic Association*, 1, 1123–1175.

Tauchen, G. (1986): "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, 20, 177–181.

Taussig, M. (1967): "Economic Aspects of the Personal Income Tax Treatment of Charitable Donations," *National Tax Journal*, 20, 1–19.

Taylor, J. B. (2000): "Reassessing Discretionary Fiscal Policy," *The Journal of Economic Perspectives*, 14, 21–36.

Thapar, A. (2008): "Using private forecasts to estimate the effects of monetary policy," *Journal of Monetary Economics*, 55, 806–824.

Tobin, J. (1972): "Inflation and Unemployment," *American Economic Review*, 62, 1–18.

Woodford, M. (2011): "Simple Analytics of the Government Expenditure Multiplier," *American Economic Journal: Macroeconomics*, 3, 1–35.