4-21-2021 9:30 AM

Three Essays on Individual and Household Responses to Information, Liquidity, and Policy Shocks

Brian C. Held, *The University of Western Ontario*

Supervisor: Bowlus, Audra, *The University of Western Ontario*
: Navarro, Salvador, *The University of Western Ontario*
: Lochner, Lance, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Economics

Abstract

This thesis consists of three empirical studies of heterogeneous individual and household behaviour in response to shocks that occur in the labour market and in the consumer marketplace. Chapter 2 contributes to the labour literature by examining the long-term effects of firings on subsequent labour market outcomes from an ability learning perspective. Chapter 3 contributes to the household savings and personal finance literature by studying how households navigate temporary liquidity shocks due to income disruption. Chapter 4 contributes to our understanding of online sales tax policy in the context of the rural-urban divide.

In Chapter 2, I develop a search model featuring public learning about worker ability. I then derive and empirically test a set of model implications for fired workers in which the severity of the effect depends upon the ex ante expectations of match success. I find that firings have a lasting effect on wages after controlling for experience, tenure, and unemployment. Further, this permanent effect is far more severe for university-educated workers relative to high school graduates. This is consistent with the model predictions, but is contrary to previous studies which found no evidence of learning among university graduates.

In Chapter 3, I perform a detailed analysis of various proposed channels by which households might be able to smooth consumption by comparing the behaviour of furloughed U.S. federal government employees to their unaffected counterparts working for fully-funded federal agencies over the course of the 2018-19 government shutdown, the longest in U.S. history. Using a difference-in-differences approach, I find that the shutdown-induced delay in the arrival of regular bi-weekly paychecks caused a significant reduction in household expenditures related to both consumption and recurring debt payments. The extent of this reduction is strongly correlated with the level of liquid savings households held prior to the shutdown. Among households with the least liquid savings, I fail to find evidence that households incur additional costs to smooth consumption-related expenditures (as captured by debit and credit card transactions), raising questions about how much effort from a policy standpoint should be made to specifically support consumption during temporary liquidity shocks.

In Chapter 4, I examine consumer purchasing behaviour in six states before and after the 2015 implementation of state sales taxes in Ohio and Michigan on purchases at Amazon.com, with a specific focus on identifying differential effects along major demographic dimensions. Using a difference-in-differences approach, I find that lower-income households and those living in more remote geographic areas are the most heavily impacted by a relative increase in the cost of online goods. Further, there is significant evidence of interaction between both income and geography. In terms of Amazon.com purchases alone, rural households absorb the entirety of the tax on online purchases, urban households absorb about half, and suburban households do not significantly increase total expenditures at all following the implementation of state sales tax collection on Amazon.com purchases.

Summary for Lay Audience

Disruptive events can have widely varying effects on workers and consumers depending upon an individual or household's unique circumstances. I examine three particular types of these events and how their impacts differ across various policy-relevant dimensions. First, I study the role job firings may play in revealing permanent characteristics of workers and find that university-educated workers suffer larger subsequent earnings reductions relative to high school graduates, a fact I attribute to larger downward revisions of expectations of worker ability. Next, I examine the consumer financial response to a short-term delay in the arrival of bi-weekly paychecks and show that without a sufficient level of liquid savings or some form of direct relief, households reduce spending on food and services in order to keep debt obligations such as mortgages current. Finally, I study how consumers behave after an increase in sales taxes on online goods and find that lower-income rural and urban households are the most disproportionately affected. I propose that this is due to higher transportation costs faced by these households when making visits to physical retailers. Taken together, this work is particularly useful in light of the recent COVID-19 pandemic and can help direct future research into the economic impact and the policy responses of governments worldwide.

Dedication

To my wife, Kerry, and daughters Ellamina and Calla, who have cheerfully made many sacrifices over the course of this adventure to our now-beloved Canada and back. I hope I've made you proud and I dedicate this work to you.

Acknowledgements

I would first like to acknowledge the unending support of my committee, whose acceptance of this work is among the greatest honours of my life. My supervisor, Audra Bowlus, imparted her passion for labour economics, helped ground me in the discipline, and was an incredible guide through this long journey. I am most appreciative of her intellect and flexibility as I moved among sub-fields over the course of my studies. Salvador Navarro was an invaluable resource both in the development of my programming skills and as a guide to best practices in econometrics. Lance Lochner was a role model for me in terms of how to think, ask questions, and drive to the point of anything I'll ever encounter.

I am very grateful to all of the individuals that made a difference during my time as a student in the Economics Department. My professors across all fields in my first and second years had a profound influence. In particular, Jim MacGee and Igor Livshits instilled in me a love for banking that serves me well to this day while Maria Goltsman and Greg Pavlov fascinated me with their presentations of game theory and mechanism design which have permanently impacted the way I view everything. I have fond memories of the kindness and care shown by Debra Merrifield, Yvonne Adams, Sandra Augustine, Sharon Phillips, and Jennifer Hope. To my fellow students, the camaraderie and great times in shared offices, classrooms, and conferences (in addition to the less professional settings) will never be forgotten.

Finally, I thank Dean and Gena Schantz, the Organization, and my parents for their love and support from beginning to end. Your help at critical times throughout this process has enabled me to reach this point.

Brian Held February 23, 2021
Columbus, Ohio

Contents

List of Figures

List of Tables

List of Appendices

Chapter 1

Introduction

This thesis consists of three empirical studies of heterogeneous individual and household behaviour in response to shocks that occur in the labour market and in the consumer marketplace. Chapter 2 contributes to the labour literature by examining the long-term effects of firings on subsequent labour market outcomes from an ability learning perspective. Chapter 3 contributes to the household savings and personal finance literature by studying how households navigate temporary liquidity shocks due to income disruption. Chapter 4 contributes to our understanding of online sales tax policy in the context of the rural-urban divide.

In Chapter 2, I develop a search model featuring public learning about worker ability. I then derive and test a set of interesting implications for fired workers in which the severity of the effect depends upon the ex ante expectations for match success. Learning about ability is widely recognized as a likely contributor in part to wage growth within jobs and over the career (Rubinstein and Weiss, 2006). Yet empirical evidence of its significance and effect on labour market outcomes is ambiguous. This is particularly true for workers that have higher expected ability such as university graduates and those with more experience in the labour market. The extent to which employer-initiated separations are idiosyncratic has important implications for the optimal mix of treatments related to human capital development and those intended to increase matching rates out of unemployment.

Starting with a search model of the labour market featuring heterogeneous workers and firms similar to that outlined in Postel-Vinay and Robin (2002), I relax the assumption of ex ante perfect information regarding the worker's innate ability and specify a production function that only depends on worker ability through a fixed skill requirement. The resulting equilibrium behaviour in response to learning generates predictions regarding negative labour market outcomes such as firings and long-term unemployment for low ability workers. Using data from the National Longitudinal Survey of Youth, I estimate a reduced-form empirical model that tests for the theoretical predictions. I then interpret the results within the context of the current literature on public learning about worker ability. I find that firings have a lasting effect on wages after controlling for experience, tenure, and unemployment. Further, this permanent effect is far more severe for university-educated workers relative to high school graduates. This is consistent with the model predictions, but is contrary to previous results which found no evidence of learning among university graduates.

In Chapter 3, I perform a detailed analysis of various proposed channels by which households might be able to smooth consumption by comparing the behaviour of furloughed U.S.

federal government employees to their unaffected counterparts working for fully-funded federal agencies over the course of the 2018-19 government shutdown, the longest in U.S. history. I separate the reduction in observed debt payments by furloughed workers during the shutdown into an insurance effect supporting consumption and a liquidity effect caused by the interruption in bi-weekly pay. This is an important and active policy channel in the light of global governmental responses to the COVID-19 pandemic which featured temporary debt relief on an unprecedented scale. Understanding the role of both liquid savings and debt relief in smoothing consumption provides a solid footing for evaluating and improving this particular policy response.

Using a difference-in-differences approach, I find that the shutdown-induced delay in the arrival of regular bi-weekly paychecks caused a significant reduction in household expenditures related to both consumption and recurring debt payments. The extent of this reduction is strongly correlated with the level of liquid savings households hold prior to the shutdown. The least-liquid households experience far greater declines relative to their controls compared to other affected households. In contrast, households with a level of liquid savings satisfying a common savings heuristic are able to manage the disruption relatively well. Among households with the least liquid savings, I fail to find evidence that households incur additional costs to smooth consumption-related expenditures (as captured by debit and credit card transactions), raising questions about how much effort from a policy standpoint should be made to specifically support consumption during temporary liquidity shocks. Rather, evidence of household prioritization of debt payments suggests that resources may be better spent on enabling households to keep existing obligations current.

In Chapter 4, I examine consumer purchasing behaviour in six states before and after the 2015 implementation of state sales taxes in Ohio and Michigan on purchases at Amazon.com, with a specific focus on identifying differential effects along major demographic dimensions. As states begin collecting sales taxes on all online purchases in the wake of the U.S. Supreme Court ruling of South Dakota v. Wayfair, Inc., it is important to understand the winners and losers of uncompensated changes to sales tax policy. Suburban households generally enjoy relatively high incomes, a lifestyle centered around personal transportation, and close geographic proximity to dense brick and mortar retail. In contrast, both rural and urban households suffer from a relative lack of accessible retail options for varying reasons. On average, urban households are most likely to rely on public transportation, are most affected by congestion, and have smaller living quarters and storage options. Rural households are physically distant from the large retail centers located on the outskirts of major metropolitan areas.

Using a difference-in-differences approach, I find that lower-income households and those living in more remote geographic areas are indeed the most heavily impacted by the relative increase in the cost of online goods. Further, there is significant evidence of interaction between both dimensions. In terms of Amazon.com purchases alone, rural households absorb the entirety of the tax on online purchases, urban households absorb about half, and suburban households do not significantly increase total expenditures at all following the implementation of sales tax collection on Amazon.com purchases. The rural poor have very little in the way of local retail options and their observed substitution behaviour is limited to competing online retailers. In contrast, suburban households substituting away from Amazon appear to have a rich set of options and significantly increase their expenditures at a variety of brick and mortar retailers once Amazon loses its sales tax-related pricing advantage.

Chapter 2

Learning and Employer Skill Requirements: Theory and Evidence

As time in the labour market accumulates, individual wages become increasingly correlated with plausible proxies for an unobserved, individual-specific quality defined as ability (Farber and Gibbons, 1996; Altonji and Pierret, 2001). This is consistent with a model in which agents learn about a fixed, transferable worker component of production over the course of the career. While Altonji and Pierret (2001) and others have focused upon the signaling role of education and statistical discrimination, evidence of this type of learning also furthers the original objective of Farber and Gibbons (1996) to present learning as a determinant of wage dynamics that complements the established frameworks of human capital and search.[1] In a general setting, a typical learning model is equally capable of matching observed mean wage profiles compared to other explanations but may differ in its predictions for other objects of interest such the variance of wages (Rubinstein and Weiss, 2006). These differences in predictions as well as potential policy implications continue to motivate further study of learning in the labour market.

To the extent that learning influences labour market outcomes, an additional consideration is the speed at which it occurs. If learning matters but happens very quickly, it may have little economic significance. Lange (2007) studies early career wage movements and finds that employers do learn relatively quickly, with initial expectation errors halving in about three years. This leads to an estimate of the signaling value of education accounting for less than 15% of the estimated marginal return to schooling. In a similar vein, Arcidiacono, Bayer, and Hizmo (2010) take the log wage regression that has been the standard in this literature beginning with Altonji and Pierret (2001) and estimate it separately for high school and college graduates. They fail to find any evidence of learning at all among the college educated. Specifically, this education group fails to generate the hallmark of learning in this specification: a significantly positive coefficient for the interaction of ability and experience. This leads to the conclusion that high school graduates demonstrate their ability on the job while college perfectly reveals ability.

These results introduce some doubt as to whether learning is a widespread phenomenon in the labour market or a process restricted to certain subgroups at early stages of the career.[2]

[1] For a comparison of search, human capital, and learning see Rubinstein and Weiss (2006).

[2] Kahn and Lange (2014) present some evidence to the contrary, finding that uncertainty regarding worker

However, it is important to note that wage changes are but one potential outcome of learning about a permanent component of individual productivity. Another plausible outcome of a negative signal, in particular, is the termination of a match by the firm. Firings commonly occur and are a powerful and unambiguous indication that a match that was once seen as profitable takes on, for whatever reason, an irreparably negative value to the employer. If learning does result in a firing for some workers, it is also conceivable that it affects job prospects going forward and increases the possibility of a long unemployment duration.

In this chapter I use a search model featuring public learning about worker ability to derive unique implications of these learning-related outcomes for employment and wages. I augment the typical empirical specification of the learning literature with plausible proxies for the reception of a negative signal and present evidence supporting the model using data from the National Longitudinal Survey of Youth (NLSY). I find that learning is important for both high school and college graduates and provide an alternative explanation for previous results that led to the conclusion that college graduates do not face any uncertainty regarding their productivity on the job.

The model, based upon the search framework developed in Postel-Vinay and Robin (2002), generates outcomes of learning that have not been fully accounted for in the empirical literature, features a rich yet tractable production and learning process, and has potential extensions in several directions. Firings and long-term unemployment arise through the interaction of search behaviour and learning in an environment where jobs have a fixed skill requirement and a fixed output if the worker satisfies the ability requirement. Workers update their beliefs about their ability by learning whether or not they satisfy the ability requirement of the job they currently hold. Firings occur when it is revealed that workers do not satisfy the minimum requirements of their current jobs. Long-term unemployment arises when workers learn they do not satisfy the skill requirement of a low skill job and must wait to be matched with an even less demanding job.

Employer skill requirements are an important characteristic of the model for two primary reasons. First, skill requirements combined with learning provide a natural source of firm-initiated separations, or firings. Workers failing to satisfy a job's skill requirement produce nothing and, once revealed, their retention is not profitable for the firm at any positive wage. Second, skill requirements also allow for a fundamental change to the learning process itself, avoiding the imposition of equal rates of learning across all jobs while retaining tractability. Rather than learning occurring through the repeated observation of noisy signals of a worker's ability, skill requirements within a search framework allow one to dispense with the signal noise process altogether and still feature the gradual revelation of a worker's true ability.

Most work in the learning literature has either explicitly or implicitly relied on a standard Bayesian learning process in which agents regularly receive a noisy signal of the worker's ability (for notable exceptions see, e.g., Cunha, Heckman, and Navarro, 2007; Navarro, 2011). This is commonly inferred via direct observation of output from a production process which is strictly monotonic in worker ability and subjected to random noise. In the canonical Jovanovic (1979) learning model, agents continuously learn about an idiosyncratic match quality component of production. This type of learning has been integrated into search models (e.g., Jovanovic, 1984), but learning about idiosyncratic match quality has no permanent effects by

ability has the largest impact late in the career due to corresponding sub-optimal human capital investment.

construction and so this particular specification is of no use in explaining long unemployment durations or wage stagnation across jobs. Farber and Gibbons (1996) modifies this framework so that agents publicly learn about a permanent, transferable worker ability and derives testable implications of the standard Bayesian learning model.

In contrast, this chapter diverges from the learning literature that assumes continually-arriving noisy signals of ability by making learning occur across jobs with varying skill requirements. Learning is based upon a random signal verifying that the worker satisfies the job's skill requirement, and after receiving a signal a worker must move on to another job to update beliefs. The rate of learning varies with the rate at which workers place into jobs requiring greater skills. Further, the amount learned about an individual worker's ability in any given job depends on the ex ante probability that the worker satisfies the job's skill requirement. This implies that the information content of signals is state dependent, a feature also present in Sanders (2012).

Another common factor in the prior empirical literature on learning is a production process that is strictly increasing in the worker's imperfectly-observed ability. Firings and endogenous separations in general are not as natural of an outcome of this type of production process compared to one with skill requirements, although it is possible to account for them.[3] However, it is difficult to allow for firings caused by learning to have a permanent impact in this framework due to the relative smoothness of the updating process. If wages are believed to represent all current information about a worker's ability at any time, the fact that a worker has been fired should not contain much additional information compared to the last observed wage.

In order to avoid some of these less-desirable implications of learning through repeated signals, the information content of individual signals can be made to vary with the current job in addition to existing beliefs.[4] In this chapter, I assume that jobs have a predetermined output and require only a worker with sufficient skill to complete the task. As a result, all workers that satisfy the requirement are equally productive.[5] Albrecht and Vroman (2002) specify a search and matching model of the labour market to this effect, where workers are of two skill levels and match with jobs that vary in their skill requirements. The model implies a high degree of assortative matching and even under perfect information the skill requirements that accompany job vacancies in the matching process have important implications such as exit rates from unemployment that depend on the worker's skill level (high type workers can match with low type firms but the reverse is not allowed). Uren and Virag (2011) address within-group wage inequality by integrating skill requirements into a Burdett-Mortensen framework,

[3]Jovanovic (1979) features separations due to learning about a match-specific component of production. When learning about a permanent worker component, however, any new information may also have an effect on the worker's outside option. When jobs are ranked by their contribution to output, if the declining expected ability of the worker does manage to cross a possibly declining threshold for separations there is a counterintuitive implication that the now unemployed worker only accepts new matches with jobs of a higher type than the previous job. Further, if the firm's value of a vacancy in equilibrium is zero the firm prefers to retain the worker at any positive level of output. If the wage is contained in some way, such as by a minimum wage, firings may occur but with the same counterintuitive implications for sorting.

[4]Pastorino (2015) presents one such method, featuring a two-type model with a state-dependent learning mechanism. The information value from signals received within a match about a worker's innate ability is dependent upon where each job falls within a firm-specific hierarchy.

[5]This is a common assumption in the skill-biased technical change literature (e.g., Acemoglu, 1999; Card and DiNardo, 2002).

generating similar implications using a continuous wage distribution. I further increase the granularity of the learning process by taking advantage of the Postel-Vinay and Robin (2002) framework to allow for a continuous rather than two-type ability distribution.[6]

This chapter also relates to the literature on unemployment stigma and scarring and the effects of job displacement. Gibbons and Katz (1991) present and find empirical support for a model in which a downsizing firm's decision of which workers to lay off and which workers to retain communicates information about worker quality to the market. Vishwanath (1989) characterizes optimal search behaviour from unemployment when pre-employment signals of ability are present. Firms avoid workers with long ongoing unemployment durations, inferring that the worker is of low ability. In this chapter I broaden the scope of learning and provide simultaneous explanations for these outcomes as well as wage growth. Stevens (1997) is also notable for several reasons. Against what appears to be convention in the displacement literature, Stevens includes firings in the definition of a displacement and also looks at the effects of displacement on wages across different educational groups separately using a fixed effects model. She finds that repeated job loss is a key factor in explaining the persistent effect of displacement on wages commonly documented in the literature, and once this is controlled for the long-run effects of displacement are reduced. In addition, she finds evidence that displaced workers with university degrees suffer larger earnings losses relative to those with twelve or fewer years of education. However, firings are not looked at separately and the estimated effect of a single displacement is concluded to dissipate over time.

Empirically, I find that negative labour market outcomes theoretically associated with a downward revision of beliefs about worker ability play a significant role in explaining wages. Firings for cause and time spent in unemployment have permanently negative effects after controlling for experience, tenure, and experience gained since the most recent firing. In addition, the relatively larger penalty for firings experience by college graduates is consistent with model predictions under plausible assumptions about the distribution of ability in each population. This difference reflects how learning manifests itself in different ways depending upon the ex ante probability that a worker satisfies the skill requirement in any given job. Fresh college graduates receive a wage premium in entry level jobs owing to a higher probability of success and receive little in the way of wage increases when it is revealed they are indeed productive. However, the few college graduates revealed to be of low ability lose this benefit of higher expectations when they are fired. These workers suffer larger post-firing wage losses compared to high school graduates that started out with a lower expected ability.

The remainder of the chapter is organized as follows. Section 2 presents the model of learning and generates new implications regarding the effect of firings on subsequent wages. Section 3 presents the data, estimates an empirical model consistent with the theoretical predictions, and analyzes the results in the context of the recent empirical learning literature. Section 4 concludes.

[6]Michaud (2018) also builds a learning process into the Postel-Vinay and Robin (2002) environment but is focused on the role of private rather than public information.

2.1 Model

This model builds upon the search with counteroffers foundation of Postel-Vinay and Robin (2002). The main premise is that workers differ in their permanent productive capacity or ability and must overcome search frictions and the lack of perfect knowledge about their true productivity to match with a job for which they are best suited. In a frictionless world with no uncertainty there would be perfect assortative matching as workers would go directly to firms requiring their exact level of ability. In this model, however, agents must learn about worker ability on the job, treating the value of the worker-firm match as an experience good (Jovanovic, 1979). Following a long line of previous work departing from Jovanovic (1979), the object of learning is permanent and transferable such that the worker's entire sequence of match values will be serially correlated.

The point at which I make a key divergence from the current learning literature is in the specification of the production function. I assume that output is not directly observable and only depends on worker ability through the worker's satisfaction of a job-specific skill requirement. This is consistent with a world in which the duties and output of each job are predetermined but also dependent on outside factors such as coworker performance that impair the immediate, direct inference of the worker's contribution. Firms are primarily concerned with qualifications and whether the worker has the requisite ability to produce in a particular job. High ability workers must be placed in more demanding jobs in order to produce more than low ability workers.

In this model, employed workers receive randomly arriving signals that reveal with certainty whether or not they have the requisite ability to be productive in their current job. This is a significant divergence from the widely used learning process in which ability is inferred directly via repeated observations of noisy signals. Workers use these signals to form beliefs about their own ability, which influences reservation behaviour while searching from unemployment and on the job. The primary implication of this particular signal process for learning is that workers must keep climbing the job ladder to continue refining their beliefs about their own ability.

2.1.1 Setup

Workers are assumed to have a permanent, one-dimensional innate ability γ drawn at market entry from a continuous unimodal population distribution $\Gamma(\cdot)$. Vacant jobs with skill requirement α and corresponding fixed output are continuously and exogenously distributed according to $F(\cdot)$. New matches with vacant jobs are made by both the unemployed and those working at rate λ. Employed workers transition to unemployment for reasons unrelated to productivity at rate δ and workers permanently leave the labour force at rate μ. Agents are risk neutral and have discount factor ρ.

Real output is a function of the job's minimum required ability and the ability of the worker. All agents hold beliefs regarding worker ability based upon the population ability distribution and signals of ability received while employed. A signal revealing the worker's output in the current job arrives at rate ϕ and there is at most one *informative* signal received over the duration of a match. All worker-firm pairs eventually learn whether or not their match is productive with certainty as long as they are not exogenously separated. As in Farber and Gibbons (1996), all information is assumed to be public so that there are no strategic considerations. Competing firms hold identical beliefs about any particular worker's ability. The primary concern of the firm regarding worker ability is whether or not he meets the minimum requirement of the job in question.

Let $\Theta = \{\theta_L, \theta_H\}$ represent the infimum and supremum of possible values for γ as revealed through positive and negative suitability signals over the working career. For the new worker, the initial beliefs $\Theta_0 = \{\gamma_L, \gamma_H\}$ consist of the upper and lower bounds of the ability distribution. When matched with a prospective employee, the firm uses this information to make the hiring decision. If the worker is hired, production begins immediately according to the function

$$\Psi(\gamma, \alpha) = \mathbf{1}(\gamma \geq \alpha)Y(\alpha),$$

where $Y(\alpha)$ is a strictly increasing function of the job type. This specification implies that output is zero if a worker does not meet the job's minimum ability requirement. If the worker has not previously revealed himself to be suitable for the particular job he holds, output is unknown until revealed by an information shock. In order to characterize the current value of the match, beliefs about the worker's current productivity must be formed. A worker's expected output is given by

$$\Pi(\Theta, \alpha) = E[\Psi|\Theta] = \frac{\Gamma(\theta_H) - \Gamma(\alpha)}{\Gamma(\theta_H) - \Gamma(\theta_L)} Y(\alpha) \tag{2.1}$$

if $\alpha \in (\theta_L, \theta_H)$, where the first multiplicative term is the probability that the worker has an ability exceeding the job requirement given the population distribution Γ, θ_L is the highest skill job at which the worker has been revealed to be successful, and θ_H is the lowest skill job at which the worker was unsuccessful. If $\alpha \leq \theta_L$ or $\alpha \geq \theta_H$ there is no uncertainty and output is $Y(\alpha)$ or 0, respectively. The production function requires a minor assumption to ensure that the model remains well-defined.

Assumption 2.1.1 *$Y(\alpha)$ is such that* $\lim_{\alpha \longrightarrow \theta_H} \Pi(\Theta, \alpha) = 0$.

In the limit, the growing potential output must be outweighed by the declining probability that the worker satisfies the skill requirement. Satisfaction of this condition enables the characterization of decision rules using subsets of the support of the vacancy distribution. In particular, it must be that for any given set of beliefs the expected output is initially increasing in the job type before peaking and then decreasing as the declining probability of worker suitability begins to outweigh any increase in potential output.[7]

Upon receiving a positive signal (or immediately if it is known from prior experience that the worker is suitable) there is no more uncertainty and the worker continues to produce $Y(\alpha)$ until the match is destroyed or the worker voluntarily leaves for a better job. As an example of the belief updating process, consider a worker at job α with $\theta_L < \alpha < \theta_H$. Upon receiving a signal, the current beliefs $\Theta = \{\theta_L, \theta_H\}$ are updated to $\Theta = \{\alpha, \theta_H\}$ if the worker satisfies the ability requirement or $\Theta = \{\theta_L, \alpha\}$ if the worker does not satisfy the requirement. In either case the interval of uncertainty surrounding the true ability γ narrows.

[7] An example of compatible functional forms is a normal ability distribution and exponential function for output, which have been used in numerical simulations of the model.

2.1.2 Wage Determination and Value Functions

Wages are determined according to the search with counteroffers mechanism outlined in Postel-Vinay and Robin (2002). Upon matching, firms make take-it-or-leave-it wage offers that represent permanent contracts only to be renegotiated by mutual consent. Firms effectively hold all bargaining power and workers always receive their outside option, which changes over the course of the match in response to the arrival of offers from competing firms. Unemployed workers therefore always have their wages set as a fraction of expected output such that the value of the new job is equal to the value of unemployment. Workers continue to search while employed and wages grow when they receive outside offers from competing firms that cause their current employer to renegotiate. Workers are poached if the new match has a larger surplus than the worker's current job; this allows the poaching firm to make a wage offer that the current employer cannot match. I assume that the equilibrium value of a vacancy is zero and firms will pursue any match with a positive expected output.

There are multiple cutoffs in the support of F that dictate worker decisions. This is slightly more complicated than other models because of the role that beliefs play in the ex ante determination of the value of a match; at some point the likelihood of being suitable for a high skill job is so low that it outweighs the value of the higher potential output. Therefore a "better" job to the worker is not necessarily one with a higher skill requirement. There is a lower cutoff in the job vacancy distribution below which a worker will not leave their current job because the expected output is too low due to the low job type. This is similar to other search models featuring firm heterogeneity. However, there is also an upper cutoff above which the expected output is too low due to the low probability of the worker having the requisite skill to be successful. Intervals of the support of the vacancy distribution are defined to characterize all worker decisions. These can be interpreted as reservation sets where upon matching with a firm of a type within the set, the worker triggers the between-firm competition that results in a wage increase on the current job or a transition to a new job. As beliefs about the underlying ability evolve over time, the reservation set can either expand or contract in response to new signals.

The value of employment $V(\Theta, \alpha, w)$ depends on beliefs Θ, current job α, and current wage w. Define the total job surplus $S(\Theta, \alpha)$ as the worker's value of job α net of the value of unemployment when the wage is equal to its maximum value of the entire expected output. Define $Q(\cdot)$ as a correspondence that maps worker state variables into the support of F, where $Q(\Theta, w)$ contains the job types that trigger a wage increase or transition to a more preferred job upon matching. This set can be partitioned into two subsets based upon the eventual outcome of the match. Matches with jobs in the set $Q_R(\Theta, \alpha, w)$ result in a raise via a counteroffer by the worker's current firm, while jobs in $Q_T(\Theta, \alpha, w)$ have a higher expected output that the current job and result in a job-to-job transition to the new firm.

$$
\begin{aligned}
Q(\Theta, w) &= [x \mid S(\Theta, x) > V(\Theta, \alpha, w),\ x \in F] \\
Q_R(\Theta, \alpha, w) &= [x \mid S(\Theta, x) < S(\Theta, \alpha),\ x \in Q(\Theta, w)] \\
Q_T(\Theta, \alpha, w) &= [x \mid S(\Theta, x) > S(\Theta, \alpha),\ x \in Q(\Theta, w)]
\end{aligned}
$$

Q^C is defined as the complement of $Q(\Theta, w)$ in the support of F, containing the jobs for which the expected output is too low for the match to generate a wage increase for the worker. For the unemployed, Q is constructed in a similar way but flow unemployment utility cannot be

Figure 2.1: Outcomes of New Matches by Job Type

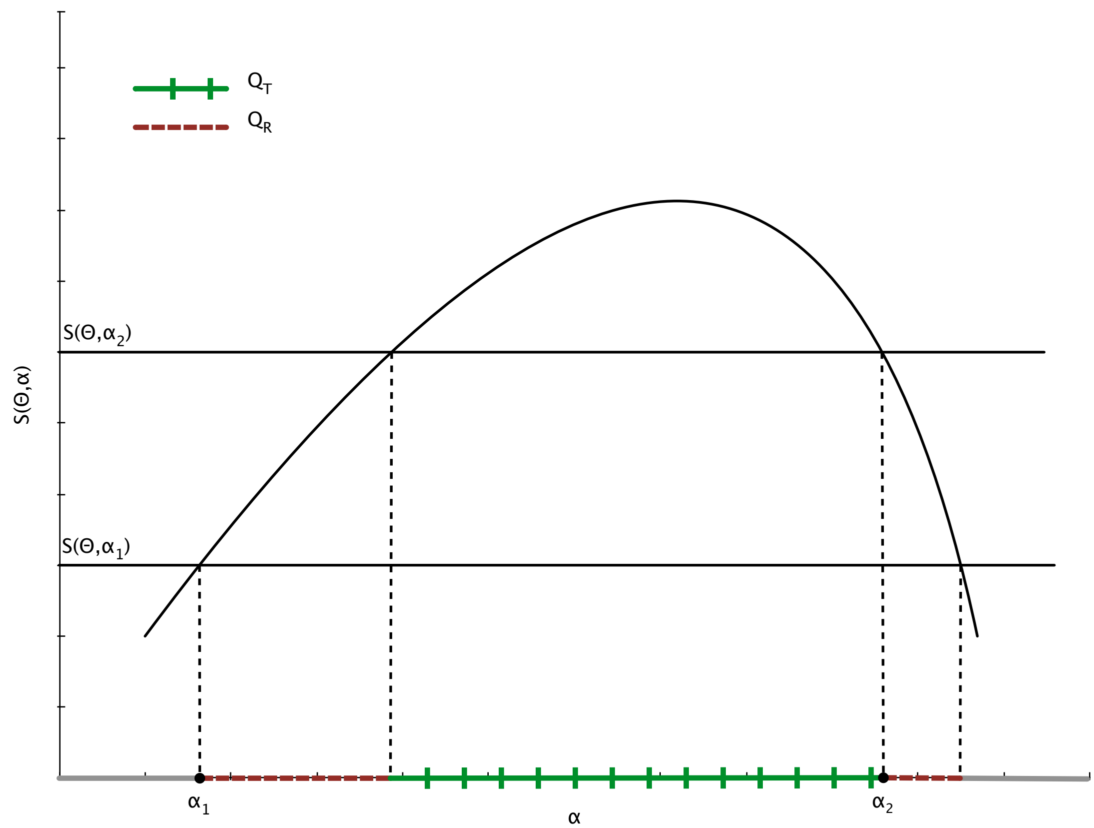

renegotiated and the set Q_R is empty. Figure 2.1 illustrates these sets for a given set of beliefs and job history. The interval on the horizontal axis denoted with long dashes corresponds to Q_R and consists of jobs with an expected output greater than the worker's previous job α_1. A match with one of these jobs results in a wage increase on the current job α_2. The interval with vertical dashes corresponds to Q_T and consists of jobs for which the expected output is greater than the worker's current job and upon matching will result in a job-to-job transition.

Unemployed workers receive a strictly positive flow benefit in unemployment that is based upon their expected ability. The value of unemployment in any period for a worker with beliefs Θ is characterized by:

$$\begin{aligned} V_0(\Theta) = b \cdot E[\gamma \mid \Theta] + \tfrac{1-\lambda-\mu}{\rho} V_0(\Theta) \\ + \tfrac{\lambda}{\rho}[\int_{Q_T} V(\Theta, x, w_0^\star) dF(x) + \int_{Q^c} V_0(\Theta) dF(x)], \end{aligned} \tag{2.2}$$

where $b \cdot E[\gamma \mid \Theta]$ is the constant value of flow utility in unemployment. Workers match with prospective employers at rate λ and if there is a positive surplus workers transition to the new job, earning a wage $w_0^\star$ that equates their new value of employment to the previous value of unemployment. All other jobs are contained in Q^C and do not result in a change to the worker's state.

The value of the worker-firm production relationship varies with job type α and the strength of beliefs regarding the worker's satisfaction of the skill requirement. Given the equilibrium value of a vacancy is zero, a firm will always want to hire a worker that has a positive probability of being productive. However, the firm must also be able to make a wage offer high enough to motivate the worker to leave their current state. This means that some matches that firms would pursue are not made, even if the worker is currently unemployed. However, since wages are paid out of expected output, workers revealed to be unproductive are immediately fired. The firm initiates the separation because the match is now unprofitable at any positive wage. The value of employment for a worker with beliefs Θ at job α with wage w satisfies

$$\begin{aligned} V(\Theta, \alpha, w) &= w + \tfrac{\delta}{\rho} V_0(\Theta) + \tfrac{1-\delta-\mu-\lambda-\phi}{\rho} V(\Theta, \alpha, w) \\ &+ \tfrac{\lambda}{\rho}[\textstyle\int V(\Theta, x', w') dF(x)] \\ &+ \tfrac{\phi}{\rho}[\Pr(\gamma > \alpha \mid \Theta) V(\Theta', \alpha, w') + \Pr(\gamma < \alpha \mid \Theta) V_0(\Theta')], \end{aligned} \tag{2.3}$$

where the final term captures the probability that a suitability signal arrives and, for a worker with $\theta_L < \alpha$, updates the beliefs to Θ' and either raises the expected output to the job's specified value or destroys the match if the worker does not satisfy the ability requirement. Wages may be renegotiated at this point if the worker's new value of unemployment exceeds the value of employment at the current wage. The second term relating to the expected value from search can be expanded to encompass the different outcomes of a match written in terms of the current state:

$$\begin{aligned} \int V(\Theta, x', w') dF(x) = &\int_{Q_R} V(\Theta, x, \Pi(\Theta, x)) dF(x) \\ &+ \int_{Q_T} V(\Theta, \alpha, \Pi(\Theta, \alpha)) dF(x) + \int_{Q^c} V(\Theta, \alpha, w) dF(x), \end{aligned}$$

where the three integrals over the previously defined subsets of the job distribution correspond to the three possible outcomes upon matching with an outside firm: receiving an offer that causes the current firm to counter the outside firm's offer; receiving an offer that motivates a transition to the new job; and receiving an offer too low to motivate a raise.

2.1.3 Model Properties

As starting wages for unemployed workers are set to give workers their value of unemployment, the expected value of the worker's ability along with model parameters completely characterize the value of unemployment as:

$$V_0(\Theta) = \frac{\rho}{\rho + \mu - 1} b \cdot E[\gamma \mid \Theta]. \tag{2.4}$$

Given the centrality of a worker's expected ability in the model, it is helpful to show how it changes in response to signals and how these changes vary with beliefs and job types. In particular, the change in beliefs can be rewritten in a form that best illustrates the varying effect of learning as a function of prior beliefs and the job type given the distribution of ability in the population. To help simplify the notation, let $\Theta_+ = \{\alpha, \theta_H\}$ and $\Theta_- = \{\theta_L, \alpha\}$ be defined as the updated beliefs after a positive and negative signal of ability, respectively. The change in the expected value of worker ability is characterized in terms of prior beliefs and the job type according to:

$$\begin{aligned} E(\gamma \mid \Theta_-) - E(\gamma \mid \Theta) &= \frac{1}{\Gamma(\alpha) - \Gamma(\theta_L)} \int_{\theta_L}^{\alpha} x \, d\Gamma(x) - \frac{1}{\Gamma(\theta_H) - \Gamma(\theta_L)} \int_{\theta_L}^{\theta_H} x \, d\Gamma(x) \\ &= \Pr(\gamma > \alpha \mid \Theta)[E(\gamma \mid \Theta_-) - E(\gamma \mid \Theta_+)] < 0 \end{aligned}$$

for negative signals, and:

$$\begin{aligned} E(\gamma \mid \Theta_+) - E(\gamma \mid \Theta) &= \frac{1}{\Gamma(\theta_H) - \Gamma(\alpha)} \int_{\alpha}^{\theta_H} x \, d\Gamma(x) - \frac{1}{\Gamma(\theta_H) - \Gamma(\theta_L)} \int_{\theta_L}^{\theta_H} x \, d\Gamma(x) \\ &= \Pr(\gamma < \alpha \mid \Theta)[E(\gamma \mid \Theta_+) - E(\gamma \mid \Theta_-)] > 0 \end{aligned}$$

for positive signals.

Each expression consists of two multiplicative terms. The first term can be interpreted as the prior probability of the counterfactual outcome (e.g., the prior probability of a positive signal for a fired worker). The second term captures the distance between the realized posterior expected ability and the counterfactual posterior expectation. For jobs in which a worker is expected to produce, a positive signal has relatively little effect on expected ability but a negative signal is highly damaging, and for the opposite case a positive signal provides a large increase while a negative signal is relatively harmless. Considered in another way, where α is relative to θ_L and θ_H in a probabilistic sense determines whether or not there will be a large effect of negative signals and a small effect of positive signals, a large effect of positive signals and a small effect of negative signals, or a similar effect for both. The width of the interval of uncertainty between θ_L and θ_H bounds the magnitude of the changes. This aspect of the model is important when comparing the effects of firings between populations with different ability distributions.

The prospect of learning in any job has a strictly non-negative effect on worker utility. While the expected value of worker ability can change significantly in response to signals, it can be shown that in expectation the post-signal value of unemployment is equal to the current value of unemployment. If the effect of a signal only changed the worker's outside option, workers would be indifferent to changes in the arrival rate of signals. However, the wage setting process generates a benefit to the worker receiving a positive signal by increasing the match surplus, raising the wage growth potential in the current job and the value of outside offers motivating a job-to-job transition.

As with Postel-Vinay and Robin (2002), it is helpful to use the match surplus when describing features of the model. As stated earlier, the match surplus is the worker's value of the job when the wage is the entire expected output net of the value of unemployment. In this state, the contribution of the arrival of outside offers to the worker's value drops out because the worker is already making the maximum wage possible in the current job, cannot receive any further raises, and will receive the same value upon moving to a different job. With this in mind, combining Equations (2.3) and (2.4) and simplifying results in the following characterization of the match surplus:

$$S(\Theta, \alpha) = \frac{\rho[\Pi(\Theta, \alpha) - bE(\gamma \mid \Theta)] + \phi \Pr(\gamma > \alpha \mid \Theta) \max\{V(\Theta_+, \alpha, \Pi(\Theta, \alpha)) - V_0(\Theta_+), 0\}}{\rho + \mu + \delta + \phi - 1}, \tag{2.5}$$

where the last term captures the expected gains from the reception of a signal, if any.

Learning's contribution to the match surplus is strictly non-negative, as the worker always has the option to renegotiate wages such that they receive their value of unemployment. To the extent that this term is nonzero, it represents the gain to the worker attributed to searching from a job in which they enjoy a share of the positive surplus rather than the zero share received out of unemployment. The upper bound is governed by the total output $Y(\alpha)$ and given the initial assumptions it is assured that $S(\Theta, \alpha)$ is well defined in all states and has a limiting value of zero in the job type.

With the surplus defined, the transition set $Q_0(\Theta)$ containing all acceptable jobs to an unemployed worker can now be characterized by its endpoints α_L and α_H, where $\alpha_L < \alpha_H$ and $S(\Theta, \alpha_L) = S(\Theta, \alpha_H) = 0$. For all allowable production functions, there exists an interior surplus maximizing job in the vacancy distribution. The two endpoints of $Q_0(\Theta)$ correspond to a low skill job that is less demanding than the optimal job and a more demanding high skill job where the probability of success is sufficiently low such that its value is the same as the low skill job. Below α_L and above α_H the expected output is too low to generate a sufficient wage to motivate the worker to transition from unemployment. Given that the surplus at these points is zero, Equation (2.5) implies that α_L and α_H must satisfy:

$$\Pi(\Theta, \alpha_L) + \frac{\phi}{\rho}\Phi(\Theta, \alpha_L, \Pi(\Theta, \alpha_L)) = \Pi(\Theta, \alpha_H) + \frac{\phi}{\rho}\Phi(\Theta, \alpha_H, \Pi(\Theta, \alpha_H)) = bE[\gamma \mid \Theta] \tag{2.6}$$

where $\Phi(\Theta, x, w) = \Pr(\gamma > x \mid \Theta) \max\{V(\Theta_+, x, w) - V_0(\Theta_+), 0\}$.

For expositional clarity $\Phi(\Theta, \alpha, w)$ is defined as the expected change in the surplus from learning the worker satisfies the skill requirement of job α, as discussed above. All other reservation sets can also be defined at this time. To allow simplicity in the characterization of these sets, define $\alpha^\star(\Theta)$ as the surplus-maximizing job given a worker's beliefs. From any job, the total surplus is increasing in the direction of $\alpha^\star(\Theta)$. Thus the current job represents one endpoint of $Q_T = [\alpha'_L, \alpha'_H]$, the set of jobs motivating job-to-job transitions. The other endpoint is the job on the opposite side of $\alpha^\star(\Theta)$ that features the same surplus. Again, the endpoints are characterized by equating the surplus between the current job and a corresponding job such that:

$$\Pi(\Theta, \alpha'_L) + \frac{\phi}{\rho}\Phi(\Theta, \alpha'_L, \Pi(\Theta, \alpha'_L)) = \Pi(\Theta, \alpha'_H) + \frac{\phi}{\rho}\Phi(\Theta, \alpha'_H, \Pi(\Theta, \alpha'_H)), \tag{2.7}$$

where $\alpha'_L = \alpha$ if $\alpha < \alpha^\star(\Theta)$ and $\alpha'_H = \alpha$ otherwise. Similarly, Q_R contains all jobs that have a surplus less than the current job but large enough to trigger a wage renegotiation. A match with a job in Q_R results in a wage increase such that the worker now receives a share of surplus equal to the surplus of the outside job. For a worker starting from unemployment, the upper and lower bounds of Q_R, defined as α_L and α_H, are the same as for Q_0. In all other working states, one of the bounds must be the worker's previous job or the last job that made an offer that resulted in a wage increase. For a worker making a job-to-job transition, Q_R for the new job has the same upper and lower bounds as Q_T for the previous job. Formally, Q_R is defined as the set $[\alpha_L, \alpha'_L] \cup [\alpha'_H, \alpha_H]$, and can be thought of as the set of all jobs with a surplus greater than the worker's share of surplus in the current job that are not contained in Q_T. It must be that $S(\Theta, \alpha_L) = S(\Theta, \alpha_H) = V(\Theta, \alpha, w)$, with α_L and α_H satisfying the same condition as the endpoints to Q_T in Equation (2.7).

I now turn to the determination of wages. Workers matching out of unemployment are offered a starting wage such that they are indifferent between accepting the offer or rejecting and waiting for another match. All subsequent wage growth can be characterized from this starting point. Combining Equations (2.3) and (2.4) and noting that the current value of employment is equal to the value of unemployment allows starting wages to be characterized by:

$$w_0^\star(\Theta, \alpha) = bE[\gamma \mid \Theta] - \frac{\lambda}{\rho}[\int_{Q_R} S(\Theta, x)dF(x) + \int_{Q_T} S(\Theta, \alpha)dF(x)] - \frac{\phi}{\rho}\Phi(\Theta, \alpha, w_0^\star). \quad (2.8)$$

Again, there is a notational simplification with respect to the change in value following a positive signal. As before, the anticipated effect of learning is strictly non-negative and captures the value of searching from a job in which the worker receives a nonzero share of the surplus. While the starting wage appears on the right hand side within this term, it enters linearly such that the above equation allows for the complete determination of starting wages. If the worker has already been revealed to satisfy the ability requirement, starting wages are simply the worker's flow value of unemployment adjusted for any expected gains from search due to the arrival of outside offers. In the perfect information case the expected value of ability is constant, signals are irrelevant, and the model is reminiscent of standard search models in which starting wages are constant in expectation over the course of the career.

Once on the job, offers arrive from outside firms and the wage will be bid up if the worker matches with a job in Q_R. Wages may increase or decrease when a firm matches with a job in Q_T, which is a characteristic of the Postel-Vinay and Robin wage setting process. This happens because workers are willing to accept lower present wages if the job offers the prospect of higher wages in the future. Regardless, wages will display an overall upward trend within any spell of continuous employment. The following subsection describes the model predictions for the data regarding the effects of new information on subsequent wages and matching rates.

2.1.4 Model Predictions

The key feature of this model is the effect signals of ability have on a worker's subsequent wages and state transitions. The immediate effect of a negative signal is a firing and transition to unemployment. An important feature is the permanent effect the signal has on the worker's beliefs and the set of matches motivating state transitions by extension. A negative shift in beliefs causes the worker's value of unemployment to fall, reducing future average wages out of unemployment. Further, a negative signal also eliminates the most demanding jobs from the set of viable matches in F and potentially lowers the arrival rate of viable matches, leading to longer unemployment durations. Finally, after the worker is eventually able to match out of unemployment, wage growth may be limited for the same reason; viable matches may not arrive as often to trigger wage increases or quits to better jobs. The remainder of this section more formally defines these model predictions in preparation for the empirical analysis.

Beliefs about worker ability are a key determinant of starting wages, wage growth, and the arrival rate of matches. These beliefs dictate the total match surplus for any job, the sets of jobs that motivate raises or job changes, and the expected share of surplus the worker may have, if any, after being revealed to be productive. By extension, the signals received on the job play a critical role in the determination of future wages through the effect on beliefs about worker ability. These signals have two effects for workers that have not yet revealed themselves to be productive in their current job. There is an immediate effect on the value of the match surplus and the worker's value of unemployment, leading either to a possible wage renegotiation if the signal is positive or to a firing if the signal is negative. The second effect has a longer horizon and concerns the set of viable matches and the expected output for each. This long-term effect determines how quickly the worker can match with new jobs and how much they might eventually expect to earn.

The effect of positive signals on within-spell wage growth is governed by the probability that the worker satisfies the skill requirement. If a worker has a job for which the probability of success is high, a positive signal adds little to the maximum attainable wage at that job. This new maximum wage is simply the expected flow output of the match and the change is characterized by $Y(\alpha) - \Pi(\Theta, \alpha)$, or the full output less the worker's ex ante expected output. This risk/reward relationship is similar to the effect signals have on expected ability. Workers that have a low probability of failure receive a smaller increase to their maximum wage compared to workers facing more uncertainty. Being revealed to satisfy the demands of a very high skill job when expectations are low has the opposite effect.[8] Also, for all previously viable jobs $\alpha > \theta_L$, the expected output increases, meaning that the worker has a larger wage ceiling in any job they may switch to in the future. Again, the magnitude of this effect depends on the probabilistic distance between α, θ_L, and θ_H, where there needs to be sufficient ex ante uncertainty about a worker's ability to produce in job α to shift expected output in other jobs in a meaningful way.

Negative signals have no implications for wage growth within the current employment spell because the worker is immediately fired and must begin searching for new jobs from unemployment. To the extent that the negative signal has decreased the expectation of worker ability, the upper bound of starting wages out of unemployment is lowered by a corresponding factor. Further, the worker can no longer hope to match with any previously viable jobs having a skill requirement greater or equal to the skill requirement of the last job. All remaining jobs for which the worker is not certain to satisfy the skill requirement have a lower expected output compared to their pre-firing levels. The change in the expectation of the expected output, or maximum wage, over all jobs in response to a signal is characterized by:

$$\begin{aligned}\Delta_+ E_\alpha[\Pi(\Theta, \alpha) \mid \Theta] = &\int_{\theta_L}^{\alpha} \frac{\Gamma(x) - \Gamma(\theta_L)}{\Gamma(\theta_H) - \Gamma(\theta_L)} Y(x) dF(x) \\ &+ \frac{\Gamma(\alpha) - \Gamma(\theta_L)}{\Gamma(\theta_H) - \Gamma(\alpha)} E_{x \in (\alpha, \theta_H)}[\Pi(\Theta, \alpha) \mid \Theta] > 0\end{aligned}$$

for positive signals, and:

$$\begin{aligned}\Delta_- E_\alpha[\Pi(\Theta, \alpha) \mid \Theta] = &- \frac{\Gamma(\theta_H) - \Gamma(\alpha)}{\Gamma(\alpha) - \Gamma(\theta_L)} \int_{\theta_L}^{\alpha} \frac{\Gamma(x) - \Gamma(\theta_L)}{\Gamma(\theta_H) - \Gamma(\theta_L)} Y(x) dF(x) \\ &- E_{x \in (\alpha, \theta_H)}[\Pi(\Theta, \alpha) \mid \Theta] < 0\end{aligned}$$

for negative signals.

[8] The extreme case is unlikely to occur within the model however because the expected output is too low to exceed the worker's value of unemployment.

Positive signals increase the expectation of expected output over the jobs distribution because all jobs requiring a skill level less than the current job are now productive with probability one, and the remaining jobs for which output is uncertain have had their probabilities of success scaled up in accordance with the new information. In contrast, a negative signal decreases the expectation because it reduces the set of viable jobs to only those with skill requirements less than the current job. The measure of jobs bounded by α and θ_H now contributes nothing to the expectation while the remaining jobs with uncertain production have a downward revision of the probability of success. These properties imply that fired workers experience permanently lower wages relative to workers that have never been fired, and the magnitude of this cost is dependent on the prior beliefs about the worker's ability and the distribution of jobs in the economy.

Proposition 2.1.2 *Given any job α and initial beliefs $\Theta_0 = \{\gamma_L, \gamma_H\}$, for any two population distributions of ability Γ_1 and Γ_2 with support $[\gamma_L, \gamma_H]$ such that Γ_2 is first order stochastically dominant:*

$$0 > \Delta_{-} E_\alpha[\Pi(\Theta_0, \alpha) \mid \Theta_0, \Gamma_1] > \Delta_{-} E_\alpha[\Pi(\Theta_0, \alpha) \mid \Theta_0, \Gamma_2] \ \ \forall \alpha.$$

Proof See Appendix.

Proposition 2.1.2 states that if the probability of success for a new worker from a high ability population is uniformly higher in all jobs compared to a new worker from a low ability population, any firing is more damaging for the worker from a high ability population in terms of the change in the worker's expected maximum wage. While stochastic dominance is not invariant to truncation and prevents a full characterization of the effects of firing in all states, the above results should hold in the most likely states. In other words, unlikely events where a worker from a low ability population has been revealed to be of extremely high ability may result in a larger expected output for this worker relative to a worker from a high ability distribution in the same state. While this impairs the generalization of Proposition 2.1.2, in most cases the property holds and fired workers from a high ability distribution suffer relatively larger losses in almost all states.

Turning now to the effect of an ability signal on matching rates, an important distinction must be made between the set of jobs at which the worker has a positive expected output and the set of jobs that have a positive surplus in relation to the current state. The former is easily characterized but it is the latter that actually drives worker behaviour. Positive signals do not affect the measure of jobs for which the worker has a positive expected output but negative signals eliminate the current job and all jobs requiring greater skill from the set. What confounds a simple relationship between a negative signal and matching rates is the concurrent effect of negative signals on the worker's value of unemployment in conjunction with the job vacancy distribution. It may be the case that a negative signal actually reduces unemployment durations by causing matches to be accepted with jobs that previously had an expected output too low to induce the worker to leave unemployment. Put another way, the new information has caused the worker to take jobs that were previously not worth accepting.

While the model does not have clear implications in general regarding the hazard rate out of unemployment for fired workers, there is a clear prediction that workers fired from the lowest skill jobs have longer expected unemployment durations. The hazard rate out of unemployment is defined as $\lambda(F(\alpha_H)-F(\alpha_L))$. When a worker is fired, the value of unemployment immediately drops and the worker is now willing to take jobs that previously had a negative surplus. Yet the worker can no longer match with the previous job and all jobs requiring greater ability. Following the characterization of Q_u above, the change in the hazard rate out of unemployment can be illustrated. The proportional change in the hazard rate out of unemployment for a newly fired worker with previous $Q_u(\Theta) = [\alpha_L, \alpha_H]$ and new $Q_u(\Theta') = [\alpha'_L, \alpha'_H]$ is given by:

$$\frac{\lambda(F(\alpha'_H) - F(\alpha'_L))}{\lambda(F(\alpha_H) - F(\alpha_L))} = 1 - \frac{(F(\alpha_H) - F(\alpha'_H)) - (F(\alpha_L) - F(\alpha'_L))}{F(\alpha_H) - F(\alpha_L)}. \tag{2.9}$$

Whether the hazard rate increases or decreases is dependent upon whether the loss of jobs at the top end of the reservation set is outweighed by the gain of lower skill jobs at the low end. However, at some point, there are no less demanding jobs to be had for a worker fired from a low skill job and the hazard rate must fall. Thus workers revealed to be unsuitable in the least demanding jobs are certain to experience longer expected unemployment durations going forward.

2.2 Empirical Analysis of Model Predictions

The model makes an unambiguous prediction about the permanent effect of firings, which directly follow a negative ability signal. Firings effectively lower the ceiling for the worker in terms of subsequent wages in two ways: they remove the highest skill jobs from the set of previously viable jobs and reduce the probability of success in any remaining jobs for which the worker has not been revealed to produce with certainty. The worker, now unemployed, is in a situation where the expected output has shrunk or disappeared completely for all previously viable jobs $\alpha > \theta_L$ without any gains elsewhere. Further, a worker fired from a job with very low skill requirements must wait to be matched with an even less demanding job, implying a strong correlation between long unemployment durations and low expected ability.

The following empirical exercise is built upon distinguishing wage observations following firings and unemployment spells, two correlates of and plausible proxies for the reception of bad signals of worker ability. The estimation methodology follows the within educational group approach of Arcidiacono, Bayer, and Hizmo (2010; henceforth ABH) for two reasons. The first is to investigate whether the data are consistent with the learning model outlined in this chapter, particularly for the college pool which previously provided no evidence of learning in the ABH specification. The second is to compare the results between the two pools to see if the magnitudes of any observed effect of firings and unemployment durations match model predictions given the population differences in the distribution of ability. If college educated workers are more likely to be successful than high school graduates, Proposition 2.1.2 implies that college educated workers should suffer larger post-firing wage losses, on average, compared to high school graduates.

2.2.1 Data

The construction of the data follows ABH with the addition of select variables that are relevant to this model. The data are comprised of non-hispanic white and black men from the the main and supplementary samples of the NLSY79 and cover the waves from 1979 to 2010. Initial labour market entry occurs when the respondent has left school for the first time. In specifications where the experience measure is years in the labour market, the measure is adjusted for any subsequent return to school full time. For individual wage observations, schooling is measured by how many years have been completed to that point in time. A college graduate is defined as a worker that has completed exactly 16 years of schooling and a high school graduate is a worker with exactly 12 years of schooling completed. Actual experience is measured in years and constructed by dividing the number of weeks spent working for pay since labour market entry by 52. Only non-military jobs and jobs that are not based in the home are included in the actual experience measure.

Figure 2.2: Kernel Density Estimates of AFQT by Education for Fired and Never-Fired Workers

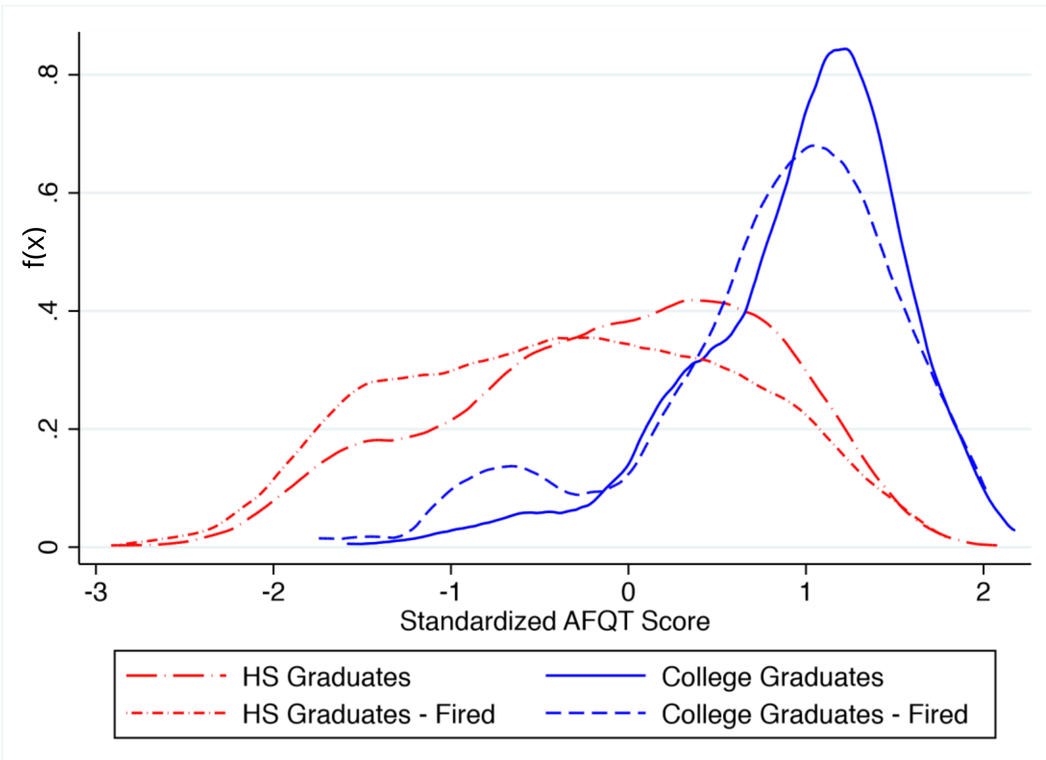

Wages are the CPS reported hourly rate of pay or equivalent measure in dollars deflated by the CPI. Only hourly wages from one to one hundred dollars are considered in the analysis. Respondent scores on the Armed Forces Qualifying Test (AFQT) are used as a proxy for ability. The AFQT score is derived from the arithmetic reasoning, word knowledge, paragraph comprehension, and numerical operations sections of the Armed Services Vocational Aptitude Battery (ASVAB). The ASVAB was administered to respondents in 1979, and the AFQT scores are normalized by age cohort.

Firings are the primary proxy for the reception of a negative signal, and I account for any reported jobs the respondent ever held that resulted in a firing. In each interview, the respondent is asked whether they are still working at any previously reported jobs. If not, they are prompted for a reason the job has ended. While the categorization of responses has changed and grown over time, firings have been accounted for in all survey years. In addition to firings, a measure of time spent in unemployment constructed in the same way as the actual experience measure is also included, taken as the number of weeks spent unemployed since market entry divided by 52.

For the purposes of this study, the AFQT is an imperfect proxy for ability. However, I operate under the assumption that the distribution of AFQT scores is roughly reflective of the true distribution of ability in each population of interest. Figure 2.2 presents kernel density estimates of the distribution of AFQT scores for each education group by whether or not individuals have ever been fired. Specifically, each education group is partitioned by whether or not the worker has ever responded that a job recorded in a preceding interview has ended due to firing. It does appear to be the case that workers reporting a firing are somewhat more likely to have low AFQT scores, but it is also clear that firings affect workers of all ability levels, as the model implies. Importantly, the density estimates follow what one might expect of the true ability distributions, with the college pool having a higher mean ability and a lower variance.

Table 2.1: Summary of Ability Indicators By Educational Attainment

	All	High School	College
Standardized AFQT			
Mean	0.172	-0.116	0.915
SD	0.973	0.915	0.623
Ever Fired (%)	33.27	38.13	19.94
Total Observed Firings (%)			
1	18.33	19.94	13.89
2	7.60	9.07	3.58
3	3.86	4.66	1.65
4	1.51	1.75	0.83
5+	1.98	2.71	0
Wage Observations			
Total	26,254	19,995	6259
Pre-Firing	2195	1866	329
Post-Firing	6714	5762	952
Individuals	2723	1996	727

Notes: AFQT scores standardized by age cohort. Educational grouping for individual ability measures is based upon the last observation in sample.

Table 2.1 presents a summary of the key ability indicators, the AFQT score and the number of firings for all individuals. There are 2723 individuals in the sample, 727 of whom are college graduates. A significant number of individuals in both educational groups have experienced a firing at some point in their careers. A majority of individuals that have ever been fired were only fired once, and as the number of total firings increase the number of observations steadily decrease.

Table 2.2: Sample Summary Statistics

	All	High School	College
Standardized AFQT			
Mean	0.356	0.095	1.042
SD	0.903	0.876	0.541
Cumulative Firings	0.372	0.443	0.185
Cumulative Years Unemployed	0.589	0.682	0.342
Log Real Wage			
Ages < 25	2.008	1.977	2.252
Ages 25-30	2.271	2.174	2.476
Ages 30-35	2.388	2.272	2.642
Ages 35-40	2.491	2.339	2.824
Ages > 40	2.630	2.431	3.053
Years Worked			
Ages < 25	2.678	2.790	1.794
Ages 25-30	5.860	6.422	4.672
Ages 30-35	9.265	9.288	9.215
Ages 35-40	12.913	12.408	14.014
Ages > 40	19.075	17.827	21.730
Urban Residence	0.762	0.732	0.843
Employed Part Time	0.033	0.038	0.021
Observations	26,254	19,995	6259

Notes: AFQT scores standardized by age cohort. Educational grouping is based upon the highest grade completed at the time of each observation. Sample is weighted by 1979 weights.

The main object of interest in this empirical analysis is the cumulative number of firings experienced up to the time of each wage observation. For a worker fired for the first time after their 1986 interview, this variable is equal to zero for all years from 1979 until their next interview. In the following interview, most likely in 1987, they report the firing and the new value for the cumulative firing variable is equal to one. An analysis of the effect of total firings on all wage observations that are not preceded by one or more firings is included as a robustness check.

I topcode the number of cumulative firings so that the small number of workers experiencing five or more firings are coded as four. This only affects the high school pool and serves to make the estimates more comparable across educational groups. The model predicts that firings occur in all jobs, and any firing from a job in which the prior probability of success was low does not significantly affect beliefs going forward. It is likely that the only difference between the workers fired four times and the workers fired five or more times is in the number of times this event has occurred. In results not tabulated here these topcoded individuals do not differ from the rest of the sample and are not driving the estimates for the effect of firings in the main specifications.

Table 2.2 presents a summary of the pooled wage data used in the estimation. The sample is weighted by the 1979 weights.[9] As might be expected, the pool of high school graduates has more total time spent in unemployment and more cumulative firings, on average, compared to college graduates.

2.2.2 The Effect of Firings on Wages

The first analysis replicates the baseline estimates of ABH and then adds the variables capturing cumulative firings and unemployment durations. To be consistent with ABH, in these specifications only observations from the years 1979-2004 for individuals with less than 13 years in the labour force are used, and potential experience is the experience measure. The sample is unweighted. All wage observations are pooled and the following log wage equation with and without the firing and unemployment variables is estimated separately for each education level:

$$\begin{aligned} w_i = \beta_0 + \beta_{AFQT} AFQT_i + \beta_{AFQT,x}(AFQT_i \times x_i) \\ + \beta_{FIRE} CUMLFIRED_i + \beta_{UNEMP} CUMLUNEMP_i + f(x_i) + \beta_X' X_i + \epsilon_i. \end{aligned} \tag{2.10}$$

The term $f(x_i)$ is a cubic polynomial of potential experience. X_i contains controls for race, a race/experience interaction, and dummies for part time status, urban residence, year, and region. Standard errors are clustered at the individual level.

Table 2.3: Pooled Log Wage Regressions, HS Graduates: ABH Specification

	(1)	(2)	(3)
Standardized AFQT	0.00779	0.0100	0.0440***
	(0.0129)	(0.0127)	(0.0119)
AFQT x Experience	0.0118***	0.0100***	0.00456***
	(0.00173)	(0.00170)	(0.000903)
Cumulative Firings		-0.0622***	-0.0600***
		(0.00975)	(0.00890)
Cumlt. Yrs. Unemployed		-0.0804***	-0.0632***
		(0.00841)	(0.00690)
Observations	11772	11772	17974
R^2	0.187	0.226	0.268

Notes: * p < 0.10, ** p < 0.05, *** p < 0.01. Clustered standard errors in parenthesis. Coefficients for cubic in potential experience, year, and demographic controls omitted.

[9] ABH do not use sample weights and so all specifications presented in the comparison table are unweighted as well. While 1979 weights are used so that the relative importance of respondents does not change within the sample, the empirical results are robust to the use of annual weights. The empirical results also do not materially change when the sample is unweighted.

Table 2.4: Pooled Log Wage Regressions, College Graduates: ABH Specification

	(1)	(2)	(3)
Standardized AFQT	0.142***	0.143***	0.129***
	(0.0354)	(0.0340)	(0.0323)
AFQT x Experience	0.00194	0.00286	0.00478
	(0.00471)	(0.00465)	(0.00305)
Cumulative Firings		-0.135***	-0.107***
		(0.0331)	(0.0281)
Cumlt. Yrs. Unemployed		-0.190***	-0.159***
		(0.0335)	(0.0274)
Observations	4112	4112	5499
R^2	0.182	0.226	0.289

Notes: * p < 0.10, ** p < 0.05, *** p < 0.01. Clustered standard errors in parenthesis. Coefficients for cubic in potential experience, year, and demographic controls omitted.

Tables 2.3 and 2.4 display the replication of the baseline estimates from ABH in Column 1, the estimates with the firing and unemployment variables are added in Column 2, and an estimation in which potential experience is not restricted to less than 13 years in Column 3 for the High School and College populations, respectively. The fundamental results of ABH relating to the effect of AFQT scores on wages are robust to the inclusion of the additional correlates of negative ability signals but it is clear that firings and unemployment durations are explaining some residual variation. The significantly negative effects of firings and unemployment on wages are an early indication in this analysis that firings may carry information about a worker's ability to produce in subsequent jobs.

I interpret the positive correlation of the AFQT/experience interaction for high school graduates shown in Table 2.3 as picking up the effect of positive signals about ability received on the job as well as some of the negative effects of firing to the extent that the AFQT score and firings are correlated. This is not a perfect proxy for the reception of positive signals because learning is subject to search outcomes, it only happens in jobs for which the worker is not certain to produce, and the signals arrive randomly. Regardless, the probability that a worker has received a signal is increasing in experience and to the extent that positive signals are significantly updating beliefs, we should see a growing correlation between ability and wages on average as workers gain experience.

Note that the estimated effect of the AFQT/experience interaction declines for the high school pool while the AFQT score becomes more informative when the sample is expanded to include more mid to late career wage observations. In a learning model, when workers reveal their ability wages stagnate at a level corresponding to their true ability. With enough post-learning observations the ability proxy becomes a better predictor of average wages as the period in which learning actually occurs constitutes less of the overall sample. That this occurs actually helps to make the case that learning is what drives the estimate of the AFQT/experience interaction rather than heterogeneous rates of human capital accumulation, the other interpretation of the interacted term. Another point against this alternative model is the lack of explanatory power the interacted term has for college graduates.

In Columns 2 and 3 of Table 2.4, college graduates suffer average wage losses of 10-13% per firing. Within a pure learning model, firings should not have a significant effect on wages if there is no uncertainty regarding the worker's ability. While this specification does not account for many confounding factors, from these initial results some simple models can be ruled out as sources of alternative explanations. A basic search model can generate a negative effect of firings by resetting the previous gains from search but will not be able to simultaneously account for a negative effect of unemployment durations. However, this specification does not do well in controlling fully for search effects and it is likely that unemployment would be picking up the effect of general separations if the data were primarily search driven. A simple model of human capital accumulation might match the effect of unemployment durations via depreciation but cannot generate an independent negative effect of firings without some sort of human capital specificity or an exogenous shock process. More needs to be done, however, to ensure that the estimated effect of firings is permanent and not driven primarily by temporary effects best explained by other models.

Now I use actual experience instead of potential experience as the experience measure and utilize all observations from 1979-2010. Actual experience is preferred as the measure of experience by the model because only workers that are employed have the potential to learn about their ability. All controls from the first analysis are retained and standard errors are clustered at the individual level.

Tables 2.5 and 2.6 contain the main empirical results for High School and College graduates, respectively. In both tables, Column 1 is a baseline regression while Column 2 adds a quadratic in tenure. The specification in Column 3 allows an additional avenue for the effect of firings to diminish over time by including a quadratic in experience accumulated since the most recent firing for all fired workers. There is a small decrease in sample size as observations with missing tenure values are dropped.

The addition of tenure in Column 2 appears to be particularly important in explaining part of the initial estimate for firings of high school graduates, yet the estimated wage penalty per firing is approximately 5% and remains a significant loss. For the college educated, the wage penalty is severe and is not as affected by the inclusion of tenure. Tenure is obviously picking up any type of firm-specific human capital, but there might be some additional search-related effects to the extent that tenure and the duration of continuous employment are correlated. In the final specification of the main analysis, I allow for any general post-firing patterns in terms of re-convergence with the pool of never fired workers. If firings do not cause permanent wage losses for most workers, it is expected that the estimated effect of a firing changes as more weight is given to the nonlinear growth in wages over the next spell of employment.

The final specification in the main tables allows for a quadratic effect of experience gained since the last observed firing. There appears to be some evidence of recovery, but the effect is not significant for college graduates. This too is consistent with model predictions. The model implies that firings are not a perfect proxy for negative signals of ability, and workers fired from very high skill jobs are not subject to a significant revision of beliefs about ability. While it is unknown whether a worker is fired from a high skill job or a low skill job, to the extent that the pool of fired workers is comprised of high-ability workers their subsequent outcomes make

Table 2.5: Pooled Log Wage Regressions, HS Graduates: Full Sample

	(1)	(2)	(3)
Cumulative Firings	-0.0708***	-0.0498***	-0.0639***
	(0.0105)	(0.0103)	(0.0112)
Cumlt. Yrs. Unemployed	-0.0682***	-0.0487***	-0.0494***
	(0.00853)	(0.00812)	(0.00818)
Standardized AFQT	0.0556***	0.0567***	0.0569***
	(0.0154)	(0.0144)	(0.0144)
AFQT x Experience	0.00232*	0.00224**	0.00226**
	(0.00119)	(0.00111)	(0.00112)
Tenure		0.0406***	0.0403***
		(0.00293)	(0.00296)
Tenure^2		-0.00112***	-0.00112***
		(0.000130)	(0.000130)
Post-Firing Experience			0.00750*
			(0.00420)
Post-Firing Experience^2			-0.000246
			(0.000151)
Experience	0.0138***	0.0137***	0.0118**
	(0.00487)	(0.00451)	(0.00467)
Experience^2	-0.000235	-0.000323**	-0.000263*
	(0.000153)	(0.000142)	(0.000147)
Constant	2.143***	2.001***	2.008***
	(0.0473)	(0.0448)	(0.0450)
Observations	19820	19611	19611
R^2	0.254	0.294	0.294

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Clustered standard errors in parentheses. Coefficients for cubic experience, year, and demographic controls omitted. Observations weighted using 1979 weights.

it appear as if there is post-firing recovery. If firings had only temporary effects on wages for all workers, it would be expected that the estimate for post-firing experience would be much stronger.

Of particular interest is the difference in magnitude of the effects between education groups. High school graduates receive relatively modest wage penalties of about 5-6% for firings compared to college graduates, who see wage losses anywhere from 11% to 17%. Given the differences between the distributions of ability for each population evident in Figure 2.2, the model provides an explanation that the ex ante strength of beliefs about worker skills determines how much the ex post beliefs change in response to new information. If college graduates start out with a high likelihood of success because of the composition of their education group, it stands to reason that there might be little evidence of learning on the job compared with high school graduates. Yet for the small number of college graduates revealed to be unsuitable for these

Table 2.6: Pooled Log Wage Regressions, College Graduates: Full Sample

	(1)	(2)	(3)
Cumulative Firings	-0.120***	-0.111***	-0.169***
	(0.0309)	(0.0304)	(0.0380)
Cumlt. Yrs. Unemployed	-0.128***	-0.111***	-0.112***
	(0.0354)	(0.0359)	(0.0359)
Standardized AFQT	0.142***	0.138***	0.137***
	(0.0347)	(0.0350)	(0.0346)
AFQT x Experience	0.00193	0.00202	0.00213
	(0.00323)	(0.00330)	(0.00325)
Tenure		0.0235***	0.0233***
		(0.00662)	(0.00655)
Tenure^2		-0.000851**	-0.000851**
		(0.000380)	(0.000367)
Post-Firing Experience			0.00842
			(0.0113)
Post-Firing Experience^2			-0.0000368
			(0.000397)
Experience	0.0581***	0.0481***	0.0475***
	(0.00884)	(0.00945)	(0.00960)
Experience^2	-0.00119***	-0.000938***	-0.000958***
	(0.000325)	(0.000345)	(0.000352)
Constant	2.105***	2.095***	2.106***
	(0.0918)	(0.0910)	(0.0904)
Observations	6164	6100	6100
R^2	0.310	0.315	0.317

Notes: * p < 0.10, ** p < 0.05, *** p < 0.01. Clustered standard errors in parentheses. Coefficients for cubic experience, year, and demographic controls omitted. Observations weighted using 1979 weights.

jobs, they experience large penalties as they now are pinned in the left tail of the ability distribution and have lost the benefit of high initial expectations enjoyed by newly minted graduates. For high school graduates, the situation is reversed to some extent. More modest initial expectations lead to learning on the job having a greater impact on average wage growth, but also to less painful adjustments to expected ability following a negative signal.

By looking at the distribution of AFQT scores by educational attainment in Figure 2.2, it can plausibly be assumed that the distribution of ability for college graduates stochastically dominates the high school ability distribution. Given this, the results are just as predicted by Proposition 2.1.2. With a low probability of failure in these types of jobs, there is not much of a monetary reward for a college graduate receiving a positive ability signal. This does not mean that learning is not present, only that the learning is basically confirming what was already believed to be the case.

Table 2.7: Pooled Log Wage Regression: Total Firings as Unobserved Heterogeneity

	High School		College	
	(1)	(2)	(1)	(2)
Total Firings	-0.0262**	-0.0455***	-0.0450	-0.0949***
	(0.0124)	(0.00819)	(0.0298)	(0.0234)
Cumlt. Yrs. Unemployed	-0.0564***	-0.0492***	-0.0865**	-0.104***
	(0.0123)	(0.00811)	(0.0418)	(0.0351)
Standardized AFQT	0.0413**	0.0562***	0.121***	0.132***
	(0.0165)	(0.0144)	(0.0367)	(0.0346)
AFQT x Experience	0.00252*	0.00230**	0.00166	0.00234
	(0.00150)	(0.00112)	(0.00371)	(0.00325)
Tenure	0.0401***	0.0394***	0.0272***	0.0227***
	(0.00342)	(0.00294)	(0.00685)	(0.00652)
Tenure^2	-0.00104***	-0.00102***	-0.00114***	-0.000831**
	(0.000143)	(0.000130)	(0.000361)	(0.000374)
Experience	0.00801	0.00674	0.0484***	0.0452***
	(0.00508)	(0.00423)	(0.0100)	(0.00896)
Experience^2	-0.000263	-0.000166	-0.000909***	-0.000725**
	(0.000162)	(0.000137)	(0.000347)	(0.000314)
Constant	2.113***	2.123***	2.123***	2.123***
	(0.0423)	(0.0361)	(0.0919)	(0.0844)
Observations	13938	19611	5168	6100
R^2	0.286	0.291	0.312	0.317

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Clustered standard errors in parentheses. Coefficients for cubic experience, year, and demographic controls omitted. Observations weighted using 1979 weights.

As a robustness check, Table 2.7 presents some results regarding the capability of the total number of firings a worker will ever experience to explain wages over the course of the career, taking Column 2 of the main results as a starting point. If the previous results are due to some type of unobserved heterogeneity, it should be indicated here. Column 1 of Table 2.7 looks for evidence that the total number of firings over the career can explain any wage differences for workers that have not yet been fired. In other words, the total number of future firings is used as an explanatory variable and only wage observations for workers that have not been previously fired are used. It appears that future firings are correlated with early wages for high school graduates, while they are completely uninformative for college graduates. These results too are consistent with the model predictions. The model predicts that fired workers experience lower wages on average relative to never-fired workers even before they are fired. This is because it is certain that the workers to be fired will never receive a positive ability signal. Given that the high school population is the group with the most within-job wage gains due to the reception of positive signals, it should not be surprising that future firings explain current wages for high school graduates but are uninformative for college graduates.

Finally, Column 2 of Table 2.7 simply replaces the cumulative number of firings with the total number of firings and replicates Column 2 of the main results. The total number of firings has a significantly negative effect for high school students, as with the previous specification, but now the college pool has a significantly large effect as well. Given the results from the previous specification, these results are not surprising for high school graduates. The total firings for college graduates also have a lower estimated effect on wages. To conclude the analysis of the robustness check, I find no evidence that unexplained heterogeneity as manifested through firings is driving the main results. Taken in isolation, the ability of future firings to predict current wages might be a problem. However, the combination of the model's predictions about pre-firing wages for populations that experience more wage growth due to learning and the results for college graduates makes a stronger case for the model.

2.3 Discussion

Starting with a simple search model featuring on the job learning about worker ability, this chapter characterizes partial equilibrium worker behaviour and wages that depend on the population distribution of ability, the distribution of job types in the economy, and individual work history. Empirically, I find that correlates of negative signals in the data have a significantly negative impact on wages. In particular, a high school graduate receives a 5% wage penalty per firing while college graduates suffer wage losses upwards of 11% per firing. Contrary to previous findings in the literature, learning appears to be important for university graduates but in a different way than for to high school graduates.

The preferred explanation is that the distribution of ability in each population plays a key role in the determination of wages over time in response to signals of ability received across different jobs over the course of the career. If college graduates have an ability distribution that is skewed toward higher values and features fewer low ability types, early-career learning for university graduates only manifests itself prominently in firings, while high school graduates learn about their ability through success as well as failure. My results confirm the findings of Arcidiacono, Bayer, and Hizmo (2010) with respect to labor market experience, but by explicitly accounting for firings and observing the major penalty a relatively small number of college graduates experience relative to fired high school graduates, it becomes clear that learning also plays a role for college graduates after all.

The findings lead in several directions. Learning has a significant role to play in explaining labour market outcomes, particularly for important yet infrequently occurring events such as displacement and long-term unemployment. The model can concurrently fit corresponding outcomes on the other side of the ability distribution such as late career wage growth and movement to the highest rungs of the job ladder. Exploring internal labour markets and job assignment within the firm in the context of this learning model would be a natural next step given how few college graduates are fired compared to high school graduates. More generally, the model can be adapted to address many different topics in which uncertainty surrounding a worker's ability to produce on the job is relevant. Possibilities include modeling the firm's vacancy decision and extensions to address macroeconomic questions.

Bibliography

D Acemoglu. Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *The Quarterly Journal of Economics*, 113(4):1055–1089, 1998.

James Albrecht and Susan Vroman. A Matching Model with Endogenous Skill Requirements. *International Economic Review*, 43(1):283–305, 2002.

Joseph G. Altonji and Charles R. Pierret. Employer Learning and Statistical Discrimination. *Quarterly Journal of Economics*, 116(1):313–350, February 2001.

P. Arcidiacono, P. Bayer, and A. Hizmo. Beyond Signaling and Human Capital: Education and the Revelation of Ability. *American Economic Journal: Applied Economics*, 2(4):76–104, 2010.

David Card and JE DiNardo. Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles. *Journal of Labor Economics*, 20:733–783, 2002.

F. Cunha, J Heckman, and S Navarro. Separating uncertainty from heterogeneity in life cycle earnings. *Oxford Economic Papers*, 57(2):191–261, December 2005.

H.S. Farber and Robert Gibbons. Learning and Wage Dynamics. *The Quarterly Journal of Economics*, 111(4):1007, 1996.

R Gibbons and L Katz. Layoffs and Lemons. *Journal of Labor Economics*, 9(4):351–380, 1991.

B. Jovanovic. Job Matching and the Theory of Turnover. *The Journal of Political Economy*, 87(5):972–990, 1979.

B. Jovanovic. Matching, Turnover, and Unemployment. *Journal of Political Economy*, 92(1): 108–122, 1984.

Lisa B. Kahn and Fabian Lange. Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures. *The Review of Economic Studies*, 81(4): 1575–1613, 2014.

Fabian Lange. The Speed of Employer Learning. *Journal of Labor Economics*, 25(1):1–35, 2007.

Amanda Michaud. An Information Theory of Worker Flows and Wage Dispersion. *American Economic Journal-Macroeconomics*, 10(8):154–183, 2018.

Salvador Navarro. Using observed choices to infer agent's information: Reconsidering the importance of borrowing constraints, uncertainty and preferences in college attendance. CIBC Working Paper 2011-8, London (Ontario), 2011.

Elena Pastorino. Job Matching Within and Across Firms. *International Economic Review*, 56 (2):647–671, 2015.

F Postel-Vinay and J.M. Robin. Equilibrium Wage Dispersion with Worker and Employer Heterogeneity. *Econometrica*, 70(6):2295–2350, 2002.

Yona Rubinstein and Yoram Weiss. Post Schooling Wage Growth: Investment, Search and Learning. *Handbook of the Economics of Education*, 1, 2006.

Carl Sanders. Skill Uncertainty, Skill Accumulation, and Occupational Choice. 2012 Meeting Papers 633, Society for Economic Dynamics, 2012.

Ann Huff Stevens. Persistent Effects of Job Displacement: The Importance of Multiple Job Losses. *Journal of Labor Economics*, 15(1):165, January 1997.

Lawrence Uren and Gabor Virag. Skill Requirements, Search Frictions, and Wage Inequality. *International Economic Review*, 52(2):379–406, 2011.

Tara Vishwanath. Job Search, Stigma Effect, and Escape Rate from Unemployment. *Journal of Labor Economics*, 7(4):487–502, 1989.

Chapter 3

Debt as Consumption Insurance: Evidence from the 2018-19 U.S. Federal Government Shutdown

Households have repeatedly shown a lack of consumption smoothing in empirical studies of income disruption (Ganong and Noel 2017). Surprisingly, this occurs even when the interruption is predictable and has no effect on permanent income (Baker and Yannelis 2017). One common explanation for this is that there exists a segment of "hand-to-mouth" households that are credit-constrained. Rather than behaving in a myopic, sub-optimal fashion, it has been suggested that these households are utilizing less-obvious sources of liquidity found at the margins of the typical household's monthly debt servicing.[1] In this chapter, I perform a detailed analysis of various proposed channels by which households might be able to smooth consumption by comparing the behaviour of furloughed U.S. federal government employees to their unaffected counterparts working for fully-funded federal agencies over the course of the longest government shutdown in U.S. history that began in December 2018.

My key contribution is in separating the reduction in observed debt payments by furloughed workers with little liquid savings during the shutdown into an insurance effect supporting consumption and a liquidity effect caused by the interruption in biweekly pay. This is achieved by making a second comparison of furloughed households with payments due during the shutdown but prior to the actual liquidity shock (the first missed paycheck) to furloughed households with payments due after the liquidity shock but prior to the announced resolution of the shutdown. Further, broad differences in relief policies across classes of debt from most accommodating to least (auto loans, student loans, and mortgages, respectively)[2,3,4] provide valuable clues regarding the short-term benefit of temporary debt relief distinct from any other inherent

[1]Over a longer time horizon, Blundell, Pistaferri, and Preston (2008) and other structural approaches find that households are more effectively able to smooth consumption. Commault (Forthcoming) resolves this divergence between structural and event-based estimates by establishing that model features such as past shocks and hand-to-mouth behaviour are correlates of log-consumption growth, which results in depressed structural estimates if left uncorrected.

[2]https://www.autonews.com/finance-insurance/auto-lenders-offer-payment-relief-us-government-workers-during-shutdown

[3]https://www.chicagotribune.com/business/ct-biz-shutdown-mortgage-finances-20190111-story.html

[4]https://blog.ed.gov/2019/01/federal-employees-manage-student-loans-government-shutdown/

consumption-smoothing benefits of debt.

In general, I fail to find evidence that households incur additional costs to smooth their consumption-related expenditures (as captured by debit and credit card transactions), raising questions about how much effort from a policy standpoint should be made to specifically support consumption during temporary liquidity shocks. Rather, evidence of household prioritization of debt payments suggests that resources may be better spent on enabling households to keep existing obligations current. Another contribution is an empirically-informed partitioning of households based upon levels of liquid savings prior to the shutdown that yields an interesting finding which lends some additional empirical support to existing theories of earmarking and mental accounting.[5] Households that have at least three months of income available in liquid savings, a common savings heuristic, show little response to the delay in compensation and provide a second reliable point of reference for the behaviour of households facing greater liquidity constraints. The role of liquid savings compared to credit card usage also relates to a broader literature on theoretical frameworks of the relationship between debt and low-yield savings.[6]

This paper is most closely related to Baker and Yannelis (2017) and Gelman et al. (2018). Both papers examine behaviour during the 2013 U.S. federal government shutdown, an almost identical event which lasted 15 days. While each shutdown is a rare opportunity to study household behaviour during a liquidity shock independent from any change in permanent income, in 2013 affected workers still received a paycheck during the shutdown of an amount about 25-40% less than usual. In contrast, the 2018-2019 shutdown studied in this paper presents an opportunity to study how the various smoothing channels highlighted by the previous literature perform over a longer period of relative illiquidity.

On December 22, 2018 the U.S. federal government shut down, putting many federal employees into a de facto furlough. At the onset, there was a great deal of uncertainty regarding how much income these workers would ultimately be entitled to and when it would be received. The shutdown became the longest in U.S. history before ending on January 25, 2019 after 35 days. In contrast to other shutdowns in recent history, affected federal employees missed an entire paycheck in its midst. The 2018-2019 shutdown did however follow convention in that affected employees were entitled to back pay upon its resolution.

Figure 3.1 shows an index of historical search activity on Google for terms containing "government shutdown" in the months leading up to the shutdown through December 17, the Monday preceding the start of the shutdown. The first signs of political conflict came in September when a potential October 1 shutdown was averted after a spending bill was signed into law on September 28. This bill funded the government through early December with a combination of full-year appropriations for some government agencies and a continuing resolution for the remainder. This led to a period of relative calm lasting until late November. A second continuing resolution for unfunded federal agencies was passed on December 6, extending the window for negotiations through December 21. No further agreement was forthcoming and the government went into a partial shutdown on December 22. The shutdown affected the FDA, NASA,

[5] Savings vehicles specifically earmarked for emergency use are particularly effective in promoting adoption (Dupas and Robinson 2013) and may play an important role in helping households weather liquidity shocks such as the shutdown studied in this paper. Lusardi et al. (2011) establish a "pecking order" of liquidity sources for households dealing with a financial shock starting with liquid savings.

[6] See Gross and Souleles (2002), Choi et al. (2003), and Telyukova and Wright (2008).

Figure 3.1: Google Search Trends

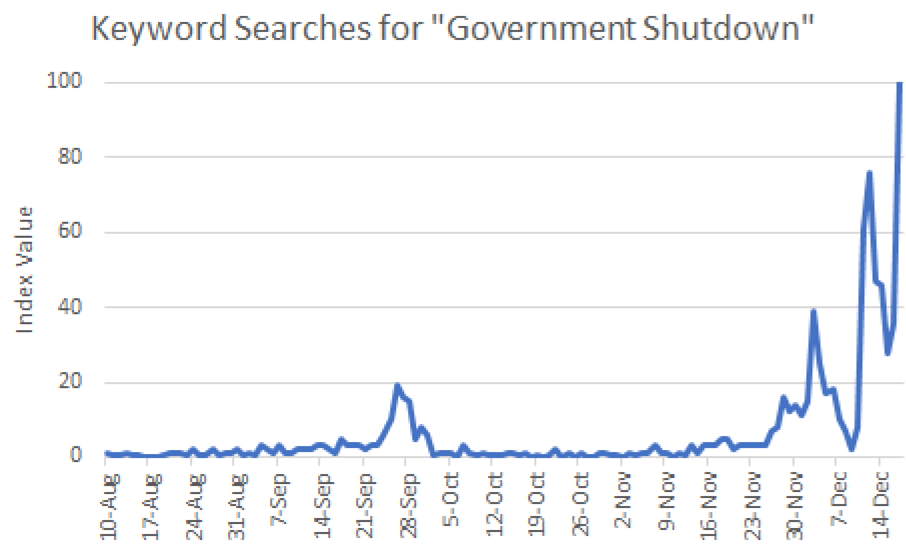

TSA, FBI, federal prisons, FAA, National Park Service, Coast Guard, and FAA, among other federal agencies.

Approximately 800,000 federal employees were directly impacted; some were furloughed and some (such as airport security officers) were deemed “essential” and required to work without pay until the resolution of the shutdown. These employees missed their first paycheck on January 11, 2019. Furloughed employees had received a bit of relief the previous day when Congress almost unanimously passed a bill entitling them to lost wages as soon as the furlough ended. A three-week funding measure was agreed to on Friday, January 25, ending the shutdown on the day most affected workers missed their second consecutive paycheck. Payments quickly went out the following week and the budget matter was resolved for good on February 14.

3.1 Data

This paper uses account and transaction data from August 2018 through April 2019 for retail banking customers residing in Ohio, Michigan, Pennsylvania, Indiana, Kentucky, and West Virginia. Starting with a sample of several hundred thousand active checking account customers for a single mid-sized regional bank, bank-administered checking account, credit card, and debit card transactions are classified based upon various bank statement transaction description fields.

Payroll is identified through bank statement transaction descriptions such as “DFAS - CLEVELAND FED SALARY” that are highly persistent over time. Recurring income is identified at the household level by observing repeated and regular payments associated with individual statement descriptions. U.S. federal government wage payments studied in this paper contain variations of “FED” and “SAL” in the description field. Households with federal employees subject to furlough or delayed compensation for essential work (the “furloughed

Table 3.1: Age and Monthly Income Summary Statistics by Furlough Group

	Not Furloughed	Furloughed
Avg. Age	45.1	44.5
Avg. Tenure	136.7	139.0
Avg. Income	\$3,849	\$4,334
Avg. Gov. Income	\$3,244	\$3,615
N	6,708	1,305

Notes: Summary of monthly data from October 2018 through April 2019 by furlough status. Income approximated by total recurring electronic (ACH) deposits.

popluation" henceforth) are identified by an interruption to their regular biweekly pay during the federal shutdown followed by an abnormally large direct deposit containing back pay observed upon its resolution. Furloughed employees receive their last paycheck at the very beginning of January and miss their next biweekly paycheck. This missed pay is not made up until the last week of January.

Debit Mastercard and credit card purchases are classified based upon bank statement descriptions that are associated with each individual debit card transaction. This detailed information about transactions for debit and credit cards issued by the bank is sent from the electronic payment card processors. Typically this information is made available to the issuing banks for routine purposes such as creating account statements and administering category-specific bonus reward programs. Merchants must provide business information to the payment processor, which is in turn provided to the banks so that individual purchases can be identified on the account statement by the consumer. The contact information includes a verified merchant name and location information such as a store identifier and/or physical address. Most retailers have store level identifiers on all transactions as well as separate descriptions for purchases made via phone or on the website. Merchant category codes allow for transactions to be classified by retail sector, such as grocers or electronics stores.[7]

Outgoing payments are also able to be classified based upon the statement description as mortgage, credit card, or other installment loan payments using a similar method to the identification of paychecks. Descriptions containing text strings such as "MTG" or "MORTGAGE" help identify mortgages specifically. Other recurring installment loans are identified by finding payees to whom a substantial portion of households make fixed regular payments. External credit card payments are identified via the account number associated with an outgoing bill payment or scheduled automatic debit.

Table 3.1 displays a comparison of the furloughed and non-furloughed populations in October 2018. This month was selected as a suitable representative of pre-shutdown attributes as the last full month before talk of a shutdown was renewed toward the end of November. The two populations are very similar in terms of age and customer tenure, but the furloughed population has higher income from government and non-government sources. One cause of

[7] For a validation of the consistency of card transaction data and store scanner data see Kaplan et al. (2020).

Figure 3.2: Normalized Income

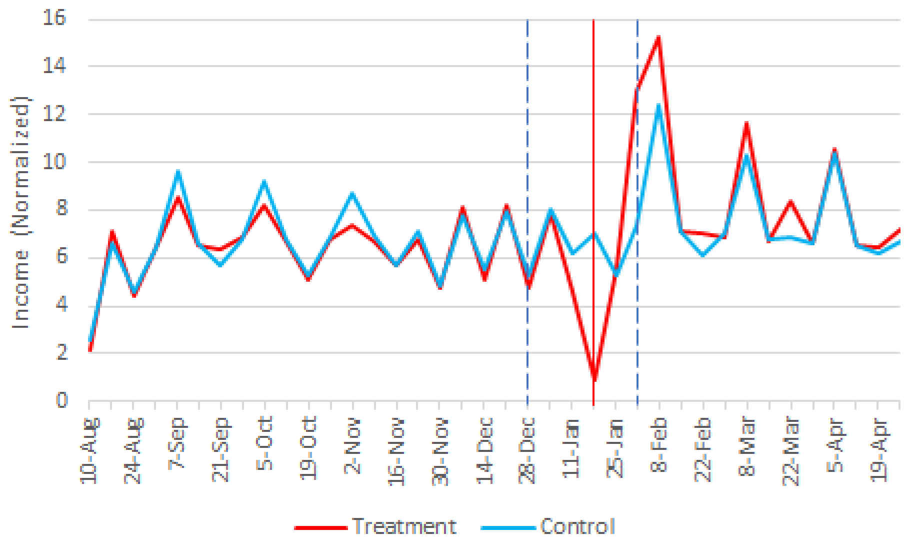

this difference is the U.S. Postal Service, a major quasi-government employer with a very large number of entry level positions that was not affected by the shutdown.

A difference-in-differences approach featuring weekly household aggregates and saturated time effects similar to that of Gelman et al. (2018) is taken. The control group consists of the federal employees that maintained their regular biweekly pay over the course of the shutdown. While the distribution of employment across the various federal agencies differs due to the full-year funding status determined by the September 2018 spending bill, the control group broadly follows the same biweekly pay date schedule as the population affected by the shutdown and also accounts for broader seasonal effects. The weekly observations of outcome variables of interest such as income and expenditure are normalized within each household to reflect days of average spending. Variation from week to week or across treatment and control groups is therefore expressed in terms of days of average income or expenditure. To reduce weekly volatility and better reflect income from other sources, the weeks on the federal biweekly pay schedule and the off weeks with no federal paycheck are normalized separately.

Figure 3.2 shows the normalized weekly income for the treatment and control groups across the entirety of sample period from August 2018 through April 2019. The first dashed blue vertical line marks the week containing the beginning of the shutdown period. The red vertical line marks the week of the first missed paycheck. The second dashed blue vertical line corresponds to the week employees returned to work and began to receive back pay for previous missed paychecks. The treatment and control groups follow the same weekly patterns over the course of the sample period except for the shutdown and weeks immediately following as the affected workers received back pay.

Figure 3.3 presents a similar view of normalized weekly expenditures for the treatment and control groups across the entirety of sample period from August 2018 through April 2019. To effectively account for the slight level differences between the treatment and control groups, I follow Gelman et al. (2018) and normalize by average days of expenditure. Again, both groups have similar week-to-week patterns in the months leading up to the shutdown. Spending

Figure 3.3: Normalized Debit/Credit Card Expenditures

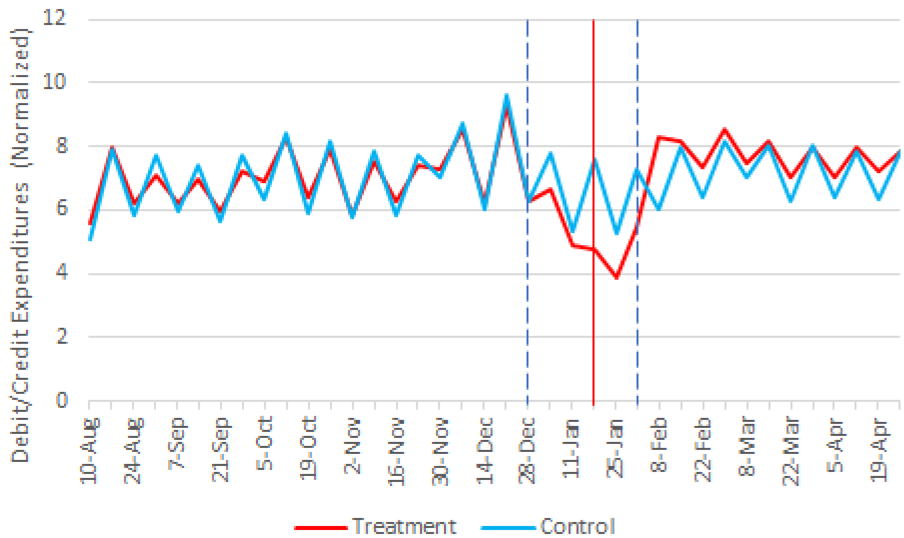

falls immediately for the treatment group during the first full week of the shutdown and does not recover until the first week of February, the first full week following the repayment of the missed paycheck on January 11. Note that holiday spending in the weeks leading up to the shutdown does not appear to be affected for the treatment group. The controls appear to be effectively capturing the strong seasonal effects at play in December through January as consumer spending rises and then retreats after the new year. In addition, the variation in spending between weeks containing pay dates and the off-weeks is immediately apparent for both the treatment and control groups.

3.2 Responses to the Shutdown

The primary objective of this analysis is to identify the methods and to what degree households affected by a temporary delay in wage income are able to smooth their expenditures. Households may draw down existing liquid deposit balances, use available credit, or delay making regularly scheduled debt payments in order to temporarily fund consumption while waiting to receive a delayed paycheck. To estimate the impact of the liquidity shock for affected federal workers, I adopt the specification proposed in Gelman et al. (2018):

$$y_{i,t} = \delta_t \times Week_{i,t} + \beta_t \times Week_{i,t} \times Shut_i + \epsilon_{i,t} \tag{3.1}$$

for each dependent variable y of interest. In each case, i indexes individual households and t indexes time in weeks. $Week_{i,t}$ is a set of indicators for each week contained in the sample data, and $Shut_i$ is the indicator for a banking household that is affected by the federal shutdown. The coefficients of interest are β_t, which identify how the federal employees affected by the shutdown in particular differ from other federal employees over the course of the sample period.

The following sections examine how affected government employees' income, retail spending, and electronic bill payments were impacted by the shutdown. These variables are all nor-

Figure 3.4: Differences in Income and Total Credit/Debit Expenditures

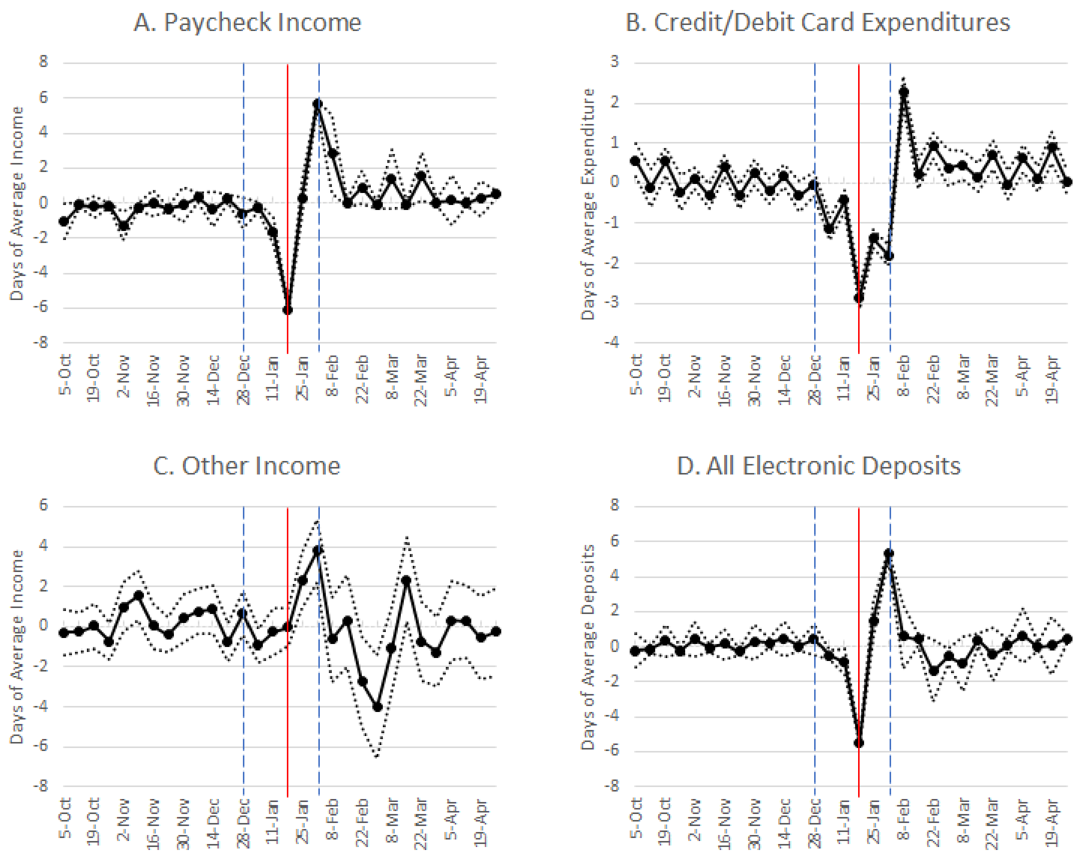

malized at the household level and scaled such that the base unit of measurement is in average daily values. For example, a household that typically spends $35 per week on coffee would have an average daily expenditure of $5. If in one week expenditure fell from that level to $20, the normalized series would move from 7 to 4 days of average expenditures. Each plot contains the week-to-week estimates for β_t with dotted lines denoting the 95% confidence interval.

3.2.1 General Effects on Income and Expenditures

The baseline results are a set of β_t estimates from Equation 3.1 for the aggregated and individual categories described in the previous subsection. Figure 3.4 presents the results for four broad measures of income and expenditure. As before, the dashed blue vertical lines mark the beginning and end of the shutdown, while the vertical red line marks the week containing the first completely missed paycheck.

Panel A shows the estimated effects on household recurring income, which during the week of the first missed paycheck during the shutdown (marked by the red line) dropped precipitously for the treatment population. Income from all other sources, presented in Panel C, shows a significant increase in the two weeks following the missed paycheck. Transfers from online savings accounts, directly-deposited cash advances from credit cards, short-term loans or transfers, and any other intermittent source of funds to the household's operating account would all be reflected in this series.

The decline in inflows following the resolution of the shutdown most certainly in part reflects a reduced need of funds transferred from other sources. A second possibility is the impact the shutdown may have had on the ability of households to file their 2018 income tax returns. With a lack of certainty regarding the repayment for lost wages at the end of the 2018 calendar year, affected households may have delayed their preparation of income tax returns until the uncertainty was resolved. As income tax refunds would be a major component of idiosyncratic inflows at this time of year, this decline in non-employment income at the end of February and subsequent spike in March is suggestive of a potential secondary impact of the shutdown.

Total debit and credit card expenditures, shown in Panel B, fell as well during the shutdown period but did not reach the lowest point relative to the control population until the week of the missed paycheck on January 11, 2019. The difference between treatment and control rebounds slightly in the following week as the treatment group decreases expenditures in the same off-week fashion illustrated by both populations in Figure 3.3. Expenditures fall again for the treatment population compared to the control in the following week as a second paycheck is missed the preceding Friday and post-shutdown payments only begin arriving late in the week. The decline in expenditures is likely muted for this reason as well as the increase in non-recurring income that began the week of the first missed paycheck.

Finally, Panel D illustrates the weekly series of all incoming funds. As recurring household pay declines for the treatment group within the shutdown period, the effect of off-week non-recurring inflows can be seen as total household inflows relative to the control is actually at its highest point aside from the week in which back payments to the furloughed and excepted federal employees are made.

3.2.2 Effects on Recurring Obligations

Figure 3.5 contains the estimates of the effect of the shutdown on outgoing payments. Panel A represents all electronic payments from a household account. This contains payments made by the household through an electronic bill payment service as well as externally-initiated debits such as those for an automatic credit card payment. There are no major outliers in the weekly series aside from the week of the first missed paycheck and the week immediately following the repayment. Mortgage and other recurring loan payments follow a similar pattern in Panels B and C, respectively. Mortgages in particular have previously been identified as an underappreciated source of household liquidity for households in this particular situation. Yet for both payment series, the furloughed households actually show a significant increase in the week prior to the first missed paycheck.

One explanation is that debt payments are actually prioritized by a large proportion of the general population and only delayed if there is no alternative. In both this analysis and in Gelman et al. (2018), there is a slight increase in mortgage payments the week prior to the first missed (or significantly diminished in the case of the 2013 shutdown) paycheck. It may be that some households, facing a near-term shortfall in liquidity, ensure their monthly debt obligations are satisfied by actually making payments in advance. This could be an optimal strategy for households that might be otherwise tempted to spend these funds on non-durable consumption goods at a later date at the cost of entering loan delinquency with the associated

Figure 3.5: Differences in Electronic Payments

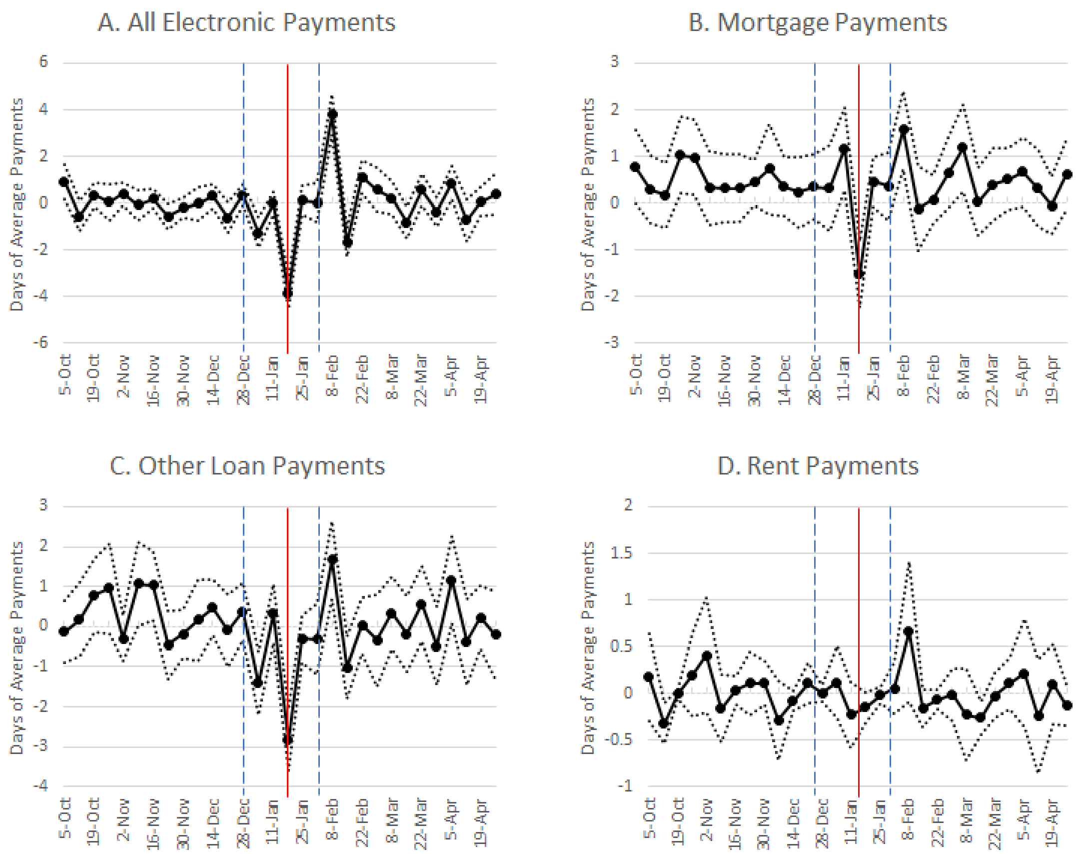

late payment fees and negative entries on credit reports.

3.2.3 Effects on Credit Card Expenditures and Payments

Figure 3.6 presents effects related to the usage of credit cards over the shutdown period. For bank-administered credit cards, expenditures do not appear to respond to the shutdown. While expenditures in Panel A do not decline, there is also little evidence that households that make use of other payment methods such as debit cards substitute into credit to fund consumption. Rather, based on internal credit card data, households appear to be more likely to reduce payments on existing balances and incur more interest charges in the following billing cycle, as shown in Panels C and D, respectively.[8]

3.2.4 Effects on Expenditures by Category

Category-level expenditures are presented in Figure 3.7 for a selected set of six spending categories. Expenditures by furloughed households at grocery and discount stores such as Wal-

[8]Gelman et al. (2018) also find that households delay card payments before increasing card balances via new spending.

Figure 3.6: Differences in Credit Card Behaviour

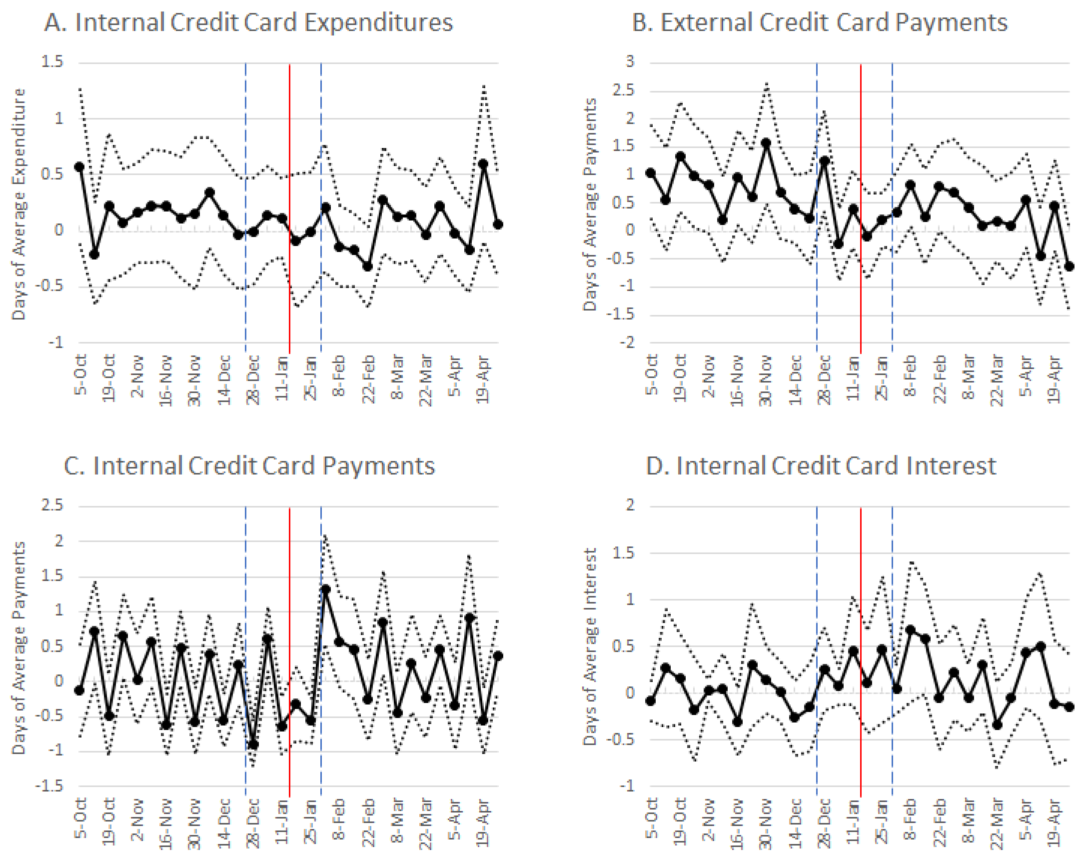

Mart Supercenters, shown in Panel A, experience a marked decline in the week of the missed paycheck relative to the non-furloughed government employees. In terms of consumer behaviour, both grocery and gas expenditures in Panel B appear to occur disproportionately in the days following a paycheck, as evidenced by the bump in estimated difference for the treatment group in the off-pay week of January 25 during the shutdown. Restaurant and Fast Food expenditures in Panels C and D, respectively, show a more uniform decline over the course of the shutdown. It is notable that expenditures in these categories bounce back to above-average levels following the shutdown as back pay is received. For at least some part of the treatment group it appears that dining out was merely delayed rather than foregone completely during the liquidity shock.

Expenditures on clothing (Panel E) and phone, cable, and video services (Panel F) are shown on the bottom row. Both decline upon the missed paycheck, but the series corresponding to phone, cable, and video charges paid via debit and credit cards experiences a much larger post-shutdown recovery. As with the other expenditure trends for recurring expenses, there is a bump the week preceding the first missed paycheck for the treatment group and a sharp decline in the following two weeks. With this pattern occurring yet again in another payment channel (debit/credit cards as opposed to electronic transfers), there is increasing evidence that the utility of delaying scheduled payments diminishes rapidly beyond two weeks. Compared to

Figure 3.7: Differences in Credit/Debit Expenditure Categories

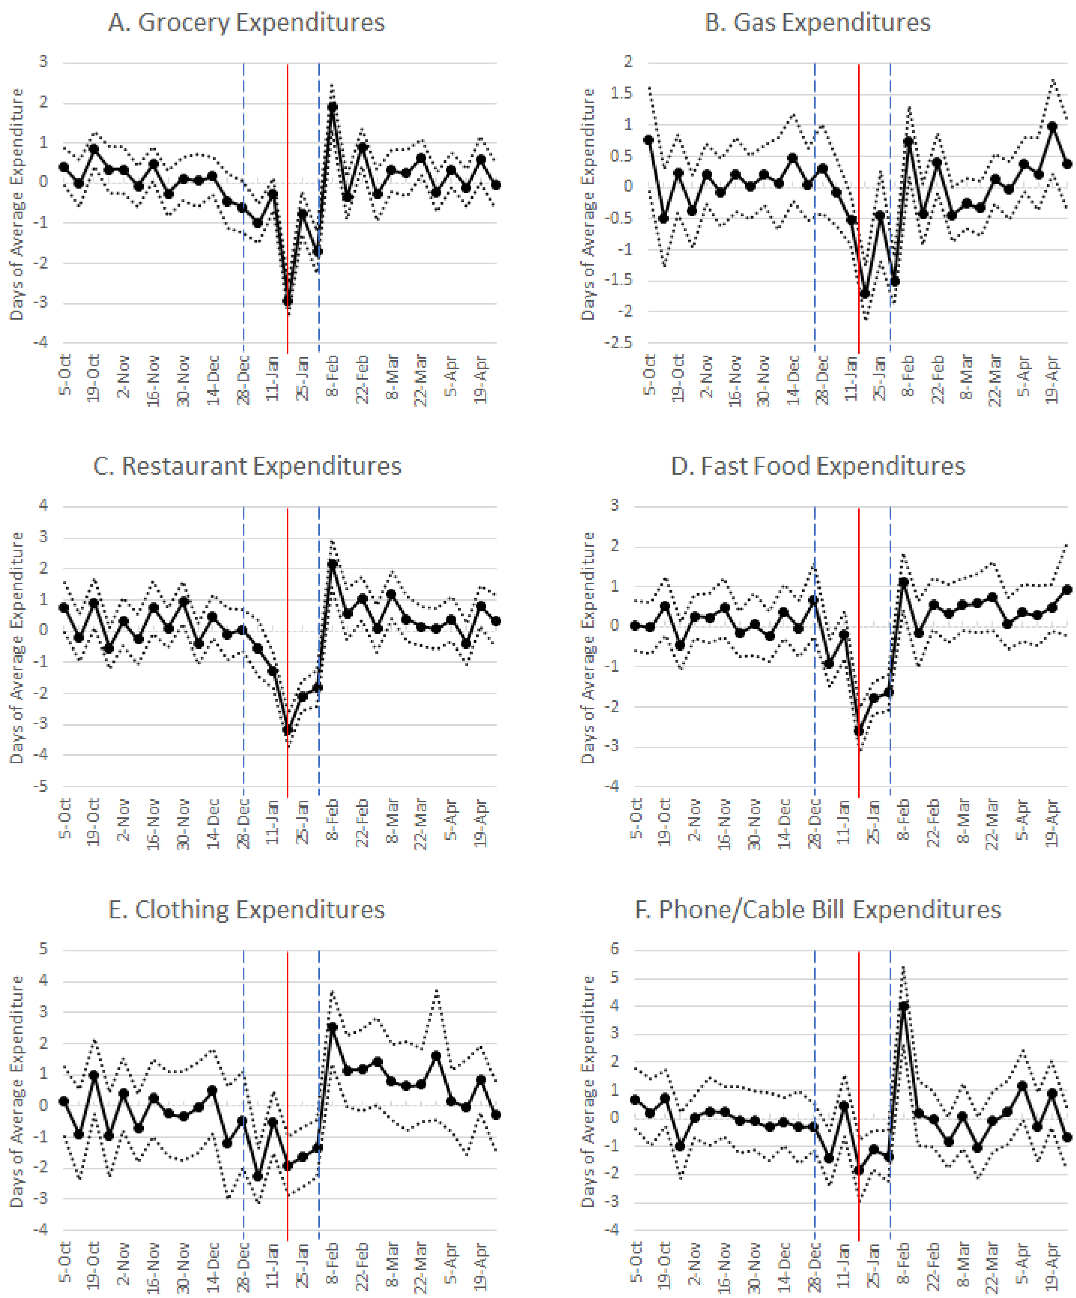

the 2013 event in which the temporary income loss occurred within week two of the shutdown, the observed reduction in payments in 2019 at week three of the shutdown indicates that the realized temporary loss of income may be far more important in determining how expenditures are prioritized than the onset of the government shutdown.

3.3 The Role of Household Liquid Savings

Household savings naturally play a pivotal role in navigating short-term liquidity crises for households. Both Gelman et al. (2018) and Baker and Yannelis (2017) recognize the importance of liquid savings in their studies of the 2013 U.S. government shutdown and find that affected households with less liquid savings make more severe cuts to their expenditures during the shutdown compared to households with a larger cash buffer. In this section, the observed household attributes in the pre-shutdown period of October 2018 and their corresponding activity in the midst of the shutdown in January 2019 are used to partition the sample of federal employees into distinct groups based upon the level of savings in the pre-period. The intent is to find a set of households for whom the government shutdown has minimal impact. Identifying households that are not only "high savings" but are also able to smooth expenditures through the liquidity shock brings additional clarity to the results.

Figure 3.8: Change in Log Expenditures by Savings Ventile

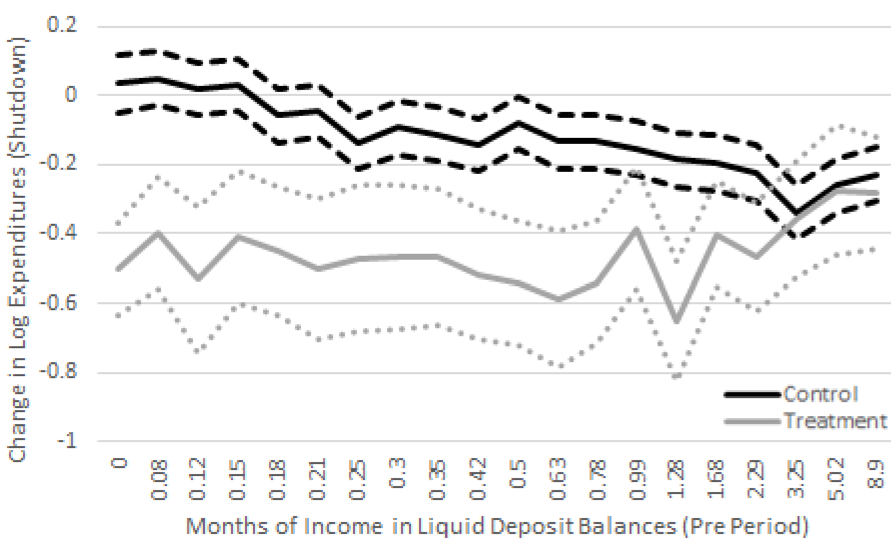

Figure 3.8 shows the relationship between liquid savings and the change in log expenditures for the furloughed population and the other federal workers that received their regular compensation throughout the shutdown. The households are grouped into 20 equally-sized bins, ordered by how many months of income were available in liquid deposit balances (checking and savings accounts administered by the bank) in October 2018. Within each bin, the average change in total log expenditures from October to January is calculated for both the treatment and control groups. The control group is represented by the black line with a dashed 95% confidence interval. The treatment group is shaded in gray with a dotted confidence interval. Generally, expenditures declined between October 2018 and January 2019 as January is typically the low point of consumer spending for the calendar year. Furloughed households experienced a sharp decline in expenditures relative to the unaffected population aside from the top 15% of savers.

Interestingly, the point at which the difference becomes statistically indistinguishable occurs approximately at the three month mark. This is a common lower bound for personal financial advice regarding how many months of income a household should hold in liquid savings. Between 15% and 20% of the government employed households in the sample hold this level of savings. On the lower end, unaffected households holding up to 0.25 months of income in liquid savings do not have a statistically significant change in expenditures between the two months. Hand-to-mouth households that have a consistent income month-to-month would exhibit this pattern, and defining the "low-savings" households based upon this point would be roughly equivalent to the lower tercile used in studies of the 2013 shutdown. Three groups are therefore defined: a low savings group of households holding less than 0.25 months of income in readily available deposits; a mid-level savings group of households with 0.25 to 3 months of income in liquid savings; and a high-savings group with more than 3 months of income available in liquid deposit accounts. These empirically-informed thresholds form the basis of the remainder of the section as the baseline results are recalculated within each savings group. Furloughed households at each level of savings will be compared to non-furloughed control households with corresponding levels of savings prior to the beginning of the shutdown.

3.3.1 General Effects on Income and Expenditures

Figure 3.9: Differences in Income and Total Credit/Debit Expenditures by Liquid Savings Group

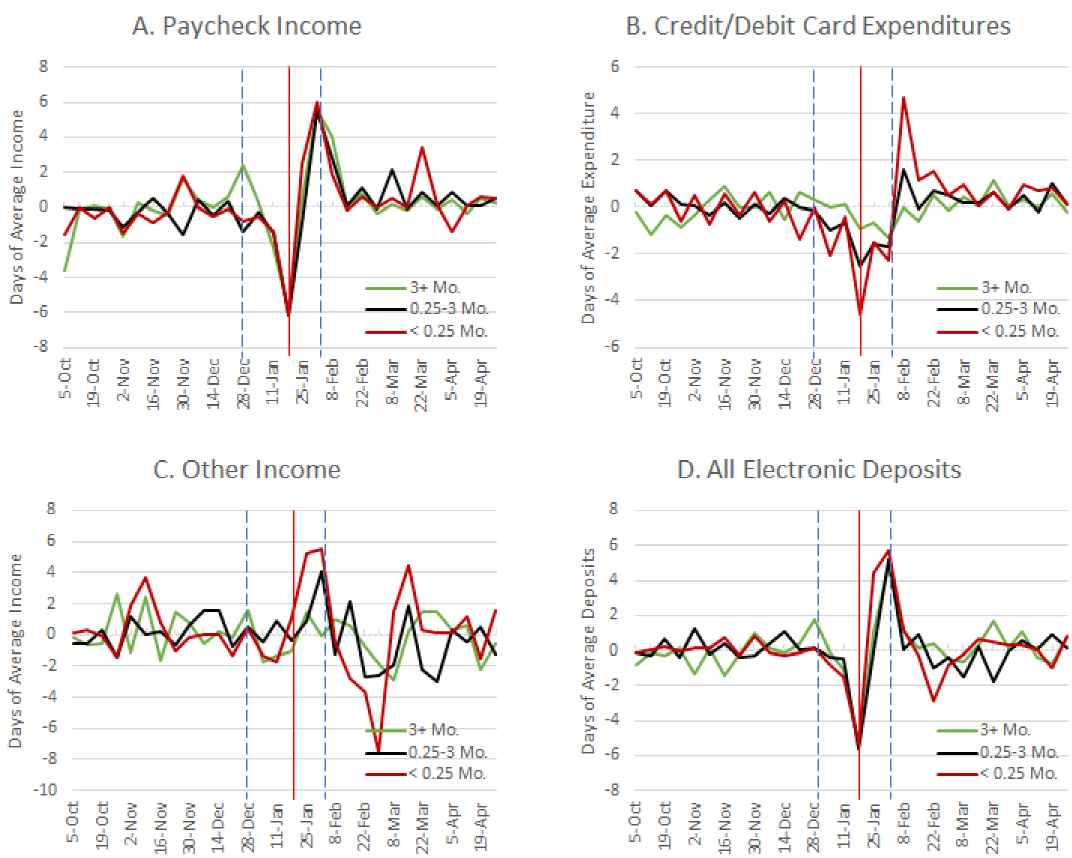

As evidenced by Panel A of Figure 3.9, the treatment group is equally affected across all savings groups in terms of the temporary reduction in income. While there is no evidence of any heterogeneity in terms of lost pay, Panel B shows substantial differences in the consumption response. The low-savings group demonstrates the largest decline in credit and debit card expenditures in the shutdown period. However, it would be a mistake to simply assume that low-savings households have uniformly low assets or lack other sources of liquidity upon which to draw. The series of non-recurring income in Panel C indicates that at least some part this group is bringing in funds from external sources. In addition, programs providing zero- or low-interest unsecured loans specifically to affected government employees were also introduced in the midst of the shutdown. Many government-affiliated financial institutions instituted these types of programs in the midst of the shutdown to provide relief to affected government employees.[9] The increasing availability of short term loans is a better explanation for the large increase in non-recurring income among low-savings households occurring late in the shutdown period. Households holding significant balances in online savings accounts for which transfers are executed without cost beyond the days it takes for funds to clear for example would be expected to move funds much more quickly given that the loss of income was certain to occur by the second week of the shutdown.

3.3.2 Effects on Recurring Obligations

Figure 3.10 presents the results for outgoing payments by savings group. As expected, Panel A shows that low-savings households display the largest amount of volatility in payments during the shutdown relative to low-savings households in the control group. In contrast, the households with at least three months of liquid savings barely appear to be affected at all. For the low-savings group, there is some evidence of heterogeneity with the significant decline in the first week of the shutdown following the first vertical dashed blue line. This was the last week in which furloughed households were paid until the end of the shutdown. For households that utilize external savings and investment accounts, this decline may represent some precautionary savings as regularly scheduled transfers were paused in anticipation of the next missed paycheck. Otherwise, looking across the individual types of recurring payments in the remaining panels, a reduction in payments is not seen until the week of the first missed paycheck. Even the households with the lowest amount of liquid savings appear to keep meeting recurring mortgage and other loan payments in Panels B and C, respectively, up until the first missed check. The increase in payment activity relative to the estimates for groups with more liquid savings in the week prior to the resolution of the shutdown (seen to the left of the second dashed blue line in Panels A, B, and C) may reflect short-term funds received through specialized programs initiated in response to the pandemic. Finally, Panel D indicates that rent payments were also delayed, although the series is highly volatile across all savings groups.

Figure 3.10: Differences in Electronic Payments by Liquid Savings Group

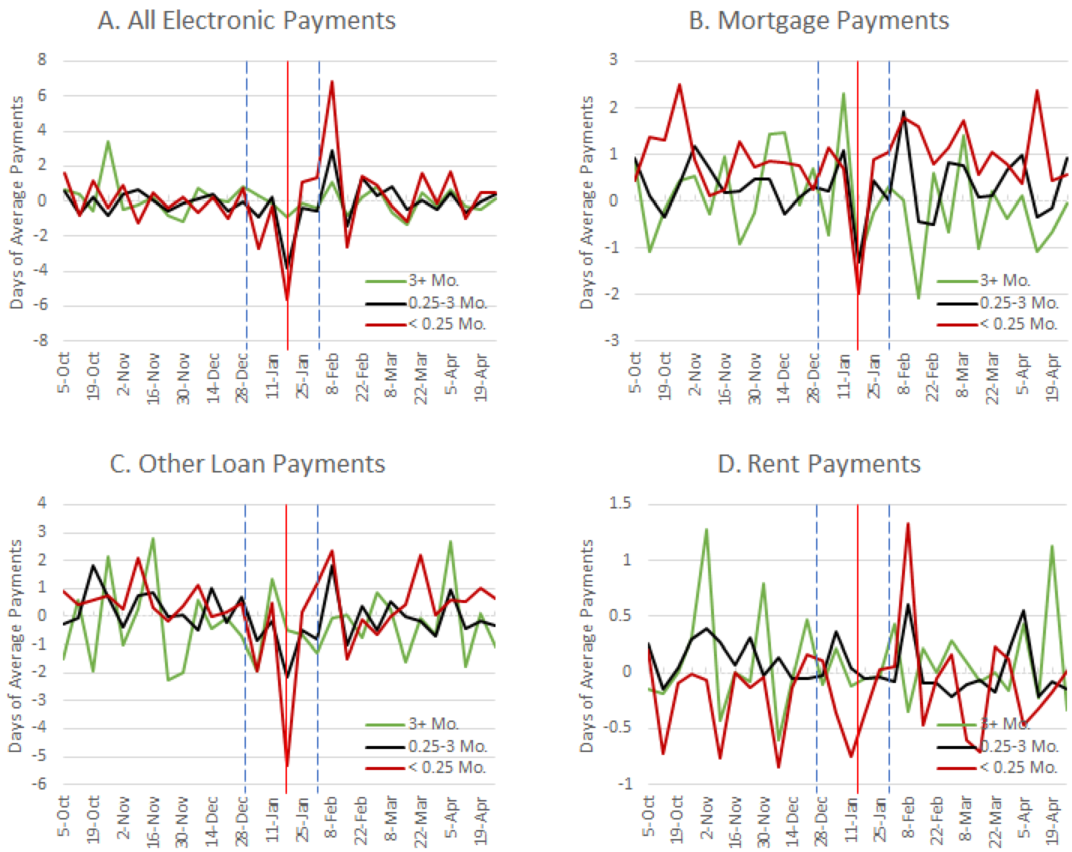

3.3.3 Effects on Expenditures by Category

As with the recurring debt payments, Figure 3.11 shows that retail expenditures by low-savings households exhibit a much greater sensitivity to the interruption of the biweekly pay schedule than households with at least three months of savings. Expenditures on clothing and other soft goods (Panel E) appear to be the first places these households make cuts. Spending at bars and restaurants also begins declining at the start of the shutdown marked by the first dashed blue line, while the remaining categories do not see as large of a response to the shutdown until the week of the first missed paycheck marked by the red vertical line. Each spending category sees a large increase in weekly expenditures by the low-savings households in the week directly following the initiation of back payments for lost wages.

In contrast to the income and debt payments detailed in the previous two subsections, the recovery for these consumption-oriented expenditures does not increase relative to the control group in the week prior to the resolution of the shutdown. This again brings into question the efficacy of postponing recurring debt payments as a means to smooth consumption through liquidity shocks. Among the households with the largest need for liquidity, it appears that debt servicing is prioritized rather than sacrificed.

[9]https://www.cnbc.com/2019/01/22/the-best-and-worst-ways-to-borrow-money-during-the-federal-shutdown.html

Figure 3.11: Differences in Credit/Debit Expenditure Categories by Liquid Savings Group

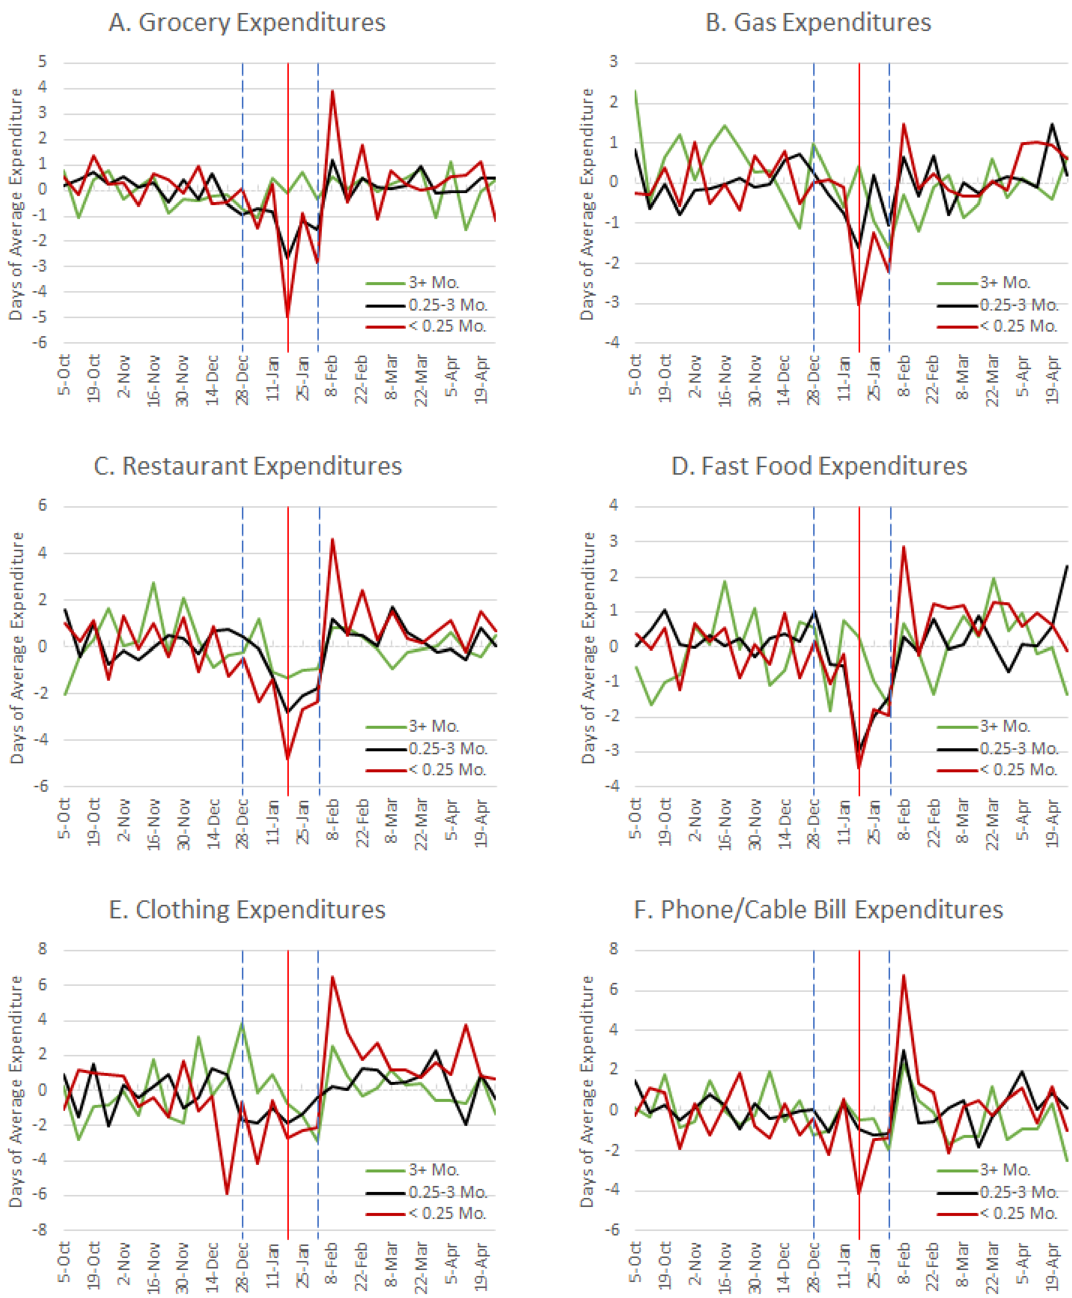

The final section identifies the specific households that defer debt payments and attempts to quantify the impact this behaviour has on smoothing consumer expenditures.

3.4 Delay of Mortgage and Other Recurring Debt Payments

For households lacking sufficient liquidity in terms of available cash balances, the ability to temporarily delay scheduled debt payments is an alternative potential source of funds. Both Gelman et al. (2018) and Baker and Yannelis (2017) document how government workers delayed payments for debts such as mortgages when income was temporarily disrupted during the 2013 shutdown. This section builds upon this prior work by identifying the households with the need and capability to delay debt payments and exploits differences in payment timing to provide a clear view of the role this particular channel plays in providing temporary liquidity. Mortgage payments, other installment loan payments such as auto and student loan payments, and payments on credit cards issued by the home bank are studied in turn. Attention will be paid exclusively to the low-savings households identified in the previous section.

The timing of the first missed paycheck on January 11, 2019 provides a unique opportunity to separate the effect of the shutdown and uncertainty regarding its resolution from the direct liquidity impact. The shutdown began immediately following the close of a complete two-week pay cycle. Checks went out on December 28th to affected and unaffected federal workers alike. Therefore, purely from a cash flow standpoint, affected households experienced almost three full weeks during the shutdown in which there was no disruption to their financial situation. Hand-to-mouth consumers were free to spend as usual right up until January 11. Unlike in 2013 when major progress had been made in the Senate the day before the reduced paycheck arrived in affected workers' accounts, there was no near-term hope for a resolution on the horizon in 2019.[10]

Households on a biweekly pay schedule commonly budget on a monthly basis. On the two occasions each year in which a third paycheck is received in one month, the income is treated more like a financial windfall compared to other paychecks (Zhang 2017). This approach to budgeting is commonly found in online content intended for general audiences.[11] Strictly following a monthly budgeting routine in which recurring bills are regularly paid from the most recent paychecks would lead households affected by the government shutdown to make their typical payments as usual for the first 10 days of January. If, however, households wished to utilize the hypothesized insurance attributes of these regular payments then payments would decline in the first part of the month. After the liquidity shock arrives with the missed paycheck on January 11th the affected hand-to-mouth households with no liquid savings no longer have the option and begin to miss payments for lack of funds. Therefore the relevant comparison to make for testing the consumption smoothing value of payment delay is between affected households during the shutdown that could have made a payment but did not, and affected households that were unable to make a payment due to the missed paycheck.

For each of the three types of payments, household payment behaviour over the entirety of the sample period is studied to infer which part of the month payments are regularly made. Of the total amount paid, households making a majority of electronic payments (payees are unable to be inferred for paper checks or cash payments) in the first 10 days of the months are classified as an "early-month" group. The remaining households with expected due dates on or

[10] A famously short meeting between President Trump and House Speaker Nancy Pelosi had occurred just two days prior on January 9.

[11] One example can be found here: https://www.discover.com/online-banking/banking-topics/5-budgeting-hacks-if-youre-paid-biweekly/

after the 11th are classified as "late-month." If households typically pay bills on or around the relevant due date, any differences in behaviour between these "early-month" and "late-month" households should be exclusively due to differences in the sequencing of repeated account transaction activity from month to month.

Table 3.2: Household Distribution of Payments - Early and Late Month

	Early	Late	None
Mortgage	1,564	2,042	4,407
Credit Card	533	1,265	6,215
Other Payments	1,562	4,412	2,039
Low Savings Households Only			
Mortgage	247	545	1,622
Credit Card	93	225	2,096
Other Payments	395	1,374	645

Notes: Summary of observed majority of household payment activity by lending product over the sample period. Early month payments are those occurring on or before the 10th of the month. Credit card payment activity is restricted to accounts administered by the source bank. Low savings households hold less than 25% of their monthly income in deposit accounts at their primary bank.

Table 3.2 contains the counts for all households observed making the majority of payments early in the month, late in the month, or not observed making a payment at all along with the previously-defined low savings presented on its own. A little under half of the households in the sample made electronic mortgage payments over the course of the sample period. Most households make at least one type of recurring payment. These could be auto loans, student loans, or possibly automated transfers into external savings or investment accounts.[12] Households with low savings have a lower rate of mortgage activity and credit card ownership but are similar to the other households for the other recurring payments category.

3.4.1 Mortgage Payments

Mortgage payments are the most significant scheduled debt payment for most households and are expected to be the first option explored by households searching for temporary payment-deferral-related sources of liquidity. Most U.S. mortgages allow for a 15 day grace period on repayment, meaning that a mortgage due on the first of the month could actually be paid as late as the 16th without incurring any additional penalties or fees. For households that normally

[12]The criteria for a payment to be included in this category are based upon size and regularity (at least five payments averaging at least $200 with coefficient of variation less than 0.25).

pay on or before the due date, delaying payment until the end of the grace period is a costless source of short-term liquidity.

Households affected by the shutdown with a mortgage coming due in early January could presumably delay mortgage payments while pursuing other less-liquid sources of cash such as transfers from savings and investment accounts held at other institutions or some of the specialized low- or zero-interest unsecured loan programs aimed at affected government workers during the shutdown. In addition, there is an option value in delaying payments if new developments such as government mandates regarding forbearance for affected workers are a possibility. Figures 3.12 and 3.13 present side-by-side estimated effects of the furlough for the early- and late-mortgage payment segments in a similar fashion to the preceding section on mortgage payments and total debit and credit card expenditures, respectively.

Figure 3.12: Variation in Mortgage Payments by Mortgage Due Date

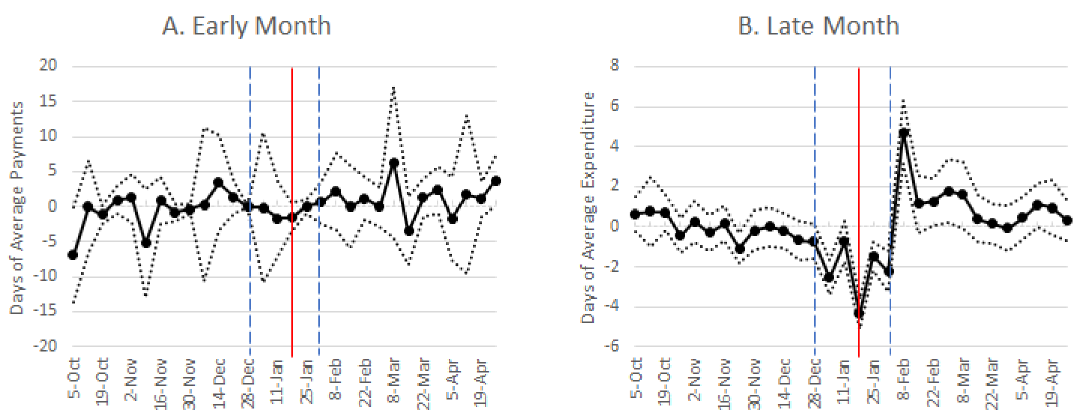

Figure 3.13: Variation in Credit/Debit Expenditure by Mortgage Due Date

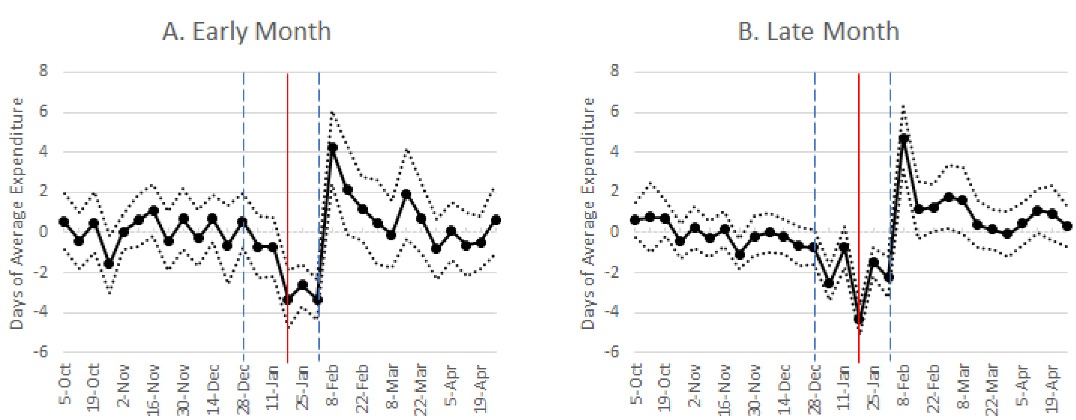

Overall, there are meaningful differences between the early- and late-month payment segments relative to their respective control groups within the shutdown period. In Panel B of Figure 3.12, affected households that typically pay mortgages later in the month show a much sharper decline in mortgage payments during the shutdown period compared to the households in Panel A that typically pay early in the month. Figure 3.13 shows that for these low-savings households (those with the highest likelihood to be hand-to-mouth consumers), the early-month group in Panel A appears to experience less of a decline in expenditures compared to the late-month group in Panel B in the first part of the month. After the missed paycheck, the part of the month the late-month group would usually be making mortgage payments, the early-month group appears to reduce spending much more than the late-month payment group. This is likely due to the mortgage payment crowding out expenditures within that part of the month, particularly for households that behave in a more hand-to-mouth fashion. For mortgage-paying households on a roughly twice-monthly pay schedule but without any liquid savings, the majority of each paycheck immediately prior to the mortgage due date is most likely earmarked for making that single payment. These households are far more likely to be able to make meaningful discretionary cuts in retail purchases following the other paycheck that does not precede the mortgage due date. Regardless of the composition of week-to-week expenditures, households typically making early mortgage payments still make mortgage payments during the shutdown and make sizeable reductions to credit and debit card expenditures upon missing the first paycheck. This is counter to what would be expected if mortgage payment deferrals were a significant avenue for consumption smoothing.

Figure 3.14: All Electronic Payments by Mortgage Due Date

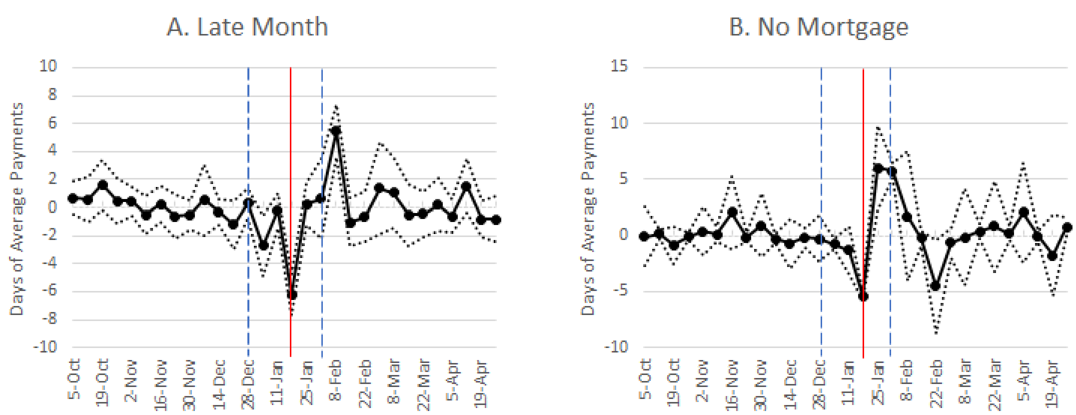

As one final mortgage-related comparison, Figure 3.14 compares the total estimated impacts on payments made via ACH and all debit and credit card expenditures for the same low-savings households with late-month mortgage payments now in Panel A to households without any observed mortgage payments at all in Panel B. In this case, there does not appear to be any difference at all between the households with a late-month mortgage payment and those without any observed mortgage payment at all in terms of total payment volume. In addition, Figure 3.15 shows no meaningful difference in purchases made via debit and credit card either.

In general, there is little in the way of evidence to suggest that the delay of mortgage payments is a meaningful channel for supporting consumption smoothing. Among the households with the greatest liquidity need, those that make mortgage payments later in the month do not appear materially different from households without a mortgage at all. Further, recall from Figure 3.13 that those that pay mortgages early in the month had more volatile expenditures than those paying late in them month, making them the least successful at smoothing debit and

Figure 3.15: Total Credit/Debit Expenditure by Mortgage Due Date

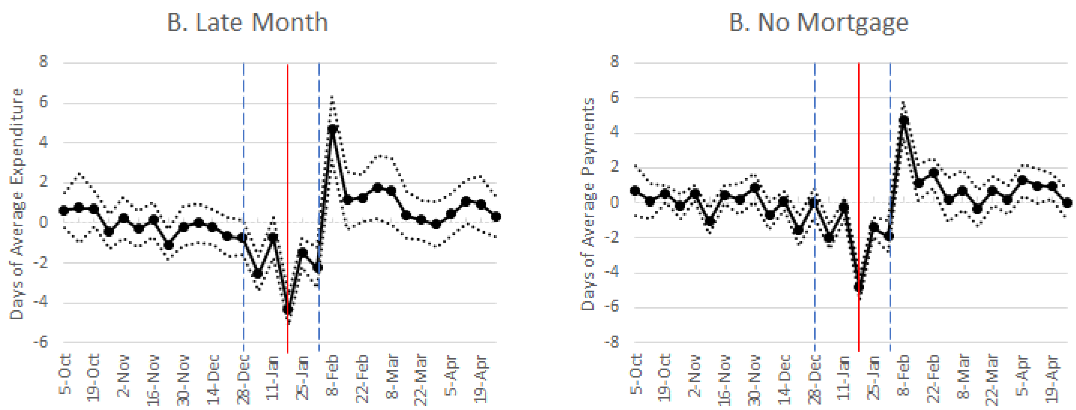

credit card expenditures out of all households.

3.4.2 Credit Card Payments

Attention now turns to the remaining types of recurring payments. While the available credit on credit card accounts is for many considered to be a primary source of emergency household liquidity, there are a few reasons why it makes sense to also look at payments as the most active margin for managing liquidity in this situation. As is generally the case with open lines of credit, there is a strong inverse relationship between credit need and credit availability. On the one hand, households with sterling credit scores and high accompanying credit limits have higher incomes, lower month-to-month credit utilization, and a demonstrated history of making on-time payments. This is precisely the profile of a household that would be least expected to have difficulty weathering a delay in expected income. On the other hand, households with less reliable income or that have previously demonstrated a need to pull consumption one or more months forward are extended less credit, if any at all.

For creditworthy households, credit cards offer a large degree of payment flexibility. A household may not be able to charge a mortgage payment to a credit card but could switch to minimum payments for a month and put funds originally earmarked to pay off a card balance toward the mortgage instead. Many other households utilize balance transfers or put major purchases on a credit card and pay it off over time. There may be little in terms of month-to-month card purchasing activity for these households but any regular payment amounts exceeding the minimum payment should likewise be available as an easily accessible, if not low cost, source of liquidity.

Similar to the mortgage payment analysis, activity over the course of the sample period is examined to infer the typical payment date for those households also using credit cards administered by their primary bank. Restricting attention to these households only, Figures 3.16 and 3.17 present the estimated effect of the furlough for the early and late credit card payment groups on credit card payments and total debit and credit card expenditures, respectively. As before, the dotted vertical lines mark the start and end of shutdown while the solid red line marks the week that began with the first missed paycheck. Of the recurring payment types, the

Figure 3.16: Variation in Card Payments by Card Payment Due Date

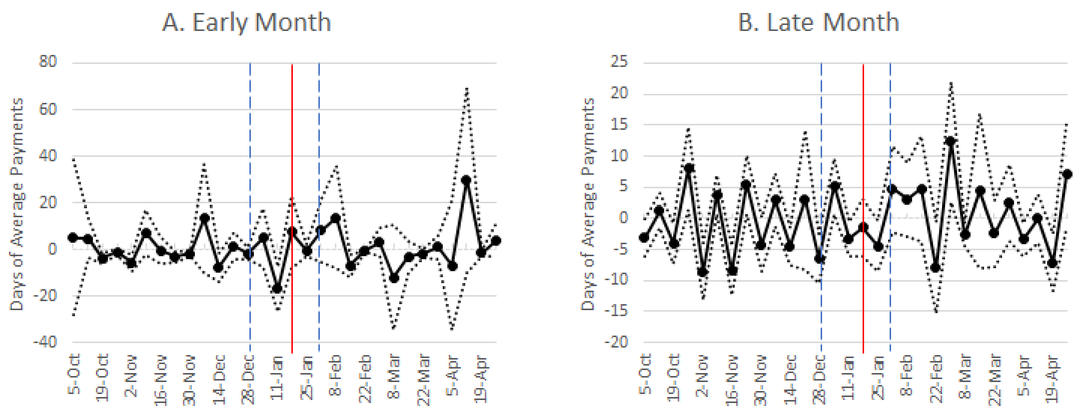

credit card series is the noisiest. By nature, credit card payments have the largest variance compared to other recurring payments. In addition, the proportion of the sample observed with a credit card is relatively small, particularly for the low-savings group. Despite these limitations, this view provides some useful insights.

Figure 3.17: Variation in Credit/Debit Expenditure by Card Payment Due Date

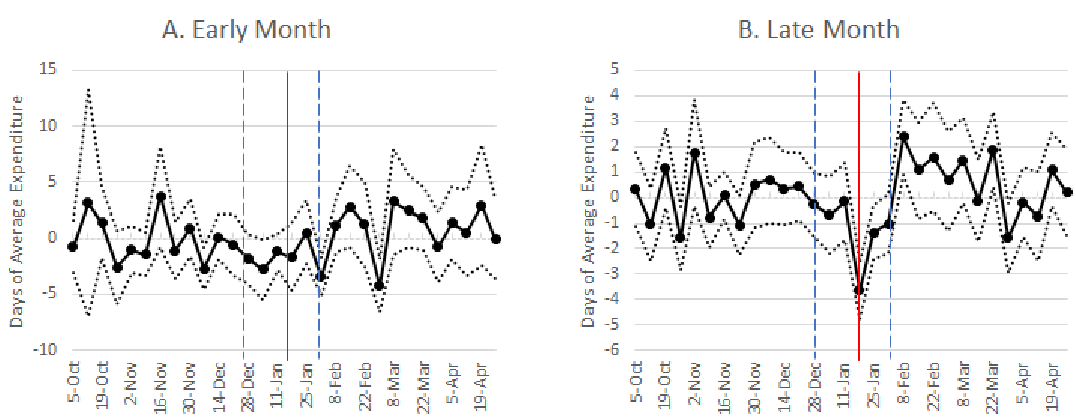

Overall, households with credit cards unsurprisingly appear to more effectively smooth consumption over the course of the shutdown. Figure 3.16 indicates that both groups of card holders reduce payments during the shutdown, the early month group in Panel A significantly so. Following the resolution of the shutdown, payment amounts do appear to directionally increase in general for credit card holders as they receive back pay. While noisy, an interesting difference occurs among these low-savings households. Affected households that typically make payments early in the month registered a sharp decline in payments the week ahead of the first missed check with a corresponding increase after the end of the shutdown. It could be the case that households weighing their option to reduce or even forego payments altogether (and incur the relevant late fees) would be more likely to do so early in the shutdown period rather than later as late fees are not prorated and carrying a revolving balance eliminates the grace period for future purchases until the balance is paid in full. Facing a material fixed cost, households with payments coming due toward the end of the shutdown may have been more inclined to temporarily reduce expenditures in other areas if expectations of a resolution were on the rise. However, Figure 3.17 is noisily estimated and does not provide much in the way of additional support for this hypothesis given the noisy estimates for the households paying early in the month.

3.4.3 Other Recurring Payments

My recurring payment analysis concludes with a look at all remaining recurring payments. The definition of these payments is the least precise of the three studied in this section; auto and student loans, insurance payments, and even regular investment purchases could meet the regularity and amount criteria. In this case, households are classified as early- or late-month payment households based up where the majority of payment volumes occur. For example, if a household typically makes an auto payment for \$500 early in the month and a student loan payment for \$350 later in the month, it would be classified as an early-month household in this section.

A cursory examination of specific payees ordered by the largest deviations of the furloughed households' share relative to the control group between the months in which the shutdown occurred and the remainder of the sample provides some insight into where most of the action was occurring. The destination with the largest drop was investment-related, followed by external transfers to other deposit accounts, managed credit payments (such as payments to credit counseling organizations), student loans, and several captive auto finance companies. This ordering makes intuitive sense considering that opportunity costs are relatively low for deferring transfers to external savings and investments over a short time frame. Funds typically earmarked for external savings and investment products are, if not easily diverted, at least done without additional monetary cost. In contrast to mortgages, many student and auto loan originators and servicers were quick to issue guidance making it easy for affected workers to receive deferrals and/or fee waivers on their loans, which many took advantage of. Finally, the prominence of payments to third-party creditors in the list is reasonable. One possibility is that credit counseling agencies are used to working with households with erratic or irregular income and were quick to accommodate the needs of furloughed government workers. The other possibility is that the credit scores for households making these types of payments is already so low that there is little incremental damage done by reducing or delaying a negotiated credit payment.

Of the three types of recurring payments studied, deferring these non-mortgage, non-credit card payments appears to have the most significant impact. Figures 3.18 and 3.19 present the estimated effect of the furlough for the early and late payment groups on the sum of these miscellaneous payments and total debit and credit card expenditures, respectively. In Figure 3.18 both groups of households affected by the shutdown appeared to take advantage of the payment flexibility offered by this channel regardless of whether payments were typically made early or late in the month, but partly by construction. Households that prefer to utilize external savings and investments naturally cease outbound transfers when income is disrupted and even look to reverse the direction if necessary. These households may be classified as low-savings, but only because their external savings are not observed. Although these accounts are less

Figure 3.18: Variation in Other Payments by Payment Due Date

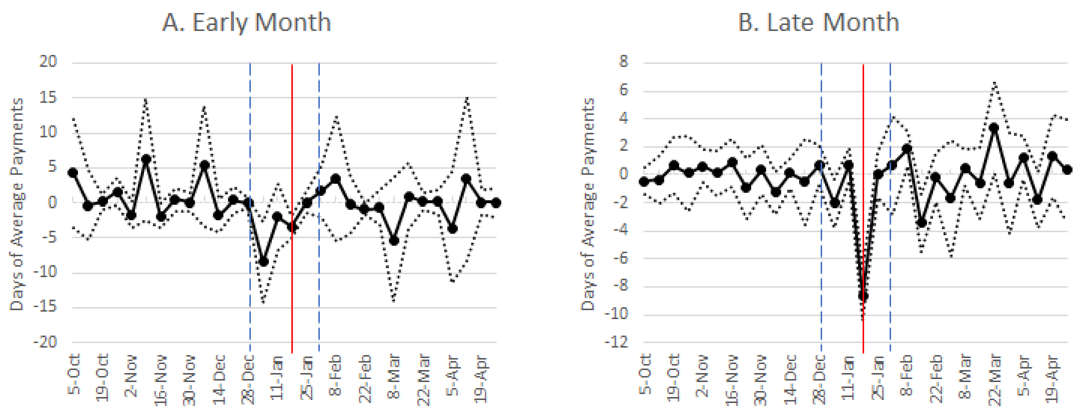

liquid in that there may be a delay in moving funds back to their operating account, pausing recurring transfers to these accounts is an immediate and costless source of liquidity.

Figure 3.19: Variation in Credit/Debit Expenditure by Payment Due Date

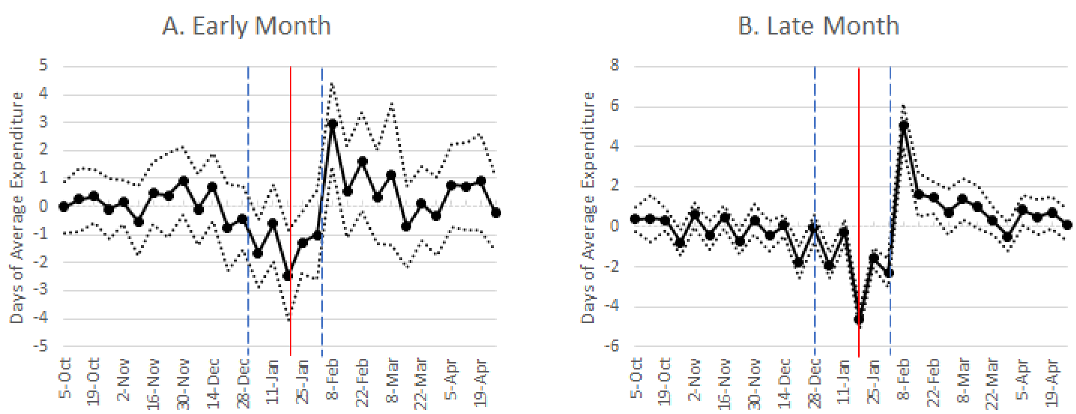

It is important, however, to recall the allowances made by auto and student lenders. These loans are a sizeable component of many low-savings households' monthly obligations and the announced deferrals in conjunction with the speed in which they were announced appear to have made a significant difference particularly for less liquid households facing payments early in the month. Comparing early-month households to late-month households in Panels A and B of Figure 3.19, respectively, the early-month households experienced a much less severe decline in debit and credit card expenditure following the interruption of pay. Among these low-savings households, having a payment deferred early halved the effect of the lost paycheck compared to the households with more payments made later in the month.

The timing comparison illustrates the importance of quickly making cash available under clear terms ahead of a liquidity crunch. Conditioning the duration of the deferral on the resolution of the shutdown likely made all the difference for liquidity-constrained households looking to temporarily cut recurring expenses to support other consumption. In contrast, mortgages failed to deliver on all accounts. Rather than deal with the 15 day mortgage grace period and uncertainty regarding continuation options, households appeared to overwhelmingly prefer auto and student loan deferment.

3.5 Discussion

This chapter studies the households of employees affected by the 2018-2019 U.S. federal government shutdown and found that most are ill-equipped to deal with temporary liquidity shocks. This shutdown was similar in many respects to the 2013 shutdown previously studied by Gelman et al. (2018) and Baker and Yannelis (2017). Household expenditures fell precipitously; the average effect on credit and debit card expenditures amounted to a decline of about three full days of average spending relative to the control group in the week following the first missed paycheck. At the onset of the shutdown more discretionary purchases such as those made at restaurants and clothing stores appeared to be the first to be cut as a precautionary savings measure; particularly for households with less liquid savings. Grocery and gas purchases then declined following the first missed paycheck on January 11.

The earlier literature studying the 2013 shutdown struck a cautiously hopeful tone regarding the role recurring debt payments may play in helping households without liquid savings to smooth consumption. Specifically, the decline in mortgage and credit card payments observed among these households the week their pay fell 40% during the 2013 shutdown was hypothesized to temporarily support consumption as households took advantage of costless buffers such as grace periods. My study of the 2018-2019 shutdown takes advantage of a more severe disruption to regular income. I confirm the broad patterns seen in 2013. Loan payments appear to fall with the missed paycheck; most significantly for households with less liquid savings. However, this paper deliberately seeks and fails to find more direct evidence of substitution from debt servicing to credit and debit card expenditures. In fact, there is a more convincing case to be made that households actually prioritize mortgage payments at the expense of new purchases in the short-run.

My key contribution is made by taking advantage of the full-calendar-month duration of the 2018-2019 shutdown to observe how affected households with mortgages, all of which would have had a payment due at some point during the shutdown, behave differently depending upon whether their monthly mortgage payment is expected to be made before or after the first missed paycheck on January 11. The fortunate alignment of government pay periods to the shutdown and a full paycheck received one week into the shutdown meant that households had about three full weeks under the shutdown before the direct liquidity impact was felt. This affords an unprecedented opportunity to separately identify delays in mortgage payments due solely to precautionary savings from those in which the loss of income is a confounding factor.

The observed decline in mortgage payments made by households affected by government shutdowns is overwhelmingly driven by the realized disruption of income as opposed to expectations. Faced with a temporary but certain-to-occur liquidity shock, households with the option to make or delay mortgage payments generally choose to pay. Relative to their respective control groups on similar mortgage payment schedules, affected households that typically make mortgage payments early in the month had the most volatile debit and credit card expenditures. The low-savings households confronting mortgage due dates after the missed paycheck exhibit the same level of consumption smoothing as those households without any mortgage at all.

Taken together, the evidence presented in this paper suggests that household confidence in meeting near-term obligations is a key determinant of how households navigate liquidity shocks. Many households held an amount in liquid savings that would have sufficiently covered typical expenditures during the liquidity shock, but only those households that followed the common heuristic of at least three months of liquid savings appeared to uniformly navigate the shutdown without a significant disruption to their usual spending patterns. When households reduce expenditures, the largest debt obligations such as mortgages appear to be prioritized. Only when households are given specific relief regarding near-term debt obligations do we see a meaningful impact on consumption smoothing, as seen with auto and student loan deferrals. These findings should be of interest to the growing literature on mental accounting.[13] It is plausible that the deliberate act of creating and funding an emergency savings account with three months of income functions as a sort of psychological primer that makes it easier to deploy savings when the need arises.

This paper makes several observations regarding households savings behaviour that may be cause for further study. The first opportunity is a more comprehensive analysis of the "three months of income" savings threshold to more precisely infer the psychological effect of satisfying a popular savings heuristic. The second relates to the study of earmarked funds. Meeting the three months heuristic appears to confer a psychological benefit; there may be similar benefits in terms of realized outcomes in situations such as these simply from having labeled partitions of funds (a feature being rapidly adopted in the banking industry) as opposed to commingled funds in a single deposit account.

Finally, regarding policy, the lack of evidence that households incur additional costs to smooth consumption related expenditures (as captured by debit and credit card transactions) raises questions about how much effort should be made to specifically support consumption through liquidity shocks such as the shutdown. Rather, resources may be better spent on enabling households to keep existing debt obligations current. Ideally this would be a comprehensive program that quickly works directly with loan servicers to implement payment deferrals that remain in place for the duration of the liquidity shock. Execution of such a program would be greatly simplified compared to one that instead sought to provide liquidity directly to affected households (e.g., through new unsecured loan programs). At the same time, policymakers must be careful not to create a windfall event that leads to a burst of new consumption.

[13] Zhang and Sussman (2018) state that "establishing a direct link between mental accounting and economic outcomes, particularly in the long-term, remains an ongoing challenge."

Bibliography

Scott R. Baker and Constantine Yannelis. Income changes and consumption: Evidence from the 2013 federal government shutdown. *Review of Economic Dynamics*, 23:99–124, 1 2017.

Richard Blundell, Luigi Pistaferri, and Ian Preston. Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921, December 2008.

James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick. Optimal defaults. *American Economic Review*, 93(2):180–185, May 2003.

Jeanne Commault. Does consumption respond to transitory shocks? *American Economic Journal: Macroeconomics*, forthcoming.

Pascaline Dupas and Jonathan Robinson. Why don't the poor save more? evidence from health savings experiments. *American Economic Review*, 103(4):1138–71, June 2013.

Peter Ganong and Pascal Noel. Consumer spending during unemployment: Positive and normative implications. *The Quarterly Journal of Economics*, 128:1123–1167, 2013.

Michael Gelman, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown. *Journal of Public Economics*, 189:103917, 2020.

David B. Gross and Nicholas S. Souleles. Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data*. *The Quarterly Journal of Economics*, 117(1):149–185, 02 2002.

Greg Kaplan, Kurt Mitman, and Giovanni L. Violante. Non-durable consumption and housing net worth in the great recession: Evidence from easily accessible data. *Journal of Public Economics*, 189:104176, 2020.

Annamaria Lusardi, Daniel J. Schneider, and Peter Tufano. Financially fragile households: Evidence and implications. NBER Working Papers 17072, National Bureau of Economic Research, Inc, 2011.

Irina A. Telyukova and Randall Wright. A Model of Money and Credit, with Application to the Credit Card Debt Puzzle. *The Review of Economic Studies*, 75(2):629–647, 04 2008.

C. Yiwei Zhang and Abigail B. Sussman. Perspectives on mental accounting: An exploration of budgeting and investing. *Financial Planning Review*, 1(1-2):e1011, 2018.

Chapter 4

E-Commerce Taxation: Who Pays?

The collection of state sales taxes on online purchases has been a contentious issue from the beginning of the e-commerce era. The U.S. Supreme Court ruling on Quill Corp. v. North Dakota (1992) set the precedent for the original piecework system where tax collection requirements were largely based upon where retailers maintained a physical presence. For the next two decades, aggrieved U.S. states losing millions of dollars annually in tax revenue and physical retailers put at a sizable pricing disadvantage lobbied federal policy makers in search of relief. The highly anticipated June 2018 Supreme Court ruling on South Dakota v. Wayfair, Inc. finally overturned these contentious limitations and opened the door for states to completely reshape their tax policy mix. Given states' rush to action following the decision, a careful examination of the implications of various tax policy options in varying stages of implementation is of great interest.

This chapter studies consumer purchasing behaviour before and after the 2015 implementation of state sales taxes in Ohio and Michigan on purchases at Amazon.com, with a specific focus on identifying differential effects along major demographic dimensions. I document a repeated theme of the rural poor absorbing the entirety of the sales tax applied to Amazon.com purchases in 2015 with little evidence of substitution to competing retailers. In contrast, the suburban households shifted a meaningful quantity of spending away from the entire Amazon platform, significantly increasing expenditures at local electronics, warehouse, and book stores. Higher incomes, more mobility, and close geographic proximity to dense brick and mortar retail give suburban consumers the widest possible choice set when it comes to retail options. Urban households also appear to experience effects similar to rural households to a lesser extent. These households are most likely to rely on public transportation, are most affected by congestion whether or not they have personal transportation, and are generally are less able to make large local shopping trips for these reasons in addition to smaller living spaces and storage.

Unsurprisingly, many states have already taken advantage of the Wayfair ruling and moved to apply state sales taxes to online sales by entities previously out of their jurisdiction. Table 4.1 lists 22 states that have enacted sales tax policy changes since the Supreme Court's decision. There are two primary ways that states have taken advantage of the new ruling. The first is initiating the collection of sales taxes from "out-of-state businesses" or "remote sellers" with no physical presence in the state that directly sold and shipped goods to the state's residents. The second relates to so-called "marketplace facilitators" such as eBay. These online marketplaces traditionally carry no inventory and merely act as a coordinator between individual buyers and sellers. The new provisions put the onus on the platform operator rather than the individual seller to collect and remit the appropriate sales tax. For simplicity, a state taking either or both of these actions is considered to have performed a "sales tax expansion."

As states begin collecting sales taxes on all online purchases, to what extent will physical retailers benefit? There will likely be some immediate gains as consumers substitute away from marginal online purchases, but retailers may benefit further from the relaxing of what may have been a binding constraint on the overall tax mix. Will states be able to favor state-domiciled retailers through a combination of higher sales taxes and lower property and/or income taxes? The answer likely depends in part upon the degree to which the online shopping experience is differentiated from brick and mortar retailers. States may have a more difficult time pushing through more ambitious changes if voters highly value a greater variety of goods or the convenience of shopping online from their connected devices.[1]

In addition to the expansion of sales tax collection, many states made changes to income and/or property tax collection that appear to at least partially offset the increase in sales tax revenues. Wisconsin, in particular, went so far as to explicitly tie the amount of income tax relief for those in its bottom two tax brackets to the additional sales tax revenue received from out-of-state retailers and marketplace providers. How should one think of this approach compared to Arkansas cutting income tax rates across the board and increasing the homestead exemption (a policy favoring homeowners and the elderly)? I seek to answer this question by studying how different types of households have previously responded to exogenous increases in the cost of online goods.

In 2015, both Ohio and Michigan began collecting state sales tax on purchases from Amazon.com, the largest online retailer. This chapter studies consumer behaviour before and after these policy changes to better understand the implications of various tax policy options available in the wake of South Dakota v. Wayfair, Inc. Detailed high-frequency consumer banking data allow for the identification of how Amazon customers in these states responded to a de facto increase in prices relative to consumers in neighboring states that experienced no change in sales tax collection practices during the sample period. The particular focus on differential effects based upon consumer wealth and breadth of local retail options is a new contribution to our current understanding of sensitivity to online pricing. Specifically, the ability to pinpoint individual consumer household locations and characterize the local retail landscape is an uncommon benefit of these unique data.

Overall, consumers do appear to behave in a way that suggests avoidance of state sales taxes is a consideration. Goolsbee (2000a, 2000b), Alm and Melnik (2005), and Scanlan (2007) all attempt to quantify the impact of state sales taxes on likelihoods to purchase based on survey data and find significant negative effects.[2] Beyond survey data, Einav et al. (2014) examines data on item views and eventual purchases by eBay users. Interested consumers clicking through to an individual item description page can be "surprised" to find an item is located in their home state and therefore subject to state sales taxes. Consumers show a

[1]In an influential early study of online purchasing behaviour, Brynjolfsson et al. (2003) examine the market for books after Amazon's entry and concludes that the consumer surplus associated with the variety of goods available for online purchasing exceeds any benefit from more competitive pricing by a factor of 7 to 10.

[2]Scanlan (2007), in particular, finds that the tax rate is itself an important consideration beyond merely whether a tax is levied or not.

Table 4.1: Post-Wayfair Tax Policy Changes

State	Sales Tax Expanded	Impact ($M, est.)	Offset (If Any)	Impact ($M, est.)
Arizona	10/1/2019	$57.0	Income	-109.0
Arkansas	7/1/2019	$32.0	Inc., Prop.	-$38.5
Colorado	6/1/2019	$43.9		
Illinois	1/1/2020	$80.0		
Indiana	7/1/2019	$63.7		
Iowa	1/1/2019	$55.8	Income	-$186.0
Kentucky	7/1/2019	$17.0	Income	-$13.7
Maine	10/1/2019	$13.1	Income	-$13.4
Maryland	10/1/2019	Unknown	Income	Unknown
Massachusetts	7/1/2019	$41.7		
Nevada	7/1/2019	$16.5		
New Mexico	7/1/2019	$43.0	Income	-$63.0
New York	6/1/2019	$170.0		
North Carolina	2/1/2020	$65.0	Income	-$79.3
North Dakota	6/21/2018	$20.0		
Rhode Island	7/1/2019	$11.5		
Texas	10/1/2019	$242.5	Income	Unknown
Vermont	6/1/2019	$13.4	Income	Unknown
Virginia	7/1/2019	$82.5	Income	Unknown
Washington	3/14/2019	$67.0	Poperty	-$6.9
West Virginia	7/1/2019	$6.3	Income	Unknown
Wisconsin	10/1/2018	$93.5	Income	-$93.5

Source: National Conference of State Legislatures

clear preference for items located out-of-state. In addition, consumer purchase probabilities on eBay increase by 2% for every 1pp increase in state sales taxes observed within the 2008-2010 sample. At the same time consumers substitute away from home-state sellers (for whom sales tax is collected) with the same 1pp tax increase lowering their purchase probability by 3-4%.

This chapter is closely related to Baugh et al. (2018), which also documents an average decline in the tax-adjusted expenditures of Amazon consumers following the implementation of sales tax collection practices. In comparison, I compile a relatively larger set of evidence of substitution to a wide range of competing retailers. This is likely due to my focus on identifying substitution patterns within specific development level segments (urban, suburban, and rural) as opposed to the population as a whole. Isolating the effects of sales tax increases at this level of granularity is a useful contribution to the broader tax policy discussion underway in many states following the Wayfair ruling.

The study proceeds as follows. Section 4.1 describes the sales tax policy environment and consumer data. Section 4.2 presents the empirical strategy. Section 4.3 contains the primary results. Section 4.4 presents several robustness checks. Section 4.5 summarizes the chapter and concludes.

4.1 Data

In 2015, the states of Ohio and Michigan both began collecting sales taxes on Amazon purchases as a result of existing or planned distribution and data centers in each state. On June 1, 2015, Ohio began collecting a 5.75% state tax with an additional varying local rate of up to 1.5%. Michigan followed on October 1, 2015 with a state tax rate of 6%. Ohio had previously required residents to self-report online purchases and remit sales taxes with state income tax returns, but enforcement was lax. Amazon, therefore, held a sizeable de facto pricing advantage over local retailers and any online merchants that did collect sales taxes up until this point. Thus, the implementation of the "Amazon Tax" represented an exogenous and immediate uniform increase in prices for Amazon customers in the affected states.

This chapter uses the same account and transaction data described in Chapter 3. Table 4.2 compares some of the broader spending patterns over selected months in the sample period representing a pre-treatment period (March 2015), the months following each state's implementation of the Amazon Tax, and one additional month well into the treatment period. The table provides a year-over-year view that encompasses the treatment for both Ohio and Michigan. These two states represent the majority of the sample while the remaining four states that are already collecting taxes on Amazon.com purchases act as a control group. Incomes are slightly higher in Ohio but average total debit Mastercard expenditures across all three groups are roughly similar.

At this level of detail, there is no discernible impact of the Amazon Tax. It is important to note, however, that the observed expenditure is the total amount charged to the payment method, not the sum of the retail prices of purchased items. To make an analogy to a standard retail receipt, total expenditures correspond to the pre-tax subtotal in the months prior to the implementation of the Amazon Tax. After the implementation of the Amazon Tax, total expenditures correspond to the final total that includes any applicable taxes to be collected. Without any change to purchasing patterns, the expectation is that total expenditures in the post-treatment period increase by the sales tax rate. To address this issue, I use both the raw expenditures and an additional tax-adjusted expenditure measure to better understand the change in purchasing patterns over time.

Transactions from Amazon and a selected set of major competing retailers are classified in the same fashion as in Chapter 3 and aggregated at the monthly level for the households in the sample. According to One Click Retail, a market research firm, the dominant product segments for Amazon are Consumer Electronics (2017 sales of $8.5B), Home and Kitchen ($5.5B), Publishing ($5B), and Sports and Outdoors ($4B).[3] The set of competitor retailer categories studied here correspond to various multi-category retailers and retailers specializing in the electronics and publishing categories. General retailers such as Wal-Mart form the largest category by sales. The other two multi-category classifications are department stores and warehouse stores such as Costco and Sam's Club. Market research by TWICE, another trade research firm, indicates that these categories along with the major specialty electronics retailers cover about 85%

[3]https://web.archive.org/web/20180104012358/http://oneclickretail.com/amazon-year-in-review-the-5-biggest-trends-of-2017/

Table 4.2: Summary Statistics by State, Selected Months

	March 2015	July 2015	November 2015	March 2016
Ohio				
Amazon Tax		X	X	X
Avg. Income	$4,859	$4,987	$4,226	$5,428
Avg. Tot. Debits	$2,069	$1,915	$1,961	$2,042
Amazon Expenditure	$17.57	$16.44	$22.04	$23.12
Amazon Share	0.85%	0.86%	1.12%	1.13%
Household Amz. Purch.	18.1%	16.9%	18.9%	21.1%
N	202,757	202,599	202,459	203,355
Michigan				
Amazon Tax			X	X
Avg. Income	$4,569	$4,676	$4,013	$5,081
Avg. Tot. Debits	$2,063	$1,923	$1,968	$2,033
Amazon Expenditure	$17.69	$16.13	$22.41	$23.84
Amazon Share	0.86%	0.84%	1.14%	1.17%
Household Amz. Purch.	17.5%	16.3%	18.7%	21.0%
N	58,769	58,780	58,674	58,723
Other States				
Amazon Tax	X	X	X	X
Avg. Income	$4,664	$4,723	$3,976	$5,059
Avg. Tot. Debits	$2,096	$1,942	$1,967	$2,054
Amazon Expenditure	$17.10	$16.02	$22.00	$22.54
Amazon Share	0.82%	0.82%	1.12%	1.10%
Household Amz. Purch.	16.8%	16.0%	18.3%	20.1%
N	60,027	59,918	59,787	60,861

Notes: Summary of monthly household data from February 2015 through January 2017 by designated geographic area. The Amazon Tax row indicates whether the geographic region was subsequent to sales tax collection on Amazon.com purchases in the specified month. Income approximated by total electronic (ACH) deposits. Total debits are the sum of all debit Mastercard transactions within the observation month. Amazon expenditures and shares are the total of the monthly debit Mastercard transactions for Amazon.com and Amazon Marketplace and the share of total debit card transaction amounts, respectively. Household Amazon Purchase rate is the proportion of households that made an Amazon.com or Marketplace purchase in the given month.

of all US electronics sales, for example.[4]

Transactions are also grouped according to the sales channel and whether or not the retailer is subject to collection of local sales taxes. Retail sales at a physical store are "local" sales, sales at an online retailer with a local physical presence are "online - taxed" sales, and other

[4]https://web.archive.org/web/20180522052412/https://www.twice.com/twice-research-center/2018-top-100-retailers-report-10-tech-dealers-dominate-industr

online retailers without a physical presence are classified as "online - not taxed." For Amazon purchases, transactions identified as sold by the "Amazon Marketplace" are assumed to be sales by third parties on the Amazon platform that do not necessarily collect state and local sales taxes. Purchases identified as "Amazon.com" purchases are shipped and sold by Amazon and are subject to their sales tax collection practices. These purchases are separated in the empirical analysis to capture any substitution away from the main Amazon.com platform that is subject to sales tax. Table 4.3 contains the full categorization of transactions.

Table 4.3: Classification of Transaction Data

Category	Retailers (w/ Sample Statement Description)
Amazon.com	Amazon.com (Amazon.com)
Amazon Marketplace	Amazon.com third-party sellers (AMAZON MKTPLACE)
General Retail - Local	Wal-Mart, Target, Kmart (WAL-MART #69324701 NOR)
General Retail - Online Taxed	walmart/target/kmart.com (WALMART.COM)
General Retail - Online Non-Taxed	jet/wish/overstock.com, AliExpress (ORDER.WISH.COM)
Warehouse - Local	Costco, Sam's Club (SAMSCLUB #8299)
Warehouse - Online Taxed	costco/samsclub.com (COSTCO.COM *ON)
Department - Local	Macy's, Kohl's, Sears, JCPenney (KOHLS 1529)
Department - Online Taxed	macys/kohls/sears/jcpenney.com (MACY*S .COM #01)
Electronics - Local	Best Buy (BEST BUY, 0008714)
Electronics - Online Taxed	bestbuy.com (BestBuyCom82715)
Electronics - Online Non-Taxed	newegg.com (WWW.NEWEGG.COM)
Bookstores - Local	Barnes & Noble (BARNESNOB 5711 Main St)
Bookstores - Online Taxed	barnesandnoble.com (BARNES&NOBLE*CO)

Finally, Amazon Prime subscriptions are captured as an independent merchant. This is a potentially important distinction because Prime subscriptions are an annual purchase and may indicate that a household is relatively more committed to the Amazon platform. A Prime subscription lowers average transaction costs through the free two-day shipping benefit on any

Figure 4.1: Distribution of Monthly Direct Deposit Income

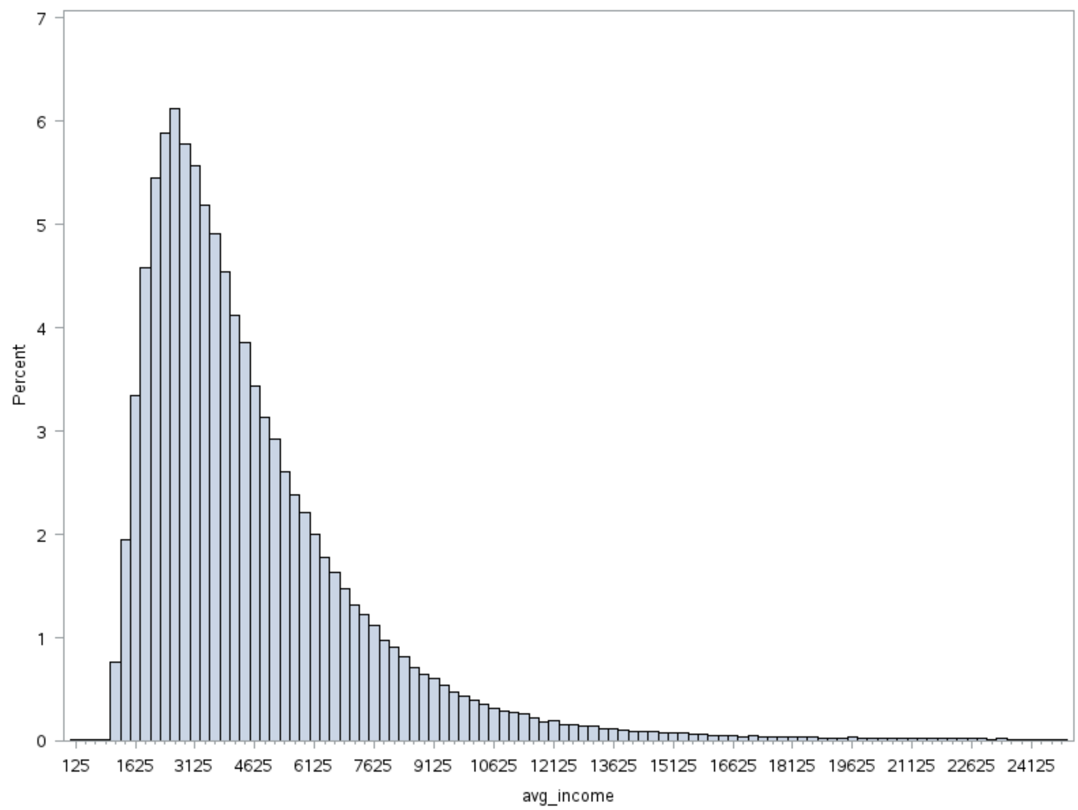

purchase. Otherwise, transactions must total \$35 to qualify for free shipping with a slightly longer delivery window. It is unclear ex ante how Prime memberships affect behaviour as the transaction size increases, which previous studies have shown to be the most sensitive to tax increases. Therefore, some of the empirical specifications allow for a separate effect for Prime subscribers.

The transaction data are supplemented with broader financial and demographic data for each household. The monthly financial data characterize account balances and cash flows. This includes checking and total deposit account balances, Automated Clearing House (ACH) debit and credit activity covering direct deposit income and scheduled bill payments, and broader data providing estimates of total household deposits across all financial institutions. The demographic information includes age, U.S. Census block group, and other neighborhood characteristics such as the total number of households, urban/suburban/rural classifications, and the number of retailers and total retail sales within 15 minutes driving time from a central location within the block group.

For this analysis, households are categorized into distinct geographic and income groups. Households are placed into income quartiles based upon average monthly direct deposit inflows over the two year sample period. These inflows are typically from wage income but regular government payments, retirement benefits, and other recurring electronic deposits are not distinguished. In addition, the payments are net of taxes and any other payroll deductions and therefore understate total compensation. As there is some ambiguity regarding the source of funds, 1,340 households with average annual deposits exceeding $300,000 are dropped from the sample. These households represent less than 1% of all households in the data. Figure 4.1 shows a histogram of the average monthly inflows for the remaining households. My analysis is primarily concerned with a household's relative income and so the effect of any measurement error due to payroll deductions should be limited. Since my focus is on explaining consumption decisions, it may even be preferable to use net pay rather than gross wages.

Table 4.4: Summary Statistics by Geographic Segment

	Rural	Suburban	Urban
Distribution by State			
OH	63.4%	61.7%	65.5%
MI	22.2%	18.3%	14.6%
Other	14.4%	20.0%	19.9%
County Pop.	49,990	188,722	386,362
Retail Sales ($M)	$746	$3,351	$6,649
Avg. Income	$4,565	$4,962	$4,261
Avg. Tot. Debits	$1,993	$2,000	$1,827
N	1,760,834	4,054,664	1,906,493

Notes: Summary of monthly household data from February 2015 through January 2017 by designated geographic area. Retail sales calculated as total retail sales within 15 min. drive time of household. Income approximated by total electronic (ACH) deposits. Total debits are the sum of all debit card transactions within the observation month.

The U.S Department of Agriculture Rural-Urban Commuting Areas form the basis for the geographic designations. This classification of census tracts is based on commuting patterns obtained from the 2006-2010 American Community Survey. For this analysis, households in census tracts designated as a metropolitan area core are classified as Urban if their tract has a business density greater than 150/mi^2 or is close to a major city center.[5] All other metropolitan area core census tracts are defined as Suburban. In addition, census tracts in micropolitan area cores (urban clusters of 10,000 to 50,000) with at least 15 businesses per square mile are also classified as Suburban. All other census tracts are designated as Rural. The geographic data are sourced from ESRI, a leading GIS software and data vendor.

The data are made up of monthly household observations from February 2015 to January 2016. However, to be included in the sample consumer (non-business) households must be active and in the data for the full 24 months from February 2015 through January 2017. Households must have total ACH credits in excess of $1000 (e.g., direct deposit of wages or social/retirement benefits) in 22 of the 24 months and average monthly debit Mastercard expenditures of at least $500.[6] These restrictions are made to limit the amount of attrition within

[5] Close in this case is defined as being within 15 miles of a city with population greater than 750,000, 8 miles for a city of at least 300,000, and 5 miles for a city of 200,000

[6] The decision to focus exclusively on debit Mastercard transactions is made for several reasons. Lower-income

the window of study. There must be at least one card transaction in each month to be included in the sample. To qualify for inclusion in the baseline sample, at least one purchase must have been made in the year preceding February 2015 on Amazon.com either directly or with a third-party Amazon Marketplace seller. Subsequent specifications relax this requirement and allow new Amazon customers to enter the sample data. Finally, households must be located in counties where the bank has a sufficient market presence. Within each county, the number of households satisfying the above criteria must represent at least 1% of total county households or all observations for that particular county are dropped from the sample.

Table 4.4 summarizes the data by geographic assignment. Michigan has a higher share of rural-designated households. Urban areas have the highest average annual retail revenues of the three segments, and about nine times that of rural areas. Even so, residents of urban areas have the lowest income and spending on average. The concern that rural households have a more limited number of retail choices within a reasonable driving distance is the primary motivation for the segmentation of the sample along this dimension.

Table 4.5: Summary Statistics by Income Quartile

	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Distribution by State				
OH	60.4%	62.5%	63.5%	65.7%
MI	19.6%	18.5%	18.1%	16.9%
Other	20.1%	19.0%	18.5%	17.5%
County Pop.	213,120	212,113	198,566	198,956
Retail Sales ($M)	$3,670	$3,629	$3,444	$3,525
Avg. Income	$2,042	$3,176	$4,586	$8,341
Avg. Tot. Debits	$1,204	$1,614	$2,074	$2,890
N	1,915,900	1,923,425	1,925,462	1,925,943

Notes: Summary of monthly household data from February 2015 through January 2017 by income quartile. Retail sales calculated as total retail sales within 15 min. drive time of household. Income approximated by total electronic (ACH) deposits. Total debits are the sum of all debit card transactions within the observation month.

Table 4.5 summarizes the sample data by income quartiles. As previously noted, Ohio has a slightly larger share of higher income households. There are no extreme differences between the income quartiles in terms of local population and retail sales. Debit Mastercard spending as a share of income ranges from 60% for Quartile 1 down to 35% for Quartile 4. This is consistent with higher savings rates for higher-income households, but may also reflect

households are underrepresented in the available credit card transaction data due to high bank underwriting standards. Another observational problem is that consumers are more likely to split their expenditures across multiple credit cards than multiple debit cards. In other words, one would be more likely to capture the full extent of card-based transaction activity by looking at a single active debit card account rather than a single active credit card. Finally, non-prepaid debit card usage dwarfs credit card usage by a factor of two and accounts for over half of all consumer and business non-cash payments (https://www.federalreserve.gov/newsevents/pressreleases/files/2016-payments-study-recent-developments-20170630.pdf).

a higher likelihood to use credit cards. As credit card transactions are not observed, total spending may be understated but does not affect inference if the Amazon Tax does not impact choice of payment method.

4.2 Empirical Strategy

The primary objective of this chapter is to establish facts regarding any disparate effects of increasing the relative cost of online purchases between locales or across differing wealth levels. In particular, households dependent upon public transportation or those in rural areas with limited local retail options may be disproportionately harmed if they cannot conveniently access the same variety of physical retail goods enjoyed by urban or more mobile populations.

The effect of the Amazon Tax is estimated using standard linear models of a selected set of spending activities such as category expenditure and purchase likelihood. The general empirical specification is:

$$y_{i,t} = \alpha + \beta_1 AmzTax_{c,t} \times Rural_i + \beta_2 AmzTax_{c,t} \times Suburban_i + \beta_3 AmzTax_{c,t} \times Urban_i + \beta_4 X_{i,t} + \beta_5 \delta_c + \beta_6 \lambda_t + \beta_7 \delta_c \times f(time)_t + e_i, \quad (4.1)$$

for each dependent variable of interest. In each case, i indexes individuals, c indexes counties, and t indexes time (month of sample period). $AmzTax$ is the treatment variable interacted with the development level (rural, suburban, urban) indicators. Because every household is assigned to one and only one development level in any particular month, no stand-alone treatment term is necessary. $X_{i,t}$ is a vector of individual characteristics comprised of a cubic age term and quadratic terms for the monthly income variable as well as total household deposit account balances. δ_c are county fixed effects and λ_t are monthly time dummies. A specification with household-level fixed effects is run in Section 4.4 as a robustness check on the baseline Amazon estimates, but county levels are chosen for the general specification in order to better identify substitution to other retailers. The final $\delta_c \times f(time)_t$ term is a county-specific linear time trend. $AmzTax$ is equal to 0 if the household is located in Ohio or Michigan in a month prior to the respective start of each state's sales tax collection on Amazon.com and 1 otherwise. The coefficients of interest are β_1, β_2, and β_3 which reveal the incremental impact of the Amazon Tax on the dependent variable. Following Abadie et al. (2017), standard errors are clustered at the state level, which is the level of aggregation at which the treatment is assigned. The preferred regression specification allows for all of the parameters to be separately estimated for each geographic or income grouping of interest (i.e., includes a full set of group interactions). In practice, separate regressions are run for each sub group. This conservative approach is possible because of the large sample size.

The identification strategy relies on the variation in treatment timing between Ohio and Michigan and uses the already treated states of Indiana, Pennsylvania, Kentucky, and West Virginia as controls. This unconventional approach is illustrated in the Goodman-Bacon (2018) decomposition of the standard difference-in-differences approach.[7] Under the strict but typical homogeneous effect and common trend assumptions these specifications can identify average

[7] Goodman-Bacon (2018, p. 6) notes that "already-treated units [..] can serve as controls even though they are treated because treatment status does not change."

treatment effects in the same way as one with a permanently untreated group. An additional specification that relies only on the differential treatment timing of Ohio and Michigan is also estimated as a robustness check.[8]

4.3 Results

Figure 4.2: Effect of Sales Tax Collection on Total Amazon.com Expenditures for Urban Households

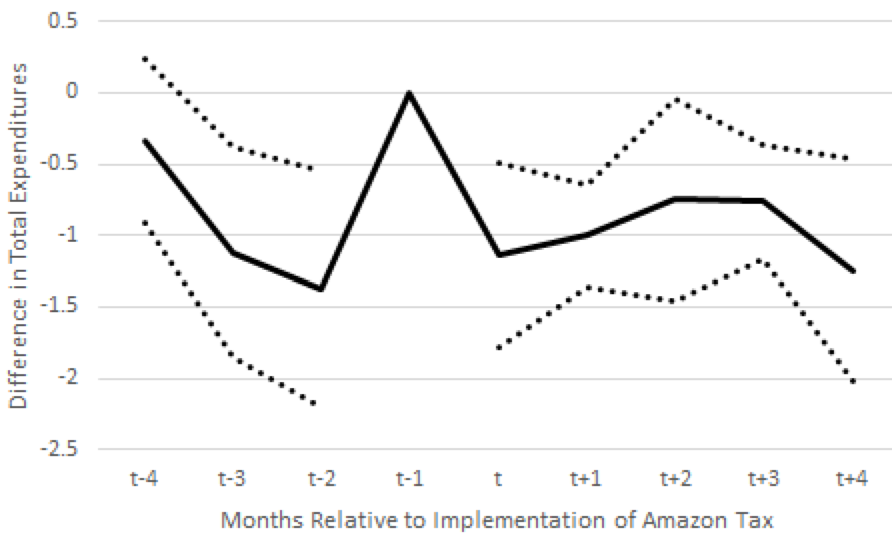

Figure 4.2 shows a coefficient plot of results from Equation 4.1 for urban households that made at least one purchase on Amazon.com or Amazon Marketplace at some point in the 12 months prior to the start of the sample period in February 2015. These are existing Amazon customers that are best positioned to reveal substitution patterns toward different retailers as they have already demonstrated a taste for online retail. The dependent variable is total monthly Amazon.com purchases (excluding untaxed Amazon Marketplace purchases). The y-axis represents the effect of sales tax collection on total expenditures. The x-axis shows the month relative to the implementation of the sales tax policy in Ohio and Michigan. t-4 is four months prior to the implementation of sales tax collection and t+4 is four months following the implementation of sales tax collection. The dashed lines form a 95% confidence interval around the coefficients. t-1 is the reference month with a zero coefficient.

The collection of sales tax does not appear to have a a significant effect on the total expenditures at Amazon.com of urban households. There is some noise, particularly around the implementation date, but for all time periods the estimate is relatively flat. The jump at t-1 does raise the possibility of households accelerating imminent purchases in order to avoid sales taxes.

[8]Goodman-Bacon (2018, p. 8) calls this the "timing-only" estimator.

Figure 4.3: Effect of Sales Tax Collection on Total Amazon.com Expenditures for Suburban Households

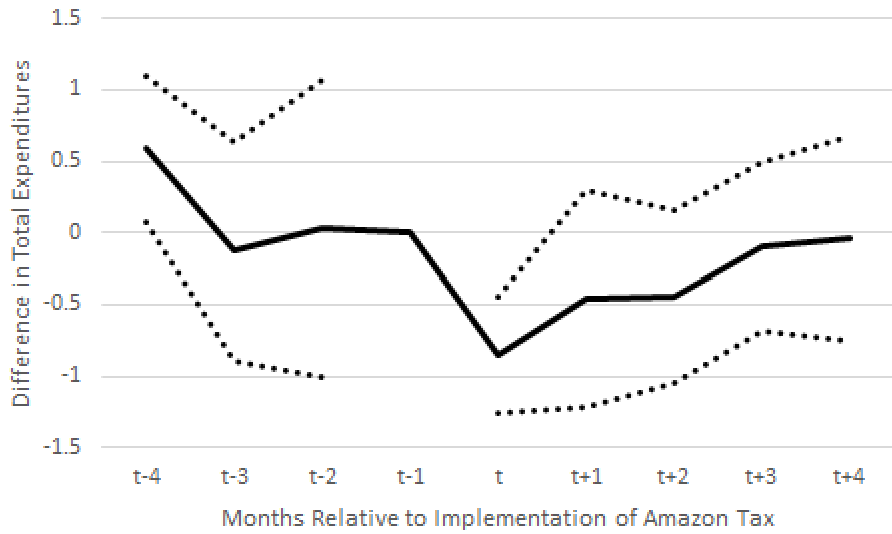

Figure 4.4: Effect of Sales Tax Collection on Total Amazon.com Expenditures for Rural Households

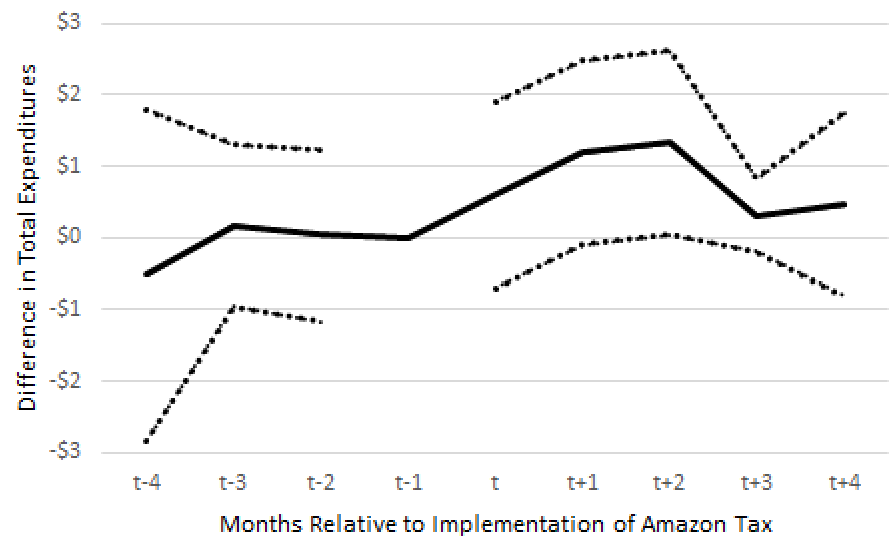

Figure 4.3 shows a similar coefficient plot for suburban households. The estimate appears to decline slightly after the implementation of the Amazon Tax, suggesting that these households may have gone to other retailers. Finally, Figure 4.4 features the plot for rural households. If anything, there appears to be a slight increase in rural household expenditures at Amazon after the tax is implemented.

Table 4.6 presents the complete results from the initial baseline regression of raw monthly household expenditures on Amazon.com. The effect of the Amazon Tax is considered separately for each location type aside from the general specification in Column 1. Each column adds control variables incrementally so that the contribution of each is more transparent, starting with only county and time fixed effects in Columns 1 and 2. Column 3 adds the additional controls for household attributes such as age, income, and total deposit balances. Column 4 introduces county-specific linear time trends, and Column 5 features the full interaction of local development level with all other model parameters.[9]

Table 4.6: Effect of Sales Tax Collection on Total Amazon.com Expenditures

	(1)	(2)	(3)	(4)	(5)
AmzTax	0.0155				
	(0.3581)				
AmzTax*Rural		0.0092	0.0252	0.0839	0.7502***
		(0.3280)	(0.3366)	(0.2691)	(0.1430)
AmzTax*Suburban		0.0875	0.0719	0.1843	-0.2760
		(0.3818)	(0.3889)	(0.2938)	(0.3459)
AmzTax*Urban		-0.1365	-0.2171	0.0571	0.3686***
		(0.4048)	(0.3807)	(0.3405)	(0.1572)
N	1,819,812	1,819,812	1,819,812	1,819,812	1,819,812
Dependent Mean	12.595				
Rural		12.859	12.859	12.859	12.859
Suburban		12.906	12.906	12.906	12.906
Urban		11.693	11.693	11.693	11.693
Controls:					
County FEs	X	X	X	X	X
Time FEs	X	X	X	X	X
Demographics			X	X	X
County Trends				X	X
Fully Interacted					X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

The estimated impact for rural households in the preferred specification featuring full development-specific controls (Column 5), is about 5.8% of the average total expenditures in the sample.[10] This is in the range of the actual sales tax rate applied and the expected effect if households did not change spending patterns at all. Urban households show a significant increase but to a lesser degree, while the change in suburban expenditures is not statistically significant.[11] It can be seen from the total average expenditures that rural and suburban house-

[9] Evidence regarding the satisfaction of the common trends assumption and adequacy of the control group is contained in the appendix.

[10] This is obtained by dividing the AmzTax*Rural coefficient estimate of 0.7502 by the dependent mean of 12.859.

[11] I find similar results for the tax-adjusted measure of expenditures, which applies the sales tax rate to Ama-

Figure 4.5: Effect of Sales Tax Collection on Adjusted Amazon.com Expenditures for Urban Households

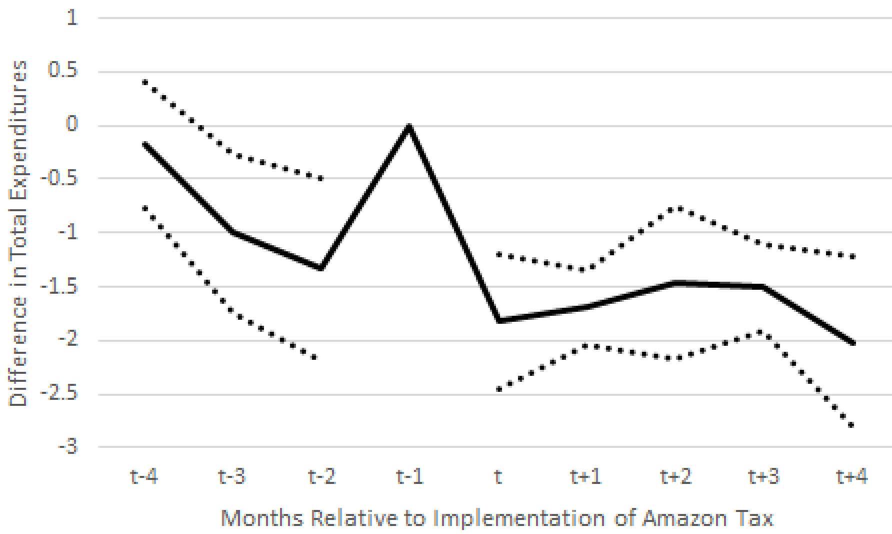

holds spend approximately the same amount at Amazon.com overall despite rural households earning approximately 8% less in take home pay as shown in Table 4.4.

4.3.1 Tax-adjusted Expenditures

An alternative specification involves an adjustment to the aggregated Amazon.com expenditures in an attempt to better reflect changes in the consumption value of goods purchased. Baugh et al. (2018) discount the expenditures observed after Amazon.com implements sales tax collection within an individual state by the appropriate local sales tax rate. This is deemed the "tax-exclusive" expenditure and can be thought of as a more consistent proxy for consumption value when making comparisons over time and across states. Another benefit of adjusting for sales taxes is that it creates a firm zero upper bound on the expected effect of the Amazon Tax if all other aspects of the shopping experience remain the same.

One potential issue with this method, however, is that it assumes no differential pricing practices. If Amazon were to target discounts specifically to customers in Ohio and Michigan after sales tax collection began, a consumption value interpretation of this measure would overstate the impact of sales taxation. Since these types of pricing practices are naturally very closely-guarded, little can be done but to note this potential source of bias. While this may lead to an overstatement of the impact of sales tax collection on the population as a whole, it is less likely to affect the inference of differential effects within the population, my primary focus. Regardless, both expenditure measures are considered as part of the baseline empirical results.

Due to the unconventional approach of using already-treated states as controls, I use a

zon.com purchases occurring prior to the start of tax collection. Suburban households reducing expenditures and no change is detected for rural households.

Figure 4.6: Effect of Sales Tax Collection on Adjusted Amazon.com Expenditures for Suburban Households

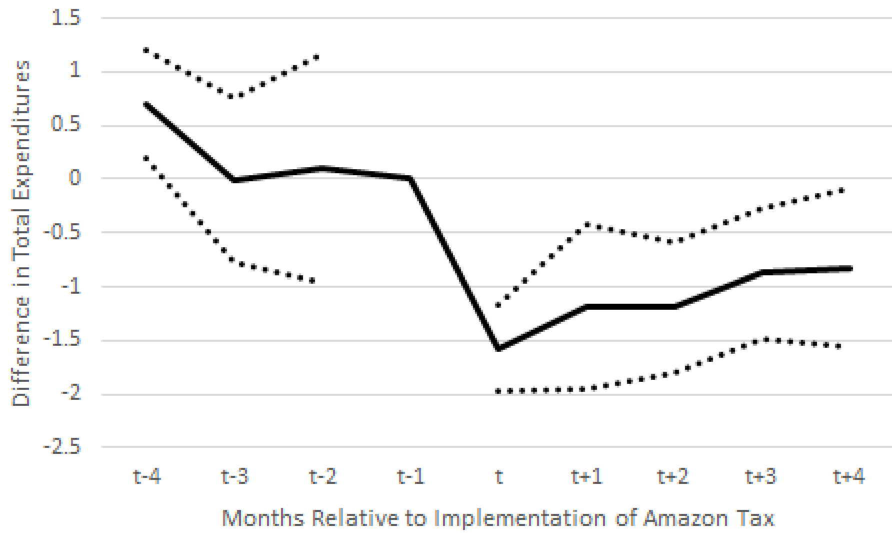

slightly different method for tax adjustment. Rather than discounting the expenditures by the appropriate local tax rate after sales tax collection begins, the local sales tax rate is instead applied to expenditures in the pre-treatment periods for residents of Ohio and Michigan. With already-treated controls, the resulting after-tax measure is easier to calculate in this case and simpler to comprehend. Sales tax is equally applied across all states and time periods, regardless of whether it was actually collected.

The baseline method is now repeated for the new tax-adjusted expenditures. Substitution away from Amazon.com appears as a negative estimated coefficient. Figures 4.5, 4.6, and 4.7 show the coefficient plots of results from Equation 4.1 for the same set of households that made at least one purchase on Amazon.com or the Amazon Marketplace at some point in the 12 months prior to the start of the sample period in February 2015 for each geographic segment, respectively. Again, a sharp spike at t-1 is observed for urban households, suggesting some anticipation of tax increases. Curiously, this does not occur for rural or suburban households. Suburban households demonstrate the clearest indication of substitution away from Amazon.com purchases following the implementation of the Amazon Tax. Rural households on the other hand appear to be completely unaffected.

Table 4.7 likewise presents complete results from the initial baseline regression of tax-adjusted monthly household expenditures on Amazon.com. The effect of the Amazon Tax is considered separately for each location type aside from the general specification in column 1. Column 1 indicates an overall decline roughly equivalent to the increase in sales taxes, consistent with the first baseline analysis of raw Amazon expenditures that found no significant effect. Again, the fully-interacted specification (Column 5) reveals some differential effects across the development levels. Rural households appear to consume the same amount of goods regardless of the tax treatment. In contrast, suburban households show strong evidence of taking their spending elsewhere. These households have the highest income of all the development groups in the sample so an explanation relying purely on price sensitivity is difficult to form.

Figure 4.7: Effect of Sales Tax Collection on Adjusted Amazon.com Expenditures for Rural Households

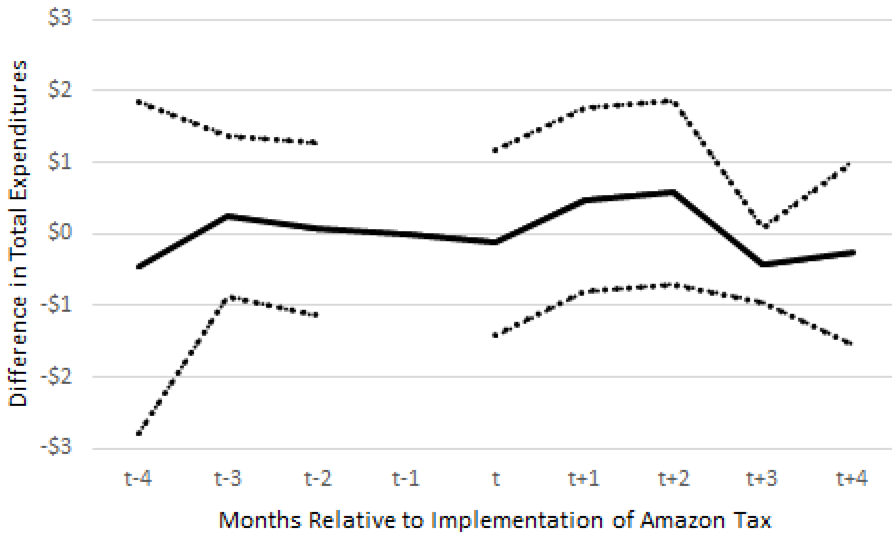

From a substitution standpoint, suburban households may have the most options. On the one hand, rural households have little in the way of geographically convenient brick and mortar retail. Urban households, on the other hand, are generally more reliant on public transportation, which may impose additional limits compared to a suburban life built around ample personal transportation. The avenues for potential substitution are analyzed subsequently.

4.3.2 Substitution to Other Retailers

Do Amazon consumer households increase their spending with other retailers once state and local sales taxes are levied? To find out, the fully interacted specification (Column 5 of previous tables) of Equation 4.1 is estimated for competitors in some of Amazon's most important product categories. Substitution to Amazon's own Marketplace is a consideration as well. The majority of Marketplace purchases are with third parties and during the sample period do not have tax applied if the seller is not based in the same state as the purchaser. For price-sensitive Amazon consumers, it might be expected that the Marketplace is the next best option. Many of the Marketplace sales are also Fulfilled by Amazon (FBA) sales which are eligible for the standard Amazon shipping policies (free with $35 total purchase or Prime membership). Amazon also introduced Seller Fulfilled Prime in 2015, allowing merchants to ship from their own warehouse or use third party logistics providers under the Prime shipping banner. Extending Prime shipping benefits to these Amazon Marketplace transactions may have provided an additional boost to Amazon's efforts to retain more price-sensitive consumers.

After the Marketplace, other electronics retailers are the natural next place to look as electronics is the category with the highest revenues for Amazon in this sample period. This segment, in particular, is instructive since local retail sales, taxable online sales, and non-taxable online sales are all separately observed by purchases at brick and mortar Best Buy stores, Best Buy online purchases, and Newegg purchases, respectively. Newegg.com is a leading specialty

Table 4.7: Effect of Sales Tax Collection on Tax-adjusted Amazon.com Expenditures

	(1)	(2)	(3)	(4)	(5)
AmzTax	-0.7226*				
	(0.3726)				
AmzTax*Rural		-0.7326**	-0.7159**	-0.6551***	0.0234
		(0.3325)	(0.3409)	(0.2698)	(0.1440)
AmzTax*Suburban		-0.6579*	-0.6736*	-0.5514*	-1.0193***
		(0.3985)	(0.4060)	(0.3055)	(0.3517)
AmzTax*Urban		-0.8548*	-0.9372**	-0.6279*	-0.3049*
		(0.4390)	(0.4130)	(0.3657)	(0.1757)
N	1,819,812	1,819,812	1,819,812	1,819,812	1,819,812
Dependent Mean	12.838				
Rural		13.127	13.127	13.127	13.127
Suburban		13.148	13.148	13.148	13.148
Urban		11.918	11.918	11.918	11.918
Controls:					
County FEs	X	X	X	X	X
Time FEs	X	X	X	X	X
Demographics			X	X	X
County Trends				X	X
Fully Interacted					X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

online consumer electronics retailer that did not have operations in Ohio or Michigan during the sample period and is therefore not subject to sales tax collection. Best Buy has both a physical and online presence. Therefore, online purchases are subject to the same state and local sales tax regime as an in-store purchase.

Substitution to other local and online publishing retailers is examined next. Barnes & Noble is the primary Amazon competitor for the households in the sample data, and transaction descriptions helpfully distinguish online sales. Warehouse stores are the last major retail sector to be included and feature separable online and in-store transactions as well. Costco and Sam's Club are the primary retailers in this space. Finally, online general retailers such as Overstock and Jet are included as potential beneficiaries of an increase in the cost of goods delivered from Amazon. Again, these retailers are domiciled outside the states of Ohio and Michigan and not subject to sales tax collection.

In general, if pure cost minimization was the dominant factor in the decision to substitute away from Amazon purchases, the most likely destination would be the competing online retailers that do not have operations in Ohio or Michigan (e.g. Newegg, Overstock, and Jet). These are the businesses that maintain the significant pricing advantage from not collecting state and local sales taxes at the time of purchase. If consumers still value the in-store shopping experience or the ability to take immediate possession of purchased goods, brick and mortar retailers may benefit the most. Finally, one scenario that would uniquely benefit taxable online retail would be if consumers highly valued the convenience of shopping from home but also the ability to quickly visit a brick and mortar retail outlet in the event of any problems with their purchase.

Table 4.8: Amazon Tax Substitution Effects on Amazon Marketplace and Electronics

	Amazon Marketplace	All Elec.	Local Elec.	Online Elec. - Taxed	Online Elec. - Untaxed
AmzTax*Rural	-0.2906**	0.3658***	0.1961	0.0838***	0.0859
	(0.1438)	(0.0806)	(0.1234)	(0.0286)	(0.0534)
AmzTax*Suburban	-0.6224***	0.4137	0.3705	0.0532	-0.0100
	(0.2291)	(0.3124)	(0.2903)	(0.0586)	(0.0748)
AmzTax*Urban	0.3110*	0.3694*	0.2651	0.0070	0.0973
	(0.1777)	(0.2200)	(0.2774)	(0.0719)	(0.0953)
N	1,819,812	1,819,812	1,819,812	1,819,812	1,819,812
Dependent Mean					
Rural	24.411	8.670	7.223	0.852	0.595
Suburban	23.647	11.228	9.613	1.117	0.498
Urban	21.873	9.441	8.077	0.921	0.443
Controls:					
County FEs	X	X	X	X	X
Time FEs	X	X	X	X	X
Demographics	X	X	X	X	X
County Trends	X	X	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

Table 4.8 captures any substitution effects within the Amazon platform toward more Marketplace expenditures and to competing electronics retailers. Expenditures significantly decline for rural and suburban households, indicating a potential shift away from the Amazon ecosystem in general. Urban households slightly increase expenditures. All changes are statistically significant but represent a very small share of average household expenditures on the Amazon Marketplace. Electronics retailers see a significant increase in expenditures among rural and urban households that is slightly more meaningful in an economic sense. However within the specific channels (brick and mortar, taxed online, and untaxed online) there is not enough power to identify any significant effects save for a meaningful increase in taxable online sales for rural households.

Table 4.9 captures any substitution effects of the Amazon Tax favoring bookstores, warehouse stores, and other general online retailers not required to collect sales taxes. Local bookstores such as Barnes & Noble are a significant beneficiary of the implementation of the Ama-

Table 4.9: Amazon Tax Substitution Effects on Bookstores, Warehouse Stores, and Untaxed General Retail

	Local Bookstores	Online Bookstores	Local Warehouse	Online Warehouse	Online Gen. Retail - Untaxed
AmzTax*Rural	0.1416***	0.0203**	-0.1849	0.0967***	0.0408
	(0.0445)	(0.0096)	(0.2957)	(0.0293)	(0.0727)
AmzTax*Suburban	0.2625***	-0.0218***	0.8806***	0.0286	0.0784***
	(0.0300)	(0.0066)	(0.2941)	(0.0269)	(0.0138)
AmzTax*Urban	0.0811***	-0.0289	1.3457***	0.2288***	0.1571***
	(0.0332)	(0.0258)	(0.3744)	(0.0555)	(0.0658)
N	1,819,812	1,819,812	1,819,812	1,819,812	1,819,812
Dependent Mean					
Rural	1.481	0.311	33.719	0.180	0.765
Suburban	2.201	0.327	46.015	0.314	0.920
Urban	2.103	0.273	35.514	0.261	1.055
Controls:					
County FEs	X	X	X	X	X
Time FEs	X	X	X	X	X
Demographics	X	X	X	X	X
County Trends	X	X	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

zon Tax. However, there is little in the way of a meaningful impact on online book sales. Warehouse stores see a significant boost from suburban and urban households in stores. Online sales to rural and urban households jump significantly as well, albeit from a very low base. Moving on to the online general retailers, the only group that does not collect sales taxes, rural households are the only development group that did not see a significant increase in expenditures.

At this point, a few patterns have emerged. First, the fact so many shifts are observed at all indicates that Amazon did indeed benefit at least in some small part from the avoidance of collecting state and local sales taxes. Second, the substitution patterns observed are in line with expectations regarding household mobility and proximity to retail options. Suburban households generally appear to prefer retail stores. Rural households, to the extent that there was any substitution, appear to continue to utilize online channels with a wider range of retailers. Surprisingly, the same rural households did not significantly increase their expenditure at untaxed general online retailers like urban and suburban households. This suggests that sales tax avoidance may not be a primary motivation for rural households.

4.3.3 Amazon Prime Membership

Amazon Prime has been a core component of Amazon's consumer strategy since its introduction in 2005 as a membership-style product providing free unlimited two-day shipping on any order sold and shipped by Amazon for an annual $79 fee. Membership benefits have grown over time to encompass video and audio streaming, e-book rentals, same-day shipping and delivery in some areas, and access to member-only sales. While Prime was once somewhat comparable to a membership at warehouse-style retailers such as Costco and Sam's Club, the current form is difficult to classify.[12] Regardless, this unique vehicle for acquiring customers and a larger share of their expenditures could have an important impact on consumer purchasing behaviour.

Table 4.10: Effects of Amazon Tax on Amazon Purchases by Prime Membership Status

	Amazon Purchases		Tax-adjusted Purchases	
	Prime	No Prime	Prime	No Prime
AmzTax*Rural	0.8234*	0.7562***	-0.7654	0.2298**
	(0.4694)	(0.1077)	(0.4732)	(0.1070)
AmzTax*Suburban	0.0413	-0.3138	-1.5680***	-0.8481***
	(0.4989)	(0.3282)	(0.5225)	(0.3266)
AmzTax*Urban	1.2608***	0.2037	-0.2423	-0.2690
	(0.2710)	(0.1792)	(0.3015)	(0.1910)
N	354,941	1,464,871	354,941	1,464,871
Dependent Mean				
Rural	27.036	9.548	27.619	9.743
Suburban	27.179	9.414	27.709	9.585
Urban	25.153	8.395	25.658	8.553
Controls:				
County FEs	X	X	X	X
Time FEs	X	X	X	X
Demographics	X	X	X	X
County Trends	X	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

Earlier studies of sensitivity to sales tax on Amazon purchases did not make a distinction between members and non-members. However, as membership perks grew over time it has become more important to account for. Such perks include: Prime Music, an unlimited ad-free streaming service was launched in the spring of 2014; the first Prime Day, a special members-

[12]In his 2016 letter to shareholders, founder Jeff Bezos called Prime an "all-you-can-eat, physical-digital hybrid that members love." (https://www.sec.gov/Archives/edgar/data/1018724/000119312516530910/d168744dex991.htm)

only sale, occurred on July 15, 2015; and Amazon undertook a major expansion of the streaming video service at the end of the 2019. Prime could almost be considered a fledgling service at the time the data for prior studies were created compared to early 2016 when the sample period ends.

To study the effect of taxes on Prime members, households that had a transaction labeled as an Amazon Prime transaction in the year preceding February 2015 are classified as members. This is an imperfect measure as some households may gift or be gifted Prime membership, but it approximates the true membership status well enough for a group-level study. Roughly one fifth of these 2015 Amazon customers are classified as Prime members. As might be expected, Prime members spend much more at Amazon, about three times as much as non-Prime members.

Table 4.10 presents the baseline estimates for the impact of the Amazon Tax on raw Amazon.com expenditures as well as the tax-adjusted measure. As with the substitution results, the fully-interacted specification is chosen and the regressions are run separately for Prime members and the remaining Amazon customers that are not Prime members. Splitting out Prime members results in some additional noise in the estimates and there does not appear to be any broad meaningful differences from other Amazon customers in terms of the response to the application of sales taxes.

4.3.4 Effects by Income

The previous section identified significant differences across development groups in the household response to the Amazon Tax. This section adds income as a complementary factor likely to influence substitution behavior interactively in conjunction with the household's location. For example, urban amenities such as public transportation could better enable low-income households to patronize brick and mortar retailers compared to rural households that would be more reliant on personal transportation.

Tables 4.11 and 4.12 repeat the baseline estimation for the impact of the Amazon Tax on raw Amazon.com expenditures and the tax-adjusted measure, respectively. Again, the population studied is consumers that made a purchase on Amazon.com in the year prior to February 2015. The fully-interacted specification is used, with separate regressions for each income quartile. The estimates for each quartile are contained in the four columns of each table, with columns 1-4 containing income quartiles 1-4, respectively.

Across income groups, rural households generally increase total expenditures on Amazon.com following the implementation of sales tax collection. Suburban households again appear to be the most likely to substitute away. In general, third income quartile households show the largest declines. These households are likely more price-sensitive than the highest income households while also commanding more resources that widen the available options compared to lower-income households.

Of the substitution options, there are several examples that illustrate some differential effects across income quartiles within development levels. Table 4.13 contains estimates of the substitution to Amazon Marketplace using the same approach as the previous two tables. Among rural households, third and fourth income quartile households appear to reduce activity on the Amazon platform altogether. In contrast, first and second income quartile households

Table 4.11: Effect of Amazon Tax on Total Amazon.com Expenditures by Income

	Income Q1	Income Q2	Income Q3	Income Q4
AmzTax*Rural	0.5100***	1.1111***	0.2653*	1.2280***
	(0.1266)	(0.0813)	(0.1406)	(0.3033)
AmzTax*Suburban	-0.3377	-0.1970*	-0.5186	0.0071
	(0.2822)	(0.1056)	(0.3967)	(0.5777)
AmzTax*Urban	0.4910**	0.2242	0.0628	0.5381***
	(0.2162)	(0.2250)	(0.5810)	(0.1879)
N	351,892	424,195	480,855	554,837
Dependent Mean				
Rural	8.524	10.040	12.340	18.371
Suburban	7.699	9.766	12.338	18.510
Urban	7.238	9.330	12.700	18.304
Controls:				
County FEs	X	X	X	X
Time FEs	X	X	X	X
Demographics	X	X	X	X
County Trends	X	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

appear to increase purchases. The same appears to be the case for urban households, although the highest income quartile also increases purchases on Amazon Marketplace post-treatment.

Warehouse stores also exhibit differential substitution effects across income within development level groups. As a suburban staple, these stores are a popular destination for households to make large bulk purchases. Table 4.14 contains estimates for the impact on expenditures at warehouse stores. Rural households appear to have a split effect, with the below-median income households clearly showing no response. This is in contrast to urban and suburban lower-income households that appear to exhibit some signs of increased expenditures in the local brick and mortar stores.

Finally, Table 4.15 contains estimates for brick and mortar electronics retailers. As with warehouse stores, below-median income households (Quartiles 1 and 2) do not show any evidence of substitution effects, while the above-median income households significantly increase purchases in stores. Taken together, the observed behavioural changes across different income groups in response to the implementation of the Amazon Tax suggests that below-median income rural households are the most constrained in terms of their alternatives to purchases within the Amazon ecosystem. These households show little ability to substitute away from Amazon.com purchases, increase purchases on the Amazon Marketplace, and do not patronize brick and mortar retailers to a larger degree as the higher-income rural households do.

Table 4.12: Effect of Amazon Tax on Tax-adjusted Amazon.com Expenditures by Income

	Income Q1	Income Q2	Income Q3	Income Q4
AmzTax*Rural	0.0562	0.5878***	-0.4496***	0.1364
	(0.1329)	(0.0812)	(0.1408)	(0.2999)
AmzTax*Suburban	-0.7900***	-0.7487***	-1.2547***	-1.0496*
	(0.2864)	(0.1108)	(0.3976)	(0.5798)
AmzTax*Urban	0.0659	-0.3130	-0.6391	-0.5623***
	(0.2029)	(0.2461)	(0.6054)	(0.1822)
N	351,892	424,195	480,855	554,837
Dependent Mean				
Rural	8.705	10.249	12.601	18.751
Suburban	7.851	9.954	12.577	18.842
Urban	7.381	9.520	12.948	18.643
Controls:				
County FEs	X	X	X	X
Time FEs	X	X	X	X
Demographics	X	X	X	X
County Trends	X	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

4.3.5 New Amazon Customers

Until now the focus has been on households that made at least one Amazon purchase in the year leading up to the start of the sample period. In this section, the behaviour of households in the data that did not make an Amazon purchase in the year preceding the start of the sample period in 2015 but satisfy all other sample inclusion criteria are studied. For this population Equation 4.1 is estimated for raw Amazon expenditures, tax-adjusted Amazon expenditures, and Amazon Marketplace expenditures once again using the same fully-interacted specification.

Table 4.16 contains the results for new-to-Amazon customers. The first column corresponds to the raw monthly Amazon.com expenditures. The second column contains the same tax-adjusted Amazon.com expenditures that have been used to this point. The third column displays the impact on Amazon Marketplace purchases. Surprisingly, rural households appear to be the most affected by the Amazon Tax, although an increase in Amazon Marketplace expenditures matches the slight decline in Amazon.com purchases. Overall, there does not appear to be any major impact to new customer acquisition for Amazon in the wake of sales tax increases, particularly outside rural areas.

Table 4.13: Effect of Amazon Tax on Amazon Marketplace Expenditures by Income

	Income Q1	Income Q2	Income Q3	Income Q4
AmzTax*Rural	0.7235	1.1648***	-1.1246***	-1.1227***
	(0.5612)	(0.2716)	(0.1898)	(0.1542)
AmzTax*Suburban	-0.6871	-1.0140***	-0.9058***	-0.1179
	(0.4530)	(0.1611)	(0.3765)	(0.2994)
AmzTax*Urban	0.7497***	0.6729***	-1.4114	1.3218**
	(0.2153)	(0.2476)	(1.2865)	(0.6253)
N	351,892	424,195	480,855	554,837
Dependent Mean				
Rural	17.247	20.080	23.455	33.368
Suburban	15.247	18.922	23.025	32.223
Urban	14.861	18.241	23.691	31.952
Controls:				
County FEs	X	X	X	X
Time FEs	X	X	X	X
Demographics	X	X	X	X
County Trends	X	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

4.4 Robustness Checks

4.4.1 Household Fixed Effects Estimates

As mentioned in Section 4.2, county fixed effects are chosen for the baseline specification in order to preserve statistical power for identifying widely dispersed substitution effects after the beginning of sales tax collection on Amazon.com. In order to satisfy any concerns about unobserved household-level attributes, a household-level fixed effects variant of the original specification is now estimated. Similar to before, the fixed effects empirical specification is:

$$y_{i,t} = \alpha_i + \beta_1 AmzTax_{c,t} \times Rural_i + \beta_2 AmzTax_{c,t} \times Suburban_i + \beta_3 AmzTax_{c,t} \times Urban_i + \beta_4 X_{i,t} + \beta_5 \delta_c + \beta_6 \lambda_t + \beta_7 \delta_c \times f(time)_t + e_i, \quad (4.2)$$

for the three core dependent variables of interest: raw Amazon.com expenditures, tax-adjusted Amazon.com expenditures, and Amazon Marketplace expenditures. The variables are the same as in Equation 4.1 except for α_i, the new household fixed effect term.

Equation 4.2 is estimated for these three dependent variables using the preferred fully-interacted specification equivalent to separate regressions for each development level. The baseline population (all active Amazon customers in the year preceding February 2015) featured in Section 4.3 remains the focus of analysis.

Table 4.14: Effect of Amazon Tax on Local Warehouse Store Expenditures by Income

	Income Q1	Income Q2	Income Q3	Income Q4
AmzTax*Rural	-0.6580***	-0.6647*	1.0766**	-0.4410
	(0.1904)	(0.3780)	(0.4970)	(0.5443)
AmzTax*Suburban	0.3001	0.5152	0.3672	1.9487***
	(0.2451)	(0.3823)	(0.7539)	(0.4861)
AmzTax*Urban	0.4796	1.2967***	1.0881***	2.0837
	(0.3540)	(0.3065)	(0.3177)	(1.6846)
N	351,892	424,195	480,855	554,837
Dependent Mean				
Rural	14.049	22.362	32.754	56.510
Suburban	17.755	27.572	43.139	77.824
Urban	16.077	24.607	38.922	67.127
Controls:				
County FEs	X	X	X	X
Time FEs	X	X	X	X
Demographics	X	X	X	X
County Trends	X	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

Table 4.17 contains the results for the households making an Amazon purchase prior to February 2015. The first column corresponds to the raw monthly Amazon.com expenditures. The second column contains the same tax-adjusted Amazon.com expenditures. The third column displays the impact on Amazon Marketplace purchases. The results are broadly consistent with the previous results. As before, suburban households show the most evidence of substitution away from the Amazon platform. Rural households again demonstrate very little evidence of substitution. Urban households are in the middle. The estimated coefficient for urban substitution effects for Amazon Marketplace purchases is no longer significant, but all coefficient estimates are similar in magnitude to those obtained using county fixed effects.

4.4.2 Increased Amazon Purchase Inclusion Criteria

Baugh et al. (2018) restrict the baseline sample of Amazon customers to those with more than $200 in annual expenditures as opposed to the more inclusive baseline requirement of any Amazon purchase used in Section 4.3. If lower-spending households are less attached to the Amazon platform, it is possible that the estimated impact of sales taxes could be overstated. Table 4.18 presents estimates for this smaller sub-population's change in raw Amazon expenditures, tax-adjusted Amazon expenditures, and Amazon Marketplace expenditures using the preferred fully-interacted specification.

Table 4.15: Effect of Amazon Tax on Local Electronics Store Expenditures by Income

	Income Q1	Income Q2	Income Q3	Income Q4
AmzTax*Rural	-0.6540***	-0.5079***	0.7148***	0.6336*
	(0.1831)	(0.2067)	(0.1724)	(0.3303)
AmzTax*Suburban	0.9249***	-0.7819***	0.3350	0.9065***
	(0.2985)	(0.0472)	(0.5243)	(0.2453)
AmzTax*Urban	-0.7045***	1.1529*	0.5360	-0.0260
	(0.1680)	(0.6457)	(0.4884)	(0.3412)
N	351,892	424,195	480,855	554,837
Dependent Mean				
Rural	4.030	5.087	6.665	11.392
Suburban	5.021	7.101	9.177	14.353
Urban	5.006	6.566	8.808	12.533
Controls:				
County FEs	X	X	X	X
Time FEs	X	X	X	X
Demographics	X	X	X	X
County Trends	X	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

Again, the results for this set of customers are broadly consistent with the previous results. As before, suburban households show the most evidence of substitution away from the Amazon platform. Rural households again demonstrate very little evidence of substitution. Urban households are in the middle, showing some evidence of within-platform substitution to the Amazon Marketplace. However, the estimated effects for Amazon Marketplace are all very small relative to the total average monthly expenditure.

4.4.3 Timing-only Estimator

The unconventional nature of the differences in differences specification in which the treatment effect is estimated by calculating the effect of "non-treatment on the treated" may lead to some questions regarding the identification of the treatment effect. Goodman-Bacon (2018) provides a general characterization of the sources of identifying variation in standard difference-in-differences (DD) models. The primary finding is that the estimated average treatment effect for models with variation in treatment timing is a weighted average of all possible pairwise comparisons at all points in time between groups that change treatment status and those that do not. Groups that are treated late act as controls for groups treated early, but groups that are treated early can also act as controls for groups treated late. This characteristic broadly impacts DD models featuring variation in treatment timing. An additional identifying assump-

Table 4.16: Effect of Amazon Tax on New Amazon Customer Purchases

	Raw	Tax-adjusted	Marketplace
AmzTax*Rural	-0.0729**	-0.1556***	0.0987*
	(0.0333)	(0.0362)	(0.0583)
AmzTax*Suburban	-0.0213	-0.0960**	0.1955***
	(0.0450)	(0.0425)	(0.0451)
AmzTax*Urban	0.0061	-0.0623	0.0951
	(0.0681)	(0.0669)	(0.0701)
N	2,013,259	2,013,259	2,013,259
Dependent Mean			
Rural	1.389	1.414	2.991
Suburban	1.479	1.502	3.073
Urban	1.396	1.417	3.015
Controls:			
County FEs	X	X	X
Time FEs	X	X	X
Demographics	X	X	X
County Trends	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

tion of time-invariant treatment effects is therefore needed, regardless of the explicit control framework. The major implication for this study is that the reliance on already-treated states to act as controls as opposed to states that remain untreated throughout the sample period (e.g., the Baugh et al. (2008) framework) does not involve any additional identifying assumptions relative to previous studies.

As an additional test of robustness and to contrast the full DD treatment effect estimate with the "timing-only" estimator presented in Goodman-Bacon (2018), the full baseline set of regressions are re-estimated for the states of Ohio and Michigan alone. In the baseline case with a single standalone treatment indicator, each state's contribution to the treatment effect estimate is based upon the variance of the treatment indicator. The states have equal weight as Ohio has four months of observed pre-treatment, while Michigan has four months of observed post-treatment in the 12 month sample.

Tables 4.19 and 4.20 present the complete results from the initial baseline regressions of raw and tax-adjusted monthly household expenditures on Amazon.com, respectively. The effect of the Amazon Tax is considered separately for each development level aside from the most general specification in Column 1. Each column adds control variables incrementally so that the contribution of each is more transparent, starting with only county and time fixed effects in Columns 1 and 2. Column 3 adds the additional controls for household attributes; age, income,

Table 4.17: Effect of Amazon Tax - Household Fixed Effects Specification

	Raw	Tax-adjusted	Marketplace
AmzTax*Rural	0.6930***	-0.0338	-0.3704***
	(0.1472)	(0.1479)	(0.1473)
AmzTax*Suburban	-0.2860	-1.0293***	-0.6120***
	(0.3279)	(0.3336)	(0.2244)
AmzTax*Urban	0.3870**	-0.2843	0.2869
	(0.1776)	(0.1961)	(0.1888)
N	1,819,812	1,819,812	1,819,812
Dependent Mean			
Rural	12.859	13.127	24.411
Suburban	12.906	13.148	23.647
Urban	11.693	11.918	21.873
Controls:			
County FEs	X	X	X
Time FEs	X	X	X
Demographics	X	X	X
County Trends	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

and total deposit balances. Column 4 introduces county-specific linear time trends, and Column 5 features the full interaction of local development level with all other model parameters.

The final specification containing the fully-interacted development level indicators is very close to the original baseline estimates from Section 4.3 for both raw and tax-adjusted Amazon.com expenditures, but generally indicates smaller levels of substitution away from Amazon.com. In terms of differential effects, the relative differences between rural, suburban, and urban households remain the same. Given that the states of Ohio and Michigan comprise 80% of the sample, the concordance between the timing-only and general DD estimator are not surprising in light of the Goodman-Bacon decomposition theorem.

4.4.4 Limited-duration Effects

There may be some concerns that the baseline specification does not allow for variation in the effect of sales taxes on Amazon purchases over time. If there is a short-term effect that decayed over time, it may not be identified using Equation 4.1. As a robustness check, the following

Table 4.18: Effect of Amazon Tax on Existing High-value Amazon Customer Purchases

	Raw	Tax-adjusted	Marketplace
AmzTax*Rural	1.2552***	-0.0691	-0.5380***
	(0.2537)	(0.2588)	(0.2299)
AmzTax*Suburban	-0.7004	-2.1061***	-1.6771***
	(0.7269)	(0.7438)	(0.3070)
AmzTax*Urban	0.4502*	-0.8822***	0.6277**
	(0.2312)	(0.2614)	(0.3045)
N	745,266	745,266	745,266
Dependent Mean			
Rural	22.972	23.468	42.137
Suburban	23.829	24.288	41.854
Urban	22.282	22.728	39.910
Controls:			
County FEs	X	X	X
Time FEs	X	X	X
Household FEs	X	X	X
Demographics	X	X	X
County Trends	X	X	X

Notes: * p < 0.10, ** p < 0.05, *** p < 0.01. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

alternative equation is specified and estimated:

$$\begin{aligned} y_{i,t} = \alpha + \beta_1 AmzTax_{c,t} \times Rural_i + \beta_2 AmzTax_{c,t} \times Suburban_i + \beta_3 AmzTax_{c,t} \times Urban_i + \\ \beta_7 AmzTax3m_{c,t} \times Rural_i + \beta_8 AmzTax3m_{c,t} \times Suburban_i + \beta_9 AmzTax3m_{c,t} \times Urban_i + \\ \beta_4 X_{i,t} + \beta_5 \delta_c + \beta_6 \lambda_t + \beta_7 \delta_c \times f(time)_t + e_i, \quad (4.3) \end{aligned}$$

where $AmzTax3m$ is a new indicator of treatment within the previous three months. Both $AmzTax3m$ and $AmzTax$ are interacted with the development level (rural, suburban, urban) indicators. All other variables remain the same from previous specifications. In this case β_1, $\beta_2, \beta_3, \beta_7, \beta_8$, and β_9 are the coefficients of interest.

Table 4.21 estimates the modified equation for the core dependent variables. The first column corresponds to the raw monthly Amazon.com expenditures. The second column contains the same tax-adjusted Amazon.com expenditures. The third column displays the impact on Amazon Marketplace purchases. As might be expected the introduction of an additional treatment indicator introduces additional noise and offsetting estimates. In particular, the general estimated effect of the Amazon Tax on tax-adjusted Amazon.com expenditures by rural households now appears to be significantly negative and at a similar magnitude as that of suburban

Table 4.19: Effect of Sales Tax Collection on Total Amazon.com Expenditures - Timing-only

	(1)	(2)	(3)	(4)	(5)
AmzTax	0.3412***				
	(0.0000)				
AmzTax*Rural		0.4348***	0.4819***	0.4538***	0.7983***
		(0.0572)	(0.0443)	(0.0154)	(0.0019)
AmzTax*Suburban		0.4034***	0.4113***	0.3776***	0.0136***
		(0.0264)	(0.0306)	(0.0180)	(0.0004)
AmzTax*Urban		0.0959***	0.0730***	0.1861***	0.5403***
		(0.0015)	(0.0199)	(0.0225)	(0.0015)
N	1,489,266	1,489,266	1,489,266	1,489,266	1,489,266
Dependent Mean	12.595				
Rural		12.967	12.967	12.967	12.967
Suburban		12.790	12.790	12.790	12.790
Urban		11.791	11.791	11.791	11.791
Controls:					
County FEs	X	X	X	X	X
Time FEs	X	X	X	X	X
Demographics			X	X	X
County Trends				X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

households, which has been viewed as the group most able to substitute away from Amazon.com purchases to this point. To further investigate, Equation 4.3 is estimated again, this time running separate regressions by income quartile as was done in 4.3.

Table 4.22 presents the estimates of Equation 4.3 for tax-adjusted expenditures on Amazon.com by consumers that made a purchase on Amazon.com in the year prior to February 2015. The fully-interacted is used, with separate regressions for each income quartile. The estimates for each quartile are contained in the four columns of each table with columns 1-4 containing income quartiles 1-4, respectively.

The time-varying substitution effect for rural households is clearly driven by those with above-median incomes. Rural households with incomes falling in the first or second quartiles show no evidence of long-term substitution away from Amazon.com. This is consistent with the hypothesis that transportation costs (broadly defined) play a role in limiting the ability of rural households to substitute as much as households in closer proximity to a greater variety of brick and mortar retail outlets. Poorer households in rural areas are likely less able to absorb the increased fuel costs and additional wear and tear on any personal vehicles to which they may have access.

Table 4.20: Effect of Sales Tax Collection on Tax-adjusted Amazon.com Expenditures - Timing-only

	(1)	(2)	(3)	(4)	(5)
AmzTax	-0.3805***				
	(0.0000)				
AmzTax*Rural		-0.2896***	-0.2417***	-0.2817***	0.0709***
		(0.0933)	(0.0793)	(0.0536)	(0.0021)
AmzTax*Suburban		-0.3246***	-0.3164***	-0.3544***	-0.7287***
		(0.0304)	(0.0344)	(0.0240)	(0.0005)
AmzTax*Urban		-0.6083***	-0.6322***	-0.4960***	-0.1136***
		(0.0317)	(0.0095)	(0.0054)	(0.0016)
N	1,489,266	1,489,266	1,489,266	1,489,266	1,489,266
Dependent Mean	12.838				
Rural		13.277	13.277	13.277	13.277
Suburban		13.090	13.090	13.090	13.090
Urban		12.071	12.071	12.071	12.071
Controls:					
County FEs	X	X	X	X	X
Time FEs	X	X	X	X	X
Demographics			X	X	X
County Trends				X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

4.5 Discussion

Online retail is currently going through one of the largest upheavals in years. The South Dakota v. Wayfair U.S. Supreme Court ruling has opened the door for states to demand online retailers or marketplace facilitators collect and remit the appropriate state and local sales taxes regardless of their physical location. States are not wasting any time in modifying existing or issuing new guidance, directives, and laws in order to take full advantage of this new potential revenue source. Given this rush to action, a careful examination of the implications of various tax policy options in varying stages of implementation should be of great interest.

The specific focus of this chapter on identifying differential effects along major demographic dimensions following the 2015 implementation of state sales taxes in Ohio and Michigan leads to a story of substitution patterns that vary across geography and wealth levels. Most important is the repeated theme of the rural poor absorbing the entirety of the sales tax applied to Amazon.com purchases in 2015 with little evidence of substitution to competing retailers. In contrast, the suburban households shifted a meaningful quantity of spending away from the entire Amazon platform, significantly increasing expenditures at local electronics, warehouse, and book stores. Higher incomes, more mobility, and close geographic proximity to dense

Table 4.21: Varying Effect of Amazon Tax on Existing Amazon Customer Purchases

	Raw	Tax-adjusted	Marketplace
AmzTax*Rural	0.0341	-1.2644***	-2.7165***
	(0.5202)	(0.5264)	(0.4445)
AmzTax*Suburban	0.0998	-1.2769**	-1.9839***
	(0.6336)	(0.6406)	(0.1946)
AmzTax*Urban	1.3701***	0.0666	0.8124***
	(0.3991)	(0.3913)	(0.2730)
AmzTax3m*Rural	1.0647*	1.0422*	1.8995***
	(0.6077)	(0.6190)	(0.4984)
AmzTax3m*Suburban	-0.6857	-0.7105	0.2629
	(0.4983)	(0.5125)	(0.3336)
AmzTax3m*Urban	-0.7821**	-0.8066***	-0.1570
	(0.3447)	(0.3275)	(0.2941)
N	1,819,812	1,819,812	1,819,812
Dependent Mean			
Rural	12.8587	13.1272	24.4110
Suburban	12.9062	13.1478	23.6465
Urban	11.6928	11.9184	21.8725
Controls:			
County FEs	X	X	X
Time FEs	X	X	X
Demographics	X	X	X
County Trends	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

brick and mortar retail give suburban consumers the widest possible choice set when it comes to retail options. Lower-income urban households also appear to experience effects similar to rural households to a lesser extent. These households are most likely to rely on public transportation, are most affected by congestion whether or not they have personal transportation, and are generally are less able to make large local shopping trips for these reasons in addition to smaller living spaces and storage limitations. These patterns are directly relevant to the decisions states are facing today with respect to their general approach to taxation.

Returning to a question originally posed at the start of the chapter, how should one now think of Wisconsin and Arizona's different approaches to offsetting a sharp increase in sales tax revenue forecasts? Given the body of evidence presented, it appears that Arizona's broad tax credit and preferential treatment for homeowners only further benefits suburban households. These are precisely the households least affected by a uniform increase in the cost of purchasing goods online. Conversely, Wisconsin's strict revenue neutral stance and offsetting credit back

Table 4.22: Varying Effect of Amazon Tax on Tax-adjusted Amazon.com Expenditures by Income

	Income Q1	Income Q2	Income Q3	Income Q4
AmzTax*Rural	1.9595	-0.4324	-0.8816*	-3.5206***
	(1.5780)	(0.7839)	(0.5260)	(0.3465)
AmzTax*Suburban	-0.3589	-0.5591	-2.5159**	-0.9767
	(1.1679)	(0.9385)	(1.2666)	(0.7106)
AmzTax*Urban	3.0696***	-1.8155*	1.0314	-1.7856
	(0.9449)	(1.0834)	(1.0250)	(1.1310)
AmzTax3m*Rural	-1.1683	1.2967*	-0.0976	2.8630***
	(1.4854)	(0.7468)	(0.6633)	(0.6479)
AmzTax3m*Suburban	-1.0494	-0.9138	-0.2158	-0.7176
	(0.7047)	(0.7881)	(0.5853)	(0.9465)
AmzTax3m*Urban	-2.7950***	1.2455	-2.6719***	0.7502
	(0.9750)	(0.8704)	(1.1483)	(0.9027)
N	351,892	424,195	480,855	554,837
Dependent Mean				
Rural	8.7053	10.2489	12.6007	18.7508
Suburban	7.8512	9.9538	12.5772	18.8417
Urban	7.3811	9.5197	12.9475	18.6429
Controls:				
County FEs	X	X	X	X
Time FEs	X	X	X	X
Household Income	X	X	X	X
County Trends	X	X	X	X

Notes: * p < 0.10, ** p < 0.05, *** p < 0.01. Standard errors are clustered at the state level. Coefficients show the effect (in dollars) of the implementation of sales tax collection on Amazon purchases on the dependent variable.

to lower-income households is, if not more utilitarian, at least less disruptive to the existing order in terms of minimizing the likelihood of overall policy winners and losers.

Regarding the newly expanded ability to reshape the state tax mix, states should tread carefully. Any simple and broad-based effort to favor local over online retailers by increasing sales taxes and lowering property taxes, for example, appears to be tantamount to a direct subsidy to suburban households. Wisconsin's policy of directly linking income tax reductions for low-income households to increases in sales tax revenues is an example of a more balanced approach. Additionally, authorities may want to consider targeted subsidies for public transportation, urban redevelopment, and rural retail if they wish to avoid creating policy winners and losers on the consumer side while increasing the relative importance of sales taxes in government budgets.

Bibliography

Alberto Abadie, Susan Athey, Guido Imbens, and Jeffrey Wooldridge. When Should You Adjust Standard Errors for Clustering? *National Bureau of Economic Research*, 2017.

James Alm and Mikhail I. Melnik. Sales taxes and the decision to purchase online. *Public Finance Review*, 33(2):184–212, 2005.

Brian Baugh, Itzhak Ben-David, and Hoonsuk Park. Can Taxes Shape an Industry? Evidence from the Implementation of the Amazon Tax. *Journal of Finance*, 73(4):1819–1855, 2018.

Erik Brynjolfsson, Yu Jeffrey Hu, and Michael D. Smith. Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science*, 49(11):1580–1596, 2003.

Liran Einav, Dan Knoepfle, Jonathan Levin, and Neel Sundaresan. Sales taxes and internet commerce. *American Economic Review*, 104(1):1–26, 2014.

Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. Working Paper 25018, National Bureau of Economic Research, September 2018.

Austan Goolsbee. In a world without borders: The impact of taxes on internet commerce. *Quarterly Journal of Economics*, 115(2):561–576, 2000a.

Austan Goolsbee. Internet commerce, tax sensitivity, and the generation gap. *Tax Policy and the Economy*, pages 45–65, 2000b.

Mark A. Scanlan. Tax sensitivity in electronic commerce. *Fiscal Studies*, 28(4):417–436, 2007.

Chapter 5

Conclusion

This thesis has examined three topics relating to heterogeneous individual and household behaviour in response to shocks that occur in the labour market and in the consumer marketplace. Chapter 2 examined the long-term effects of firings on subsequent labour market outcomes from an ability learning perspective. Chapters 3 and 4 utilized novel high-frequency consumer financial data to study how households navigate temporary liquidity shocks due to income disruption and respond to changes in online sales tax policy, respectively.

In Chapter 2, I developed a search model featuring public learning about worker ability and derived testable implications for fired workers in which the severity of the effect depends upon the ex ante expectations for match success. Using data from the National Longitudinal Survey of Youth, I estimated a reduced-form empirical model of wages and found that firings have a lasting effect on wages after controlling for experience, tenure, and unemployment. Further, this permanent effect is far more severe for university-educated workers relative to high school graduates.

This result is consistent with the model predictions, but is contrary to previous results which found no evidence of learning among university graduates. This new evidence supports the increasing attention in the literature being paid to this type of learning as a complement to search and human capital in explaining employment and wage dynamics throughout the career. Understanding the extent to which employer-initiated separations reflect new information about worker productivity has important implications for the optimal mix of treatments related to human capital development and those intended to increase matching rates out of unemployment.

Chapter 3 performed a detailed analysis of various proposed channels by which households might be able to smooth consumption over the course of a liquidity shock by comparing the behaviour of furloughed U.S. federal government employees to their unaffected counterparts working for fully-funded federal agencies over the course of the 2018-2019 government shutdown, the longest in U.S. history. Using a difference-in-differences approach, the reduction in observed debt payments by furloughed workers during the shutdown is decomposed into an insurance effect supporting consumption and a liquidity effect caused by the interruption in bi-weekly pay. Particular attention is paid to the role played by the level of liquid savings held prior to the shutdown.

I failed to find evidence that households incur additional costs to smooth consumption-related expenditures over a shock of this duration (as captured by debit and credit card transactions), raising questions about how much effort from a policy standpoint should be made to specifically support consumption during temporary liquidity shocks. Rather, evidence of household prioritization of debt payments suggests that resources may be better spent on enabling households to keep existing obligations current.

In Chapter 4, I examined consumer purchasing behaviour in six states before and after the 2015 implementation of state sales taxes in Ohio and Michigan on purchases at Amazon.com, with a specific focus on identifying differential effects along major demographic dimensions. Lower-income households and those living in more remote geographic areas are the most heavily impacted by the relative increase in the cost of online goods. Further, there is significant evidence of interaction between both dimensions. In terms of Amazon.com purchases alone, rural households absorb the entirety of the tax on online purchases, urban households absorb about half, and suburban households do not significantly increase total expenditures at all following the implementation of sales tax collection on Amazon.com purchases.

As states begin collecting sales taxes on all online purchases in the wake of the U.S. Supreme Court ruling of South Dakota v. Wayfair, Inc., it is important to understand the winners and losers of uncompensated changes to sales tax policy. Suburban households generally enjoy relatively high incomes, a lifestyle centered around personal transportation, and close geographic proximity to dense brick and mortar retail. In contrast, both rural and urban households suffer from a relative lack of accessible retail options for varying reasons. On average, urban households are most likely to rely on public transportation, are most affected by congestion, and have smaller living quarters and storage options. Rural households are physically distant from the large retail centers located on the outskirts of major metropolitan areas. Policymakers may want to consider targeted subsidies for public transportation, urban redevelopment, and rural retail if they wish to mitigate the potentially disproportionate impact of changes to tax policies on lower-income households in both rural and inner-city areas.

Appendix A

Proof of Proposition 2.1.2 (Chapter 2)

Proof Proposition 2.1.2 states that given any job α and initial beliefs $\Theta_0 = \{\gamma_L, \gamma_H\}$, for any two population distributions of ability Γ_1 and Γ_2 with support $[\gamma_L, \gamma_H]$ such that Γ_2 is first order stochastically dominant:

$$0 > \Delta_- E_\alpha[\Pi(\Theta_0, \alpha) \mid \Theta_0, \Gamma_1] > \Delta_- E_\alpha[\Pi(\Theta_0, \alpha) \mid \Theta_0, \Gamma_2] \;\; \forall \alpha.$$

The proof involves a simple comparison of the expanded notation taken at the given set of beliefs. Starting with the definition:

$$\begin{aligned} \Delta_- E_\alpha[\Pi(\Theta, \alpha) \mid \Theta] = &- \frac{\Gamma(\theta_H) - \Gamma(\alpha)}{\Gamma(\alpha) - \Gamma(\theta_L)} \int_{\theta_L}^{\alpha} \frac{\Gamma(x) - \Gamma(\theta_L)}{\Gamma(\theta_H) - \Gamma(\theta_L)} Y(x) dF(x) \\ &- E_{x \in (\alpha, \theta_H)}[\Pi(\Theta, \alpha) \mid \Theta] < 0, \end{aligned}$$

expand the notation and note $\Gamma_i(\gamma_L) = 0$ and $\Gamma_i(\gamma_H) = 1$. Then:

$$\Delta_- E_\alpha[\Pi(\Theta_0, \alpha) \mid \Theta_0, \Gamma_i] = -\frac{1 - \Gamma_i(\alpha)}{\Gamma_i(\alpha)} \int_{\gamma_L}^{\alpha} \Gamma_i(x) Y(x) dF(x) - \int_{\alpha}^{\gamma_H} (1 - \Gamma_i(x)) Y(x) dF(x). \quad \text{(A.1)}$$

By first order stochastic dominance, $\Gamma_1(x) \geq \Gamma_2(x) \;\; \forall x$, with strict inequality at some x. Comparing Equation A.1 term by term:

$$-\frac{1 - \Gamma_1(\alpha)}{\Gamma_1(\alpha)} \geq -\frac{1 - \Gamma_2(\alpha)}{\Gamma_2(\alpha)},$$

and:

$$\int_{\gamma_L}^{\alpha} \Gamma_1(x) Y(x) dF(x) \geq \int_{\gamma_L}^{\alpha} \Gamma_2(x) Y(x) dF(x),$$

and finally:

$$-\int_{\alpha}^{\gamma_H} (1 - \Gamma_1(x)) Y(x) dF(x) \geq -\int_{\alpha}^{\gamma_H} (1 - \Gamma_2(x)) Y(x) dF(x).$$

Therefore:

$$\Delta_- E_\alpha[\Pi(\Theta_0, \alpha) \mid \Theta_0, \Gamma_1] \geq \Delta_- E_\alpha[\Pi(\Theta_0, \alpha) \mid \Theta_0, \Gamma_2] \;\; \forall \alpha,$$

and strict inequality is obtained by noting the integration is piecewise over the entire support.

Appendix B

Differences in Time Trends (Chapter 4)

Visual and/or statistical evidence supporting the common trends assumption necessary for the identification of average treatment effects in difference-in-differences paper is typically focused on trends in the data prior to the treatment of the treated. However, due to the identification strategy of this paper the period in which the treatment and control groups are in a common state is the post-period. Therefore the relevant test is for differences in post-period time trends rather than pre-period trends. This section presents charts of average monthly expenditures on Amazon.com along with a statistical test using baseline results from Section 4.3.

Figures B.1, B.2, and B.3 show the average monthly expenditures within each treatment group for rural, suburban, and urban households, respectively. The time axis is inverted so that the area of focus for common trends is on the left. This is done to assist the reader as most difference-in-differences papers examine trends in the pre-period. Due to variation in treatment timing, Ohio and Michigan are presented separately with vertical lines indicating the month in which each state began collecting sales taxes on Amazon.com purchases.

Table B.1 regresses the residuals from the regression of Amazon.com expenditures presented in Column 5 in Table 4.6 in the post-treatment months on an indicator for treatment, time, and the interaction of the two. For the controls, the months after the average treatment date of July 2015 are included. There is a near-zero difference between the states that implemented the tax and those that already had a tax implemented. These results provide evidence of the adequacy of the five control states.

Figure B.1: Rural Household Monthly Expenditures on Amazon.com by Treatment Group

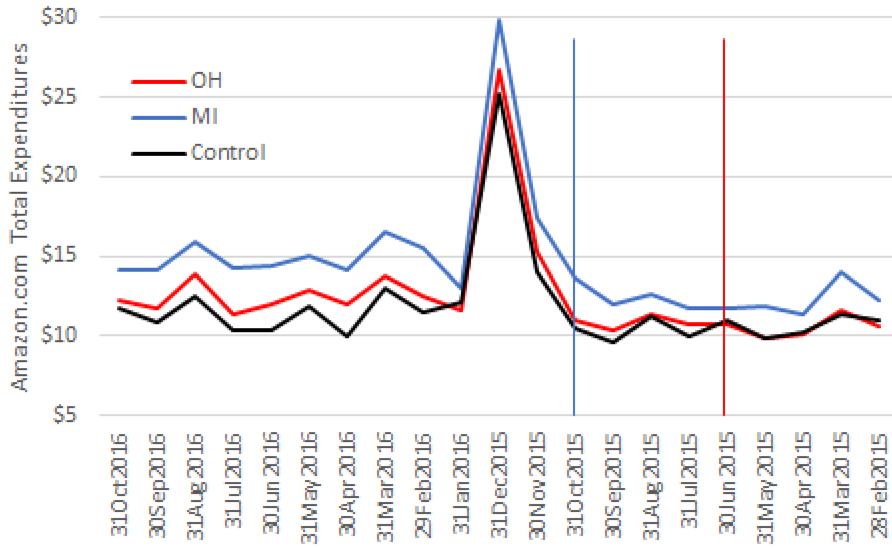

Figure B.2: Suburban Household Monthly Expenditures on Amazon.com by Treatment Group

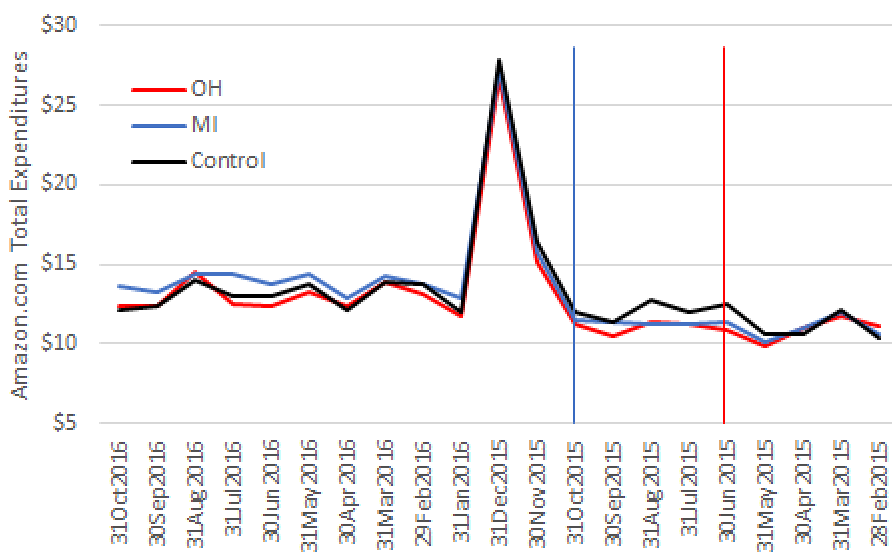

Figure B.3: Urban Household Monthly Expenditures on Amazon.com by Treatment Group

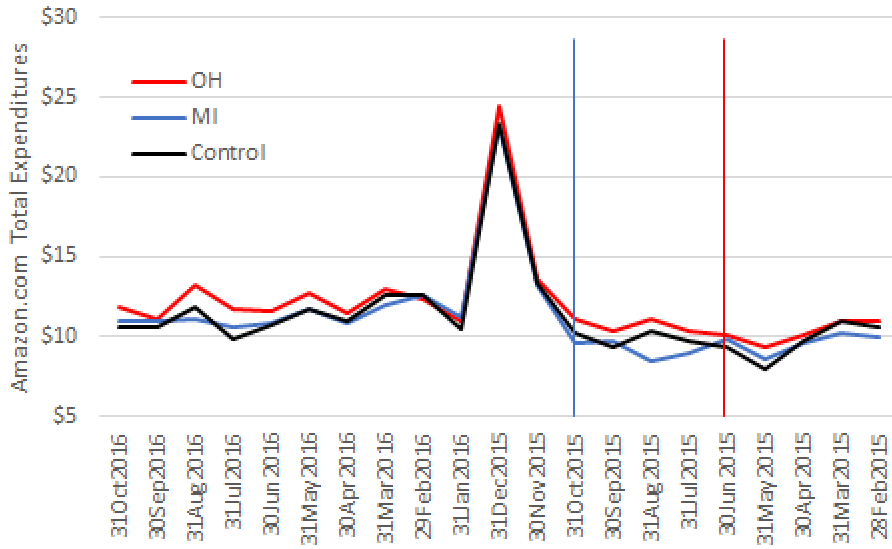

Table B.1: Test for Differences in Post-period Time Trends

	Rural	Suburban	Urban
AmzTax*Time	-0.0009	0.0277	0.0009
	(0.0775)	(0.0443)	(0.0608)
AmzTax	0.0075	-0.2144	-0.0070
	(0.5548)	(0.3233)	(0.4404)
Time	0.0000	-0.0000	0.0000
	(0.0609)	(0.0319)	(0.0440)
N	251,779	651,248	307,407

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The outcome variable is the residual of the fully-interacted baseline specification of the Amazon.com expenditure regression. Control states are assigned a treatment date equal to the average treatment date in the sample (July 2015). The coefficient of interest is AmzTax*Time, which captures any differences in post-period trends of Amazon.com expenditures between the states implementing a tax on Amazon.com purchases and states that had already been collecting the tax from the start of the sample period.

Curriculum Vitae

Name:	Brian Held
Post-Secondary Education and Degrees:	The University of Iowa Iowa City, IA 2002 - 2007 B.A.
	The University of Western Ontario London, ON 2009 - 2021 Ph.D.
Related Work Experience:	Data Scientist Huntington National Bank 2016 -
	Data Scientist JPMorgan Chase & Co. 2014 - 2016
	Teaching Asst., Research Asst. (2012-2014), Instructor (2012-2013) The University of Western Ontario 2009 - 2014
	Math Instructor Edgewood-Colesburg (IA) Jr./Sr. High School 2007 - 2009

Conference Presentations:

Held, B., 2013. Evidence of an Ability Learning Process Among University Graduates. 47^{th} Annual Conference of the Canadian Economics Association, Montreal, QC.