

Coming in at a Trickle: The Optimal Frequency of Public Benefit Payments *

Cameron LaPoint[†]

Shogo Sakabe[‡]

July 2, 2020

[Latest version here](#)

Abstract

The question of how governments should choose the frequency of payments has received little attention in the literature on the optimal design of public benefits programs. We propose a simple model in which the government chooses the length of the interval between payments, subject to a tradeoff between the administrative cost of providing more frequent benefits and the welfare gain from reducing deviations from full consumption smoothing. In our empirical application, we examine consumer and retailer responses to bimonthly payments from the Japanese National Pension System. We exploit variation in the duration of payment cycles using a unique retail dataset that links consumers to their purchase history. Our difference-in-differences style approach shows a clear spike in spending on payment dates for customers who are of retirement age relative to those who are not. While within-store average prices increase by 1.6% on payday, this effect is almost entirely due to consumers substituting towards higher quality goods rather than a retailer response. We use these reduced form estimates to parameterize the model and conclude that the optimal frequency of Japanese public pension payments is less than one month, implying the government could improve welfare by increasing payment frequency.

Keywords: optimal payment frequency, pension payments, consumption smoothing, incidence, retail panel, store-level prices

JEL classifications: D03, D91, E21, H21, H55

*We would like to thank Michael Best, Wojciech Kopczuk, Michaela Pagel, Jón Steinsson, Takashi Unayama, Tsutomu Watanabe, David Weinstein, and participants in the Applied Micro Methods Colloquium at Columbia for helpful advice and comments. We thank Louisa Liu for excellent research assistance. We are grateful to Magee Co., Ltd. for providing the retail scanner data used in this project, and to the Research Institute of Economy, Trade and Industry (RIETI) for providing access to the JSTAR data. We also thank the Center on Japanese Economy and Business (CJEB) at Columbia Business School for generous financial support.

[†]Yale School of Management. 165 Whitney Avenue, New Haven, CT 06511. Email: cameron.lapoint@yale.edu; Web: <http://cameronlapoint.com>

[‡]Department of Economics, Columbia University, 1022 International Affairs Building, 420 W 118th Street, New York, NY 10027. Email: s.sakabe@columbia.edu

1 INTRODUCTION

Most advanced economies offer public benefits programs to provide regular transfer payments to their citizens. One important question that has been neglected in the literature on the optimal design of these programs is how the government should choose the frequency of disbursement to eligible households. Several studies have found evidence that infrequently distributing benefits to people who live “payment to payment” may have adverse consequences. For instance, [Dobkin & Puller \(2007\)](#) document that recipients of Supplemental Security Income (SSI) drive spikes in drug-related hospital admissions at the beginning of the month. Similarly, [Mastrobuoni & Weinberg \(2009\)](#) find that Social Security benefit recipients without savings not only consume fewer calories the week before they receive checks relative to the week afterwards, but also that nonsavers’ diets contain more fat and cholesterol towards the end of the pay period.

In light of this evidence, we ask whether governments can improve household welfare by distributing transfer payments more frequently. To answer this question we propose a simple model that defines the optimal frequency of pension payments as a function of aggregate statistics about the benefits system and individual preferences. In our model, the government optimally chooses the length of the interval between benefit payments. This decision is subject to a tradeoff between the administrative cost of providing more frequent benefits and the potential welfare gain from reducing deviations from full consumption-smoothing behavior.

In our empirical application, we apply our model to the Japanese National Pension System (JPS), which distributes bimonthly annuity payments that are a function of average monthly earnings while employed. Upon reaching retirement eligibility, contributors to the system can begin receiving pension payments every two months on the 15th of each disbursement month (February, April, June, etc.). However, if the scheduled delivery date falls on a Saturday, Sunday, or public holiday, payments are instead sent on the first previous non-holiday weekday. This timing rule, combined with annual variation in calendar weekdays, induces variation in the length of periods in between payments that is unrelated to pensioners’ spending decisions.

We exploit this variation in the duration of payment cycles using a unique retail point-of-sale dataset from a Japanese marketing firm. Our sample includes the price and quantity of each good purchased in over 500 grocery stores across Japan over a three-year period between 2011 and 2014. Transactions in our data are tied to loyalty point cards for which we observe a unique member ID and the birth month/year of the registered member. We are thus able to track individuals over time while observing exact prices paid for goods within each transaction at daily frequency.

Leveraging the high-frequency nature of the retail panel, we find that regular shoppers eligible for pension payments increase overall grocery expenditures by 10% within two days of the scheduled delivery date. We also estimate a “duration elasticity,” or the extent to which expenditure responses are more pronounced following longer intervals between payments in our sample time period. We use this duration elasticity and our estimates of the administrative costs of the JPS to calibrate

our optimal payment frequency model. To estimate the shape of the administrative cost function, we exploit a reform to the Japanese pension system in 1988 that reduced the interval between payments from three to two months. The results from this natural experiment suggest the marginal cost to the government of increasing the payment frequency is negligible. We find annual costs rose by 4.3%, or by 0.14% per day the pay cycle decreased in length, for municipalities which contained JPS branch offices that bore the bulk of the increased administrative burden from the reform relative to municipalities without a branch office.

At the same time, we provide evidence that shoppers in our data can be described by a combination of quasi-hyperbolic discounters and “payday liquid” consumers who exhibit a spike in consumption on payday but maintain an otherwise smooth consumption profile. For each type of agent, the implied marginal welfare loss is sufficiently non-trivial to justify an increase in the frequency of pension payments. We also consider alternative motivations for payday expenditures, including near-rationality and liquidity constraints. Following [Kueng \(2015, 2018\)](#), we use total expenditures over the pay cycle as a proxy for permanent income and document a slightly negative relationship between permanent income and payday expenditures, but only among shoppers with below-median permanent income. This result indicates that some shoppers may face short-term liquidity constraints, but we caution that such patterns would also be consistent with a higher incidence of self-control problems among low-income shoppers ([Parker 2017](#)).

In our baseline optimal frequency model we assume that consumers pay prices that do not vary with pay cycle length. To test the validity of this assumption, we examine whether retailers capture some of the incidence of pension payments by price discriminating against elderly customers on or around scheduled delivery dates. Accounting for this possibility is crucial to assessing the welfare change associated with changes to the interval between payments, as changing the pay cycle length could alter the menu costs retailers incur by changing their prices. While our model can easily incorporate retailer pricing responses by characterizing the price of consumption as a function of interval length, we find limited empirical support for the price discrimination hypothesis.

Using the Japanese old-age pension system as our research setting helps avoid issues that have plagued the large literature on measuring the consumer response to anticipated payments. First, small living spaces in Japan make it difficult to buy in bulk or store groceries over long periods, so the transactions we observe more closely approximate “instantaneous consumption” rather than a savings mechanism. Second, universal health insurance coverage reduces the need of retirees to save their pension income for uncertain medical expenses. Third, as [Stephens & Unayama \(2011\)](#) document in 1986-1994 household survey data, over 80% of income for public pension recipients is due to these benefits. Since pensioners have to subsist for long intervals (two months) between which they receive little to no income, our study provides a particularly stark case where increasing the frequency of benefits might improve welfare.

[Shapiro \(2005\)](#) first introduced the idea that a government could improve welfare by increasing the frequency of benefit payments. This idea is echoed by [Parsons & Van Wesep \(2013\)](#), who

study how firms offer workers contracts with different pay cycle duration. Shapiro finds that caloric intake declines among SNAP recipients over the month in between payments and concludes that this consumption path is consistent with quasi-hyperbolic discounting. Using the same dataset, [Mastrobuoni & Weinberg \(2009\)](#) analyze the consumption path of SSI recipients over the pay cycle and draw similar conclusions in favor of hyperbolic discounting. [Dobkin & Puller \(2007\)](#) study the link between hospital admissions and SSI receipt and argue that recipients of SSI drive spikes in drug-related hospital admissions at the beginning of the month, while [Foley \(2011\)](#) shows that crime incidence increases towards the end of monthly welfare cycles. Moreover, the link between adverse health outcomes and non-smooth consumption paths suggests calculations of the utility loss from permanent income hypothesis (PIH) deviations, such as those conducted in [Cochrane \(1989\)](#) and [Browning & Crossley \(2001\)](#), may be severely underestimated.

More generally, this project relates to the vast literature testing the permanent income hypothesis in the context of the consumption response to *regular* income sources (i.e. recurring payments, such as pension payments, that are delivered more than once per year).¹ The most closely related study to our paper is [Stephens & Unayama \(2011\)](#) who look at the Japanese pension system during an earlier period in the 1980s when pension payments were distributed quarterly. Consistent with our results, they find suggestive evidence that consumption growth in the month of check receipt is lower after the reform which shifted from quarterly to bimonthly payments.

A more recent strand of this literature uses high-frequency data from financial planning applications to document consumer responses to various income sources. [Gelman et al. \(2014\)](#) find that total spending rises 70% above the daily average on the day a regular paycheck arrives. [Olafsson & Pagel \(2018\)](#) find similar results in Iceland and argue that consumers act as if they have a license to spend at the beginning of a new pay cycle. [Baker \(2018\)](#) shows that liquidity constraints can explain heterogeneity in the responsiveness of consumption to income shocks. Using a similar research design to ours, [Baugh & Wang \(2018\)](#) exploit within-household variation in the length of Social Security pay cycles to show that households are more likely to experience financial shortfalls during longer pay cycles. While they argue that inattention to changes in the length of the pay cycle can explain their results, we provide evidence that consumers in our setting are highly attentive to even small changes in payment arrival dates.

Our analysis of expenditure patterns of the elderly also relates to the literature on the retirement consumption puzzle. [Bernheim et al. \(2001\)](#) show that total expenditures drop at retirement for consumers in all but the highest income and wealth quartiles. [Aguiar & Hurst \(2005\)](#) contend that this drop does not imply sub-optimal saving for retirement, as at-home food production offsets the decline in food expenditures. [Battistin et al. \(2009\)](#) exploit a discontinuity in pension

¹The vast majority of these studies use quarterly or monthly panel data and examine payments that occur at the cutoff separating two discrete time observations. For example, with monthly panel data, estimates of consumption growth between June and July due to receipt of a bonus at the end of June are attenuated if households spend a portion of the bonus in June and the other portion in July. We overcome this issue by using high-frequency scanner data that allows us to observe expenditure responses at a precise number of days between scheduled payment dates.

eligibility in Italy to estimate a 9.8% drop in consumption due to male retirement. However, they argue this drop becomes statistically insignificant once they account for adult children moving out when parents retire. [Stephens & Unayama \(2012\)](#) find little evidence of a drop in consumption among Japanese retirees, but show that this is primarily due to working households receiving large lump-sum retirement bonuses from their employer. Although we cannot directly test for liquidity constraints in our data, we document similar consumption patterns for shoppers who visit stores at different average shopping frequencies. This suggests that part of the response of expenditures to payday is due to factors other than short-term liquidity.

Finally, whether retailers capture some of the incidence of public benefit payments by engaging in price discrimination against eligible consumers is an open question. This question is of direct relevance to the welfare calculations we conduct here, since increasing the frequency of payments may also hinder retailers' ability to engage in price discrimination if menu costs are sufficiently large. [Warner & Barsky \(1995\)](#) collect daily price data for retail stores between November and February and document that markdowns tend to occur during more intensive shopping periods. This finding is echoed by [MacDonald \(2000\)](#) and [Chevalier et al. \(2003\)](#), who argue that average prices fall during seasonal demand peaks due to coincident declines in retail margins. [Nevo & Hatzitaskos \(2006\)](#) instead offer consumer substitution towards an increased supply of cheaper products as an explanation. In contrast, [Hastings & Washington \(2010\)](#) show that the price index for the basket of food goods consumed by SNAP-recipient households in Nevada falls by 3% over the month. [Goldin et al. \(2019\)](#) extend the analysis to 48 states and a large number of stores and find no evidence of a retailer pricing response on SNAP delivery dates.

We find limited evidence that retailers raise prices to capture the incidence of pension payments. Average prices paid within a store increase by 1.6% on payday, but the number of unique goods purchased increases by 6%. Using a simple price index that isolates a change in average prices due to consumer substitution from the retailer's pricing response, we demonstrate that the observed response of store-level prices is almost entirely due to consumers substituting towards an expanded expenditure basket which includes higher quality goods, rather than retailers raising prices to capture the incidence of pension payments. This finding is consistent with evidence from [Hastings & Shapiro \(2018\)](#) on the expenditure patterns of SNAP recipients and the literature on mental accounting ([Thaler 1999](#); [Farhi & Gabaix 2015](#)), which argues individuals earmark income sources for specific spending categories. Overall, our results imply that altering the pay cycle length is unlikely to affect retailers' ability to capture the incidence of benefit payments.

The remainder of the paper is organized as follows: [Section 2](#) presents a simple optimal payment frequency model. [Section 3](#) provides background on the Japanese pension system and describes the data we use. [Section 4](#) presents our main empirical results on the expenditure response to pension payment receipt. [Section 5](#) describes our methods for constructing price indices and shows results for the response of real expenditures to payment receipt. [Section 6](#) discusses how we calibrate the model to determine the optimal payment frequency. [Section 7](#) concludes.

2 A SIMPLE OPTIMAL PAYMENT FREQUENCY MODEL

In this section, we outline the modeling framework for the government’s choice of the optimal interval between regularly occurring public benefit payments. While we present the model in the context of a public pension system, the model can be applied to a variety of contexts where the government disburses payments at regular, anticipated, intervals.

2.1 BASIC FRAMEWORK

Consider periods of time t that occur within the interval $[0, T]$, where each unit of time is measured in days. A fraction p of people in the economy are pensioners who receive a flat pension benefit every T days equal to $b(T) = \bar{B} \cdot T$. The other $1 - p$ fraction of people in the economy are workers, who instead of receiving the pension benefit, earn an arbitrary wage $w(t)$ and pay a lump-sum tax of $\tau(b)$ that is used to finance the pension system.²

In addition to the direct cost of $p \cdot b(T)$ of delivering benefits to retirees, the government faces an administrative cost function $\mu(T)$.³ While the shape of $\mu(T)$ will ultimately depend on the particular application, for now we assume the cost function is both strictly increasing in the frequency of payments and weakly convex, so $\mu'(T) < 0$, $\mu''(T) \geq 0$. We discuss our strategy for estimating the administrative cost function in [Section 6](#).

At each point in time, assume the government must follow a balanced budget rule, so that for each $t \in [0, T]$ the following must hold:

$$(1 - p) \cdot \tau(b) = p \cdot b(T) + \mu(T) \implies \tau(b) = \frac{p \cdot \bar{B} \cdot T + \mu(T)}{1 - p} \quad (1)$$

Working households have instantaneous utility of consumption $u(C(t))$. These non-pensioners face a standard consumption-saving problem of

$$\max_{\{C(t)\}} \int_0^T u(C(t)) dt \text{ s.t. } C(t) = S(t) + w(t) - \frac{\tau(b)}{T} \quad (2)$$

where for simplicity we normalize $r = 0$. In the absence of liquidity constraints that would place restrictions on the asset position $S(t)$, the non-pensioner’s optimal consumption path is completely flat over the interval: $C(t) = C^*, \forall t \in [0, T]$.

Let $u^r(C(t))$ denote the instantaneous utility of retirees. Pensioners are given $b(T)$ to spend over the payment cycle. The optimal consumption path of retirees is therefore also constant on the

²The portion of JPS benefits that are not contingent on employment are financed through monthly lump-sum payments. The value of the monthly payment in 2018 was 16,340 JPY (\approx \$143).

³These administrative costs to implementing the pension system may include costs associated with: authorizing benefits, delivering benefits (either electronically or through post), redeeming and reconciling benefits, investigating and prosecuting fraud, managing pension funds.

interval $[0, T]$. However, we suppose that, following [Shapiro \(2005\)](#), the pensioner instead chooses a consumption path given by

$$C(t) = \exp(\theta - f(t)) \quad (3)$$

where $f(t)$ is a potentially non-monotonic function representing the deviation of the chosen consumption path at each time t from a constant value. Because the source of this deviation will again depend on the particular application, in the general setup we remain agnostic on the underlying behavioral phenomenon driving the household away from the consumption smoothing benchmark. The only restriction we impose on $f(t)$ here is that $f(0) = 0$. Since the only income pensioners receive within the payment cycle is $b(T)$, the budget condition pins down the value of the constant θ :

$$\int_0^T \exp(\theta - f(t)) dt = b(T) \implies \exp(\theta) = \frac{b(T) \cdot f'(0)f'(T)}{f'(T) - f'(0) \cdot \exp(-f(T))} \quad (4)$$

To determine the potential welfare gain that the government might achieve by changing payment frequency, we can ask the share $(1 - \lambda)$ of its benefit the pensioner household would be willing to give up to achieve the optimal constant consumption profile. This is the fraction λ that solves the following equation:

$$\int_0^T u^r(\exp(\theta - f(t))) dt = \int_0^T u^r(\lambda \bar{B}) dt \quad (5)$$

[Shapiro \(2005\)](#) solves this expression for λ in the special case where $f(t) = \nu \cdot t$ and $u^r(C(t)) = \frac{C(t)^{1-\rho}}{1-\rho}$, so that the welfare loss is independent of the benefit level. Our aim here is more general: we wish to find the optimal frequency T that minimizes the welfare loss, given the total costs of funding the pension system.

Suppose there exists a constant γ that represents the cost $\tau(b)$ of a dollar spent by the government in the same units as total utility of the non-pensioners over the pay cycle. Such a constant corresponds to the marginal cost of funds (MCF), and under a lump-sum tax imposed only on the non-pensioners, $\gamma = 1$.⁴ Given a constant MCF, the optimal frequency problem a utilitarian government faces is given by

$$\min_T \left\{ -p \cdot \lambda(T) + \gamma \cdot (p \cdot b(T) + \mu(T)) \right\} \quad (6)$$

where $\lambda(T)$ is the solution to (5). For any strictly concave $u^r(\cdot)$, the welfare loss $1 - \lambda$ will be strictly convex in T . Therefore, in such a case the FOC of this problem is necessary and sufficient

⁴More generally, for a distortionary tax τ , the MCF can be computed as $\gamma = -\frac{\partial \mathcal{U}^* / \partial \tau}{\partial R^* / \partial \tau}$ where \mathcal{U}^* is the total optimized level of utility for the non-pensioners, and R^* is revenue collected at the optimum.

for a solution:

$$\frac{p \cdot \lambda'(T^*)}{\gamma} = \mu'(T^*) + p \cdot \bar{B} \quad (7)$$

We have now obtained T^* as an implicit function of observables. These include the daily average benefit amount among claimants over the pay period \bar{B} and the fraction p of the population who are claimants.

Next we turn to two special cases of the model, which are motivated by observations in the literature on expenditure responses to income receipt. One case features benefit recipients who are quasi-hyperbolic discounters, and another assumes recipients behave as if they have a license to spend on payday but otherwise maintain a smooth consumption profile over the pay cycle.

2.2 QUASI-HYPERBOLIC CONSUMERS

We describe a special case of our optimal frequency model where pensioners are quasi-hyperbolic (hereafter, QH) discounters, and hence the deviation function $f(t)$ is linear in time. The analysis here builds on the special cases examined in [Shapiro \(2005\)](#) and [Mastrobuoni & Weinberg \(2009\)](#), who posit that consumer preferences take the form:

$$u(c_0) + \beta \sum_{t=1}^T \delta^t u(c_t) \quad (8)$$

where β is the QH discount factor, and δ is the standard daily exponential discount factor. Both papers focus on the case where $\delta = 1$, $\beta < 1$, and the felicity function $u(\cdot)$ is isoelastic.⁵ Under these conditions, one can show that the decrease in log consumption over time is:

$$\frac{\partial \log(c_t)}{\partial t} = \frac{1}{\rho} \cdot \log \beta - \frac{1}{T-t+1} + \frac{1}{T-t+\beta^{-1/\rho}} < 0 \quad (9)$$

where ρ is the inverse intertemporal elasticity of substitution (IES). Under these assumptions for preferences, for values of β close to one, the time path for consumption given by (9) is approximately linear over the pay cycle.⁶ Thus following [Shapiro \(2005\)](#), in the setup of our optimal frequency model we capture QH discounting behavior for the case where $f(t) = \nu \cdot t$ and ν is the constant daily rate of decline in consumption over the pay cycle. The linear rate ν corresponds to some combination of β and ρ that rationalizes an observed total percentage decline in consumption over the pay cycle. For instance, [Shapiro \(2005\)](#) finds that consumption declines by 0.4% daily ($\nu = 0.004$) for the average household participating in Maryland's food stamps program.

⁵The case of $\delta = 1$ is a reasonable approximation to the daily exponential discount factor for analyzing consumption decisions over a relatively short interval between payments.

⁶Exponential discounters with $(\beta, \delta) = (1, 1)$ have a flat consumption profile, and for $\beta = 1$, $\delta < 1$ the exponential discounters have a linear decline in consumption over the cycle with slope $\log(\delta)/\rho$.

For log utility ($\rho = 1$), the compensating variation term from (5) is

$$\lambda(T) = \frac{1}{B} \cdot \exp\left(\theta - \nu \cdot T/2\right) \quad (10)$$

and the FOC to the government’s problem is given by

$$-\frac{p \cdot \nu}{2B} \cdot c_0 \cdot \exp\left(-\nu \cdot T/2\right) = \mu'(T) + p \cdot \bar{B} \quad (11)$$

We provide an expression for the welfare loss with $\rho \neq 1$ in [Appendix A](#).

In [Section 6](#) we numerically solve for the welfare loss and the optimal payment frequency from a calibrated version of our model with pensioners as QH discounters. For empirically valid combinations of ν and ρ , the welfare loss from non-smoothing is relatively small in the Japanese pension system context. For example, if the average daily consumption decline over the pay cycle is 0.4%, then for $\rho = 10$, the welfare loss with a bimonthly pension system ($T = 60$) is 2.3% of consumption; in the log utility case the loss is only 0.24% of consumption.⁷ [Figure 1](#) summarizes how the welfare loss varies with the interval T for different assumed values of the inverse IES.

We note two features about the welfare loss under QH discounting that the figure illustrates. First, the welfare loss is increasing in the government’s choice of the interval length T . Intuitively, for longer pay cycles the integral between the optimal smooth path and the path under QH discounting will be greater, as log deviations from the smoothed level of consumption grow linearly over time. Second, the welfare loss is increasing in the inverse IES. A higher value for ρ means consumption is less substitutable between periods, so an individual is willing to pay more *ex ante* to get closer to the smooth consumption path.

2.3 PAYDAY LIQUID CONSUMERS

We now examine a version of our model where we assume pensioners exhibit “payday liquidity.” Several recent studies exploiting high-frequency data from financial planning applications have documented that consumers spend considerably more around the day they receive a regular paycheck or benefit payment from the government. For instance, [Gelman et al. \(2014\)](#) analyze 60 million transactions in the U.S. and find that, on average, total spending rises 70% above the daily average on the day a regular payment arrives. [Olafsson & Pagel \(2018\)](#) report similar results with user data from a financial planning app in Iceland, with the poorest tercile of households in their sample spending 70% more on a payday, and the richest tercile spending 40% more on paydays. Notably, both papers show that the spike in consumption on paydays cannot be fully explained by measures of liquidity. Households instead act as if they have payday liquidity, or a license to spend at the beginning of a new pay cycle.

⁷Most non-experimental estimates for the inverse IES fall in the range $\rho \in [1, 3]$. However, [Best et al. \(2018\)](#) use notches in mortgage interest rate schedules in the UK to argue that $\rho = 10$ is plausible.

In our framework, we consider a case where pensioners heuristically spend more at $t = 0$ when they receive their payment, but their consumption path is otherwise smooth over the rest of the pay cycle. Due to the discontinuous nature of the consumption path in this setting, in this subsection we proceed in discrete rather than continuous time for ease of exposition.⁸ We assume pensioners consume over the time interval $[0, T - 1]$ as follows:

$$C_t = \begin{cases} (1 + x) \cdot \bar{c} & \text{if } t = 0 \\ \bar{c} & \text{if } t \in [1, T - 1] \end{cases} \quad (12)$$

As before, pensioners receive $\bar{B} \cdot T$ to spend over the time period $[0, T - 1]$, so that the budget constraint pins down the value of \bar{c} .

$$\begin{aligned} \sum_{t=0}^{T-1} C_t &= (1 + x) \cdot \bar{c} + \sum_{t=1}^{T-1} \bar{c} = (T + x) \cdot \bar{c} = \bar{B} \cdot T \\ \implies \bar{c} &= \begin{cases} \frac{\bar{B} \cdot T}{T + x} & \text{if } T > 1 \\ \bar{B} & \text{if } T = 1 \end{cases} \end{aligned} \quad (13)$$

\bar{c} represents the constant value of consumption in all days within the pay cycle besides the payday where consumption spikes. The value $x = c_0/\bar{c} - 1$ represents the magnitude by which consumption spikes on payday.⁹

Many utility functions can produce the consumption path in (13) as the solution to utility maximization. For instance, a log felicity function with a $(1 + x)$ weight on utility in $t = 0$ will generate a spike on payday of $x\%$ relative to \bar{c} .¹⁰ However, if C_t were consistent with utility maximization, then there would be no welfare loss from non-smoothing, and our work here would be done. We thus consider a more interesting case of behavioral consumers who heuristically spend more on payday, even though this is inconsistent with utility maximization.

The discrete time analog of the compensating variation equation in (5) is:

$$\sum_{t=0}^{T-1} u(C_t) = T \cdot u(\lambda \bar{B}) \quad (14)$$

⁸It is possible to directly map the results in discrete to continuous time using the concept of Cauchy principal values to evaluate integrals of functions that are discontinuous over the time interval.

⁹While in practice consumption is not literally flat over time intervals with no paydays, it is roughly flat net of a full set of calendar effects (day-of-the-week, week-of-the-month, month-year). We show results to this effect later in the paper. In that sense, the consumption path we model here is that of average daily consumption, or an ergodic process of consumption.

¹⁰More generally, an isoelastic felicity function with a $(1 + x)$ weight on utility in $t = 0$ will generate a spike on payday of $[(1 + x)^{1/\gamma} - 1]$ % relative to \bar{c} .

To facilitate comparison to other cases of our model, we continue to assume an isoelastic felicity function with inverse IES ρ . The expression for the welfare loss is then:

$$1 - \lambda(T) = \begin{cases} 1 - \frac{\bar{c}}{B} \cdot (1+x)^{1/T} & \text{if } \rho = 1 \\ 1 - \frac{\bar{c}}{B \cdot T^{1/(1-\rho)}} \left[(1+x)^{1-\rho} + (T-1) \right]^{\frac{1}{1-\rho}} & \text{if } \rho \neq 1 \end{cases} \quad (15)$$

Figure 2 plots the welfare loss formula in (15) for the isoelastic payday liquid consumer as function of T for different values of ρ . In this figure we assume a conservative estimate for the spike in payday consumption ($x = 0.1$) that we obtain in Section 4.2 for perishable food consumption. We normalize the welfare cost to be 0 when there is only one period. The welfare loss with payday liquid pensioners is declining in the government's choice of T . Because the welfare loss is entirely concentrated in the initial drop in consumption between periods $t = 0$ and $t = 1$, increasing the length of the pay cycle merely subdivides this loss over a larger total amount of consumption. Hence, for any non-negative cost function $\mu(T)$, the optimal payment frequency is a corner solution. Even if the government incurs zero costs to distributing payments more frequently, it could minimize the welfare loss by choosing the largest T^* such that pensioners can subsist ($\bar{c} > 0$).

So far we have assumed that the spike in consumption on payday is independent of the interval between payments. If instead $x'(T) > 0$ this would capture the notion that consumers splurge more on a payday when more time has passed since they last received a payment. In other words, there may be more “pent-up demand” when pay cycles are longer. One can show that in this scenario decreasing T can improve welfare if the loss from the increase in the spike magnitude associated with an increase in T exceeds the welfare gain from subdivision as T increases. With log utility that condition reduces to the following:

$$\underbrace{\left(1 + x(T)\right) \cdot \log\left(1 + x(T)\right)}_{\text{loss from spike magnitude as } T \uparrow} > \underbrace{Tx'(T)}_{\text{gain from subdivision as } T \uparrow} \quad (16)$$

In our empirical setting, we locally estimate the slope $x'(T)$ by exploiting exogenous calendar variation in pension delivery dates and find that $x'(T) = 0.0023$, so the spike on payday increases by 0.23 percentage points for each additional day in the pay cycle. For this value of $x'(T)$ the condition in equation (16) holds for any $\rho > 0$. Figure 3 shows the welfare loss as a function of T under this calibration. In contrast to the case where the spike was assumed to be a constant, the welfare loss is now increasing and concave in interval length.

We note that the welfare loss for payday liquid consumers is much smaller than that for QH discounters when $T = 60$ as in the current Japanese Pension System. This is true regardless of whether the spike is a function of T or a constant. For example, at $T = 60$, for $\rho = 10$ and $x = 0.1$ the welfare loss for payday liquid consumers is 0.06%, compared to 2.29% of consumption for QH discounters. When we allow the spike on payday to depend on the interval length, the loss increases to 0.1% of consumption. The small magnitude of the loss is unsurprising given that the welfare

loss for payday liquid consumers is entirely concentrated in the initial drop in consumption. Yet, what matters for the optimal frequency is the *marginal* welfare loss. We show in [Section 6](#) that the marginal welfare loss is sufficiently non-trivial in each of these cases to justify an increase in the frequency of payments in the JPS.

To summarize this section, we have characterized two special cases of our optimal payment frequency model that feature prominently in the literature: QH discounting and payday liquidity. We present detailed derivations in [Appendix A](#) and assess the empirical relevance of these two special cases and alternative theories in [Section 4.4](#).

3 DATA AND BACKGROUND

In this section we provide background on the Japanese National Pension System and describe the retail panel data and municipal-level data we use to construct administrative cost measures.

3.1 DATA

Retail panel data. We use retail scanner data provided by Magee Co., Ltd. to measure how expenditures and retail prices respond to shoppers receiving payments from the Japanese National Pension System (JPS).¹¹ The data record sales receipts from 1,120 grocery stores of 19 grocery chains across Japan from April 2011 to October 2014. Each sales receipt consists of the purchase date and hour, a consumer-identifier code, a store-identifier code, and barcode-level prices and quantities of goods purchased. A unique characteristic of this dataset is that it includes consumer identifiers assigned at the time consumers sign up for a shopper loyalty program. This member ID has information about each consumer’s birth year/month and gender. If a consumer brings her membership card when she shops, her identifier code is recorded on the sales receipt, and it enables us to keep track of her purchasing history over time.

Although the data include all transaction records for each store during the sample period, we focus on transactions involving consumers who own and regularly used their membership cards, as this allows us to assign pension eligibility status. We define regular shoppers as those who use their point cards at least two times per month (four times per pay cycle) starting from the first month they appear in the data. This leaves us with approximately 1,000,000 shoppers out of a total of roughly 4,000,000 unique shoppers with point cards. [Figure 4](#) shows the number of unique point card shoppers who make a purchase at least k times per month. The number of shoppers stabilizes at around 500,000 for shoppers who go to the store on at least a weekly basis ($k \geq 4$). We obtain qualitatively similar results when we deviate from our baseline sample restriction of $k = 2$.

¹¹To our knowledge, the only other paper that has used these data is [Shoji \(2018\)](#), who investigated consumer stockpiling behavior in advance of a consumption tax hike in 2014.

We make two additional sample restrictions at the store level to obtain our final sample of 511 stores. First, we drop stores from the sample that exit the panel prior to the last month of our sample. These stores might exit either because the store closes for business or because they choose to stop providing data to the marketing firm. Second, we restrict to stores that offer loyalty point card programs. Additionally, our results are unaffected when we restrict to a more balanced panel of stores which record transactions during all non-holiday dates in our sample.

We make use of the detailed system of classification codes provided by Magee Co. Each barcode in our retail panel is matched to a set of 1-digit to 4-digit categories; for instance a tomato would fall into the “agricultural products,” “vegetables” and “fruits and vegetables” codes. To facilitate comparison of our analysis to other papers in the literature and classify products by durability, we aggregate these codes into 13 subcategories, plus a category called “raw foods” which contains all fresh, non-packaged food items in our sample. We describe the contents of these categories in more detail in [Appendix B](#).

[Table 1](#) provides summary statistics on average monthly expenditures all goods expenditures and raw foods expenditures, and shopping patterns of the consumers who regularly visit stores in our sample. Overall, these tabulations suggest that our sample restrictions generate a set of shoppers who make frequent store visits and conduct most of their grocery shopping at stores we can observe in our data. Roughly 28% of the shoppers in our sample have reached the normal retirement age of 65 as of the beginning of the sample, and the average consumer makes 9 trips to a store each month, with an average of 3 days between each visit. Average monthly expenditures among shoppers in our sample are 23,100 yen per month (approximately 230 USD).

Japanese Study of Aging and Retirement (JSTAR) data. Since our primary dataset is at the retail level, we can only observe the small number of characteristics that shoppers report when they sign up for a loyalty point card. To gain a more detailed snapshot of pensioner demographics we examine survey responses to the Japanese Study of Aging and Retirement (JSTAR) conducted by the Research Institute of Economy, Trade and Industry (RIETI). The JSTAR is a biannual panel survey of individuals aged 50 or older that is modeled after the Health and Retirement Study in the U.S. We pool responses from the 2007, 2009, and 2011 waves of the survey, each of which includes data from interviews of approximately 4,000 individuals across ten municipalities in Japan.

We use the JSTAR data to help interpret our empirical results along three dimensions. First, we determine the distribution of ages at which pensioners begin claiming their benefits. Since we identify payments based on a shopper’s age, accounting for the potential endogeneity of claimant ages is critical to our research design. We use responses related to elderly retirement decisions to shed light on the drivers of early or late claiming of public pension benefits. Second, because our dataset covers primarily grocery store purchases, we cannot exclude the possibility that a large fraction of pension benefits are spent through other unobserved transactions. We use questions on typical elderly expenditure patterns in an attempt to quantify the share of total expenditures we capture in our data and determine the extent to which we underestimate the true response of

expenditures to payday. Third, we use detailed questions related to wealth and income to assess whether elderly households are liquidity constrained. These questions supplement our analysis in [Section 4.4](#) in which we proxy for liquidity constraints by looking at the sensitivity of payday responses to measures of permanent income.

Public expenditures on the elderly. We collect annual statistics from the Cabinet Office Historical Data on municipal expenditures towards elderly welfare and aggregate statistics on administrative costs from annual reports published by the national pension system. Such expenditures include local outlays for means-tested benefits targeting the elderly and any administrative costs the local government incurs from administering the national pension system. While the vast majority of public pension payments are processed through the national office and wired via direct deposit to claimants' bank accounts, the national government outsources day-to-day administrative functions to local branch offices.¹²

We merge these data with a list of locations for 245 branch offices tied to the national pension system which are responsible for processing and mailing benefits. All municipalities have a processing center (i.e. a city hall) that issues pension certificates, but larger municipalities ($\approx 14\%$) have a separate branch office which both provides data used to process payments and determines pension eligibility. Therefore, municipalities with a branch office are more likely to bear many of the administrative costs of implementing the pension system that would vary with benefit frequency. By comparing elderly expenditures in locations with a branch office to those without an office, we estimate in [Section 6.1](#) how the government's administrative costs for providing pension payments changed following a 1988 reform in which payments switched from a quarterly to bimonthly disbursement schedule.

3.2 BACKGROUND ON THE JAPANESE OLD-AGE PENSION SYSTEM

The Japanese National Pension System (JPS) is a two-tiered old-age insurance system consisting of (i) a National Pension (NP), a flat-rate basic pension with required contributions for residents aged 20 to 59, and (ii) Employee Pension Insurance (EPI), an earnings-related pension with compulsory coverage for those employed full-time by private companies (public employees receive a similar, albeit separate, earnings-based pension). Although the two JPS tiers are distinct in terms of how they determine contribution amounts, both tiers are implemented together as one system. Coverage under the NP and EPI is voluntary only for a few groups, including those with very low income, non-resident citizens, and elderly immigrants.

Full NP eligibility begins at age 65 for those with at least 25 years of coverage under the system. Participants have the option of claiming benefits early starting at age 60 in exchange for a permanently lower annual payment or delaying receipt (possible until age 70) in exchange for a

¹²Such administrative functions include providing consultations, processing applications, providing residency and payment records to the national office, confirming eligibility, investigating fraud, and reconciling benefits.

permanently higher annual payment. In 2012, the annual full benefit amount for the flat-rate annuity portion of the JPS was 780,100 JPY (\approx \$9,000 in 2012) for those with 40 years of contributions.

Annual EPI benefits are annuitized via the formula

$$EPI = (A - NP) + B + C$$

where NP is the flat-rate annuity the claimant receives under the basic NP tier of the system. A is a fixed amount times the number of contribution months, B is a function of average monthly earnings while employed, and C is an additional allowance for dependents. EPI eligibility begins at 65 but a “bridge” pension equal to $A + B + C$ is available for those aged 60-64. The EPI is subject to an earnings test where payments are reduced or suspended if the sum of the bimonthly EPI payment and wages exceeds a threshold of 460,000 in 2012 JPY (\approx \$4,050).¹³ Taken together, these features of the system imply that the vast majority of Japanese begin receiving some benefits from JPS once they turn 60. Our tabulations from the JSTAR data confirm this. Over 50% of pensioners in the JSTAR sample report claiming some public pension benefits by age 60 (the early retirement age), and over 95% begin claiming by age 65 (the normal retirement age).

Eligible claimants receive pension payments every two months on the 15th of each disbursement month (February, April, June, etc.). If the scheduled delivery date falls on a Saturday, Sunday, or public holiday, payments are instead sent on the first previous non-holiday weekday. This timing rule, combined with annual variation in calendar weekdays, induces variation in the length of periods in between payments that is unrelated to pensioners’ spending decisions, conditional on seasonality and intra-week shopping patterns. Our sample period covers 22 scheduled payments, seven of which were rescheduled due to overlap with weekends or holidays. The average interval length among these pay cycles is 59.8 days, with interval lengths ranging from 57 days to 62 days.

4 EXPENDITURE RESPONSE TO PAYMENTS

In this section we describe our main results and high-frequency difference-in-differences empirical strategy for identifying the effect of pension payment receipt on nominal expenditures.

¹³The system generates some income redistribution within age cohorts through B by down-weighting income earned in the 10 years prior to retirement; this feature limits the extent to which pensions replace income from retirement bonuses. More details on benefit formulas can be found at <http://www.nenkin.go.jp/international/english/healthinsurance/employee.html>.

4.1 IDENTIFICATION STRATEGY

We adopt the following modified version of the regression used in [Olafsson & Pagel \(2018\)](#) to estimate the effects of pension receipt on expenditures:

$$\frac{X_{i,c,t}}{\bar{X}_{i,c}} = \sum_{j=-7}^{+7} \beta_j \cdot \text{Payment}_{i,t+j} + \delta_{dow} + \phi_{wom} + \psi_{my} + \xi_h + \eta_i + \epsilon_{i,c,t} \quad (17)$$

where $X_{i,c,t}$ is expenditures of shopper i on goods within category c on date t , and $\bar{X}_{i,c}$ is average daily expenditures of the shopper on goods within the category. δ_{dow} are day-of-week fixed effects, ϕ_{wom} are week-of-the-month fixed effects, ψ_{my} are month-by-year fixed effects, and η_i are individual fixed effects. ξ_h is a dummy equal to 1 if date t falls on a public holiday. The dummy $\text{Payment}_{i,t+j}$ is an indicator equal to 1 if shopper i is scheduled to receive a payment at time $t+j$. The day-of-week dummies control for intra-week patterns of spending. Week-of-the-month dummies capture lumpy expenditures that take place regularly within some week of the month.¹⁴ The month-by-year dummies account for business cycle conditions and seasonality. We cluster standard errors at the individual level.

Since we do not directly observe pension receipt in our data, we set $\text{Payment}_{i,t+j}$ equal to 1 if the shopper is older than the normal eligibility age of 65 for the national pension and $t+j$ is a scheduled pension delivery day. We also report results using earlier and later age cutoffs for treatment, spanning the earliest (60) and the latest possible age (70) at which beneficiaries can begin claiming benefits. However, the change in the point estimates is negligible as we increase the age threshold. We find this unsurprising given that in the JSTAR data over 90% of retirees begin claiming prior to age 65, and roughly half begin claiming at age 60. The coefficients β_j measure the fraction by which a shopper’s expenditures deviate from the average daily level of spending in a one-week window around the scheduled pension delivery date.¹⁵

The model in (17) is a high-frequency difference-in-differences regression that compares the expenditure paths of pension eligibles to non-eligibles. Our identifying assumption is that there is no omitted variable that would differentially impact eligibles and non-eligibles on a scheduled payment date, conditional on time-invariant shopper characteristics and fixed effects that capture typical shopping patterns at monthly, weekly, and daily frequencies. One potential issue is that individuals in each of the two groups of shoppers might differ in their propensity to visit a store on a given date because non-pensioners are more likely to be working during the week than pensioners and thus conduct more of their shopping on weekends. This is a concern due to the fact that JPS

¹⁴These lumpy expenditures may be due to concurrent deadlines for mortgage, utility, and other bill payments, or retailer sales campaigns which offer discounts at certain times of the month. We discuss the possibility of the latter in [Section 5](#).

¹⁵In other words, because we do not directly observe pension receipt, our estimates $\hat{\beta}_j$ are intent-to-treat (ITT) estimates of the effect of pension receipt on expenditures. While some shoppers in our sample that are under age 60 might be eligible for spousal, survivor, or disability benefits which arrive on the same payment date, any expenditure response among the under age 60 population to these types of benefits would bias our estimates downward.

reschedules delivery dates that would otherwise fall on Saturdays or Sundays. Hence, the coefficients β_j may capture the effect of pension receipt plus differences across the treatment and control groups in the propensity to visit a store on a payment date, even if no payment were actually delivered.

We address this concern by augmenting the regression in (17) in two ways. First, we interact the individual fixed effects with the full set of day-of-week dummies to account for intra-week shopping patterns that are specific to each shopper. Second, we introduce a new variable, $Period_{i,c,t}$, defined as the number of days at date t since shopper i last purchased within category c . For instance, if a shopper makes a purchase at time t , but their last purchase was in $t - 30$, we set $Period_{i,c,t} = 30$; in $t + 1$ we then set $Period_{i,c,t+1} = 1$. We then create a set of dummies for each quartile of $Period$ and interact these dummies with individual fixed effects.¹⁶ The interpretation of $Period$ depends on the category of expenditures under consideration. For instance, for highly aggregated category such as all raw foods expenditures or total expenditures, this interaction will capture individuals' propensity to shop on certain days of the month. For less aggregated categories such as salad products, this interaction will instead capture the frequency with which shoppers consume goods in that category. Thus we believe the interaction with day-of-week effects is a more appropriate strategy for analyzing expenditure patterns of specific categories of goods.

Finally, we emphasize that because payment dates are rescheduled when a payday falls on a weekend or holiday, the presence of weekends and holidays generates exogenous variation in the day of the month that income arrives. Our specification isolates this exogenous variation even without interacting individual fixed effects with day-of-month fixed effects because all pension claimants receive benefits on the same day.

4.2 MAIN RESULTS FOR EXPENDITURES

Figure 5 plots the event study coefficients β_j and 95% confidence intervals obtained from estimating equation (17) for total goods expenditures. There is a clear spike in expenditures equal to 9% of average daily expenditures on the scheduled pension date. There are also positive spikes corresponding to 3-4% of average daily expenditures three days prior to and two days after the scheduled pension date. We speculate that these smaller responses within a few days of the delivery date are due to variation in the ability of personal banks to process transfers from the government to individual claimants' accounts. Overall, the cumulative response of expenditures within a one week window (3 days before until 3 days after the payment date) is 10% of average daily expenditures.

Shoppers spend more on payday for certain discretionary categories of goods such as alcohol, pre-packaged or prepared meals, and desserts. Figure 6 shows heterogeneity in the response of expenditures when we estimate the event study equation in (17) using expenditures on different categories of goods to construct the dependent variable. With the exception of the processed

¹⁶Since we always restrict to shoppers who make at least one purchase within a category each month, $Period$ falls between 1 and 30 days, and these quartiles correspond roughly to weeks since last purchase.

fruits/vegetables and processed fish categories, in all other major goods categories we observe a spending response that ranges between 5% to 10% of average daily expenditures on goods in that category. The result that discretionary categories of grocery spending react more strongly to regular income receipt reflects the notion that a large fraction of pension recipients exhibit payday liquidity, or act as if they have a license to splurge on more infrequently purchased goods.

Much of these observed total responses of goods expenditures occur along the intensive margin. In other words, in even months when payments are delivered, shoppers who receive a payment spend more at the store around the payment date but are not significantly more likely to make a trip to a store compared to the same date in odd months when no payment is received. [Table 2](#) provides estimates of the intensive and extensive margin response for total expenditures and major goods categories. To obtain estimates of the intensive margin response to pension receipt, we estimate a version of equation (17) defined at the individual-trip level (i.e. excluding observations within the panel where we observed zero spending). For the extensive margin, we instead replace the dependent variable in equation (17) with an indicator equal to unity if the shopper’s expenditures were strictly positive on a given date. For all goods expenditures, the intensive margin estimates indicate that shoppers spent 13.4% more on payday visits relative to average expenditures on other days when they visited a store. On the extensive margin, the probability that a consumer eligible for pension payments visits a store is roughly 1% higher on payday. The results by goods category mirror our event study results: both the intensive and extensive margin responses are more pronounced for discretionary goods spending.

4.3 HETEROGENEITY BY PAY CYCLE LENGTH

The results in the previous subsection on pensioners’ expenditure response to payday reflect average responses across pay cycles with slightly different lengths. Since payment dates are rescheduled when a payday falls on a weekend or holiday, pay cycle lengths in our sample time period range from 57 to 62 days. As noted in our presentation of the payday liquidity model of the consumption path in [Section 2.3](#), there may be more pent-up demand when pay cycles are longer. In other words, the magnitude of the spike in expenditures may be greater when benefit recipients must wait longer in between consecutive payments. This is of direct relevance to the government’s problem of setting the optimal payment frequency, because the presence of pent-up demand implies that shorter pay cycles help limit the extent of consumption non-smoothing that is concentrated on payday.

We test for this pent-up demand mechanism in shopper behavior by augmenting equation (17) with terms that interact the payment dummy with a polynomial function of the pay cycle length:

$$\begin{aligned} \frac{X_{i,c,t}}{\bar{X}_{i,c}} = & \beta_1 \cdot Payment_{i,t} \times Length_p + \beta_2 \cdot Payment_{i,t} \times (Length_p)^2 \\ & + \beta_3 \cdot Payment_{i,t} \times (Length_p)^3 + \delta_{dow} + \phi_{wom} + \psi_{my} + \xi_h + \eta_i + \epsilon_{i,c,t} \end{aligned} \quad (18)$$

where $Length_p$ represents the length in days of pay cycle p . It is useful to analyze this regression specification through the lens of the payday liquidity model. When $Payment_{i,t} = 0$, because individual i is either not eligible for pension payments or i is eligible but t is not a payday, then in the payday liquidity model payday consumption is determined by the smoothed value: $C_0 = \bar{c}$. In contrast, when $Payment_{i,t} = 1$, equation (18) says that payday consumption is some multiple of the smoothed value: $C_0 = \bar{c} \cdot (1 + \beta_1 T + \beta_2 T^2 + \beta_3 T^3) \equiv \bar{c} \cdot (1 + x(T))$.

Table 3 provides evidence from estimating equation (18) that shoppers tend to spend more on payday after a longer pay cycle. In the baseline case (column 1) where we assume a linear interaction between payment and pay cycle length, this translates to a 0.151 percentage point increase in the magnitude of the spike per extra day pensioners have to wait for a payment to arrive. The estimated specifications in even columns allow for a direct effect of interval length on expenditures. The fact that this direct effect is negative and highly significant suggests that there are spillover effects of the interval length to shoppers who themselves are not eligible for pension receipt.¹⁷ However, allowing for this direct effect of pay cycle length barely changes the fitted values for $x(T)$, which we report in rows 5 through 7 of the table for different pay cycle lengths we observe in our sample.

Although the results in column (5) and column (6) of Table 3 provide evidence of non-linearities in pent-up demand following the end of a pay cycle, we note that due to the limited variation in length of the fourteen pay cycles that occur within our sample time period, this polynomial specification is not particularly well-defined. From this exercise, we take the coefficient on the linear interaction $Payment \times Length$ in column (3) to obtain our preferred estimate of the expenditure spike function: $x(T) = 0.0023 \cdot T$. We use this estimate to calibrate the payday liquidity version of our optimal frequency model in Section 6.

4.4 HETEROGENEITY IN NON-SMOOTHING BEHAVIOR

The evidence presented so far shows that individuals exhibit a clear spike in spending on or around the receipt of predictable and regular pension payments. What are the underlying economic mechanisms driving these responses? While our findings are consistent with features of the payday liquidity and quasi-hyperbolic discounting versions of the theoretical framework in Section 2, other explanations, such as near-rationality or liquidity constraints could play a role.

Distinguishing between these mechanisms is important for the welfare analysis we conduct in this paper. Near-rational consumers induce a welfare loss from excess spending that is proportional to the size of the payment relative to their permanent income (Kueng, 2018). Recipients for whom the payment is a large fraction of permanent income will smooth consumption due to the higher potential welfare loss, and those for whom the payment is a small fraction of permanent income will not smooth consumption but incur very small welfare losses. In any case, if pay cycle length

¹⁷For instance, if a non-claimant is shopping on payday then their consumption might be crowded out by the pent-up demand of their claimant spouse who just received a payment.

is unrelated to the size of the payment relative to permanent income, near-rationality implies the government cannot influence the size of the welfare loss by altering payment frequency.

Our framework applies to a regular stream of payments with predetermined amounts. Moreover, in our empirical application, unanticipated variation in consumption needs over a pay cycle is likely to be limited, as universal health insurance coverage in Japan reduces consumption risk from illness and most retirees own their homes. For these reasons, if liquidity constraints exist, they will be highly idiosyncratic in nature. It is therefore difficult to envision how the government’s choice of payment frequency could influence welfare by relaxing short-term liquidity constraints.

To further shed light on the motivations for shoppers’ spending around payday, we look at heterogeneity in the cross-section of shoppers by average total expenditures over two-month pay cycles in our sample time period. We follow [Kueng \(2015, 2018\)](#) in using total expenditures to proxy for permanent income, as we do not observe other potential proxies for permanent income, such as earned income streams.¹⁸ We restrict to shoppers aged 65 or over who are eligible for pension payments and who are regular shoppers, defined as those who visit a store at least four times per month, or one week on average. These restrictions leave us with roughly 115,000 shoppers.

We run our main specification in equation (17) separately for each of these shoppers using intensive margin raw foods expenditures as the outcome variable and collect the coefficients on the *Payment* dummy. The raw foods category refers to highly perishable goods, so that expenditures within this category are reasonable proxies for instantaneous consumption.¹⁹ [Figure 7](#) plots the distribution of these individual payday responses binned by quartiles of average total expenditures over the pay cycle. The figure shows that spending responses are more variable among low permanent income shoppers; such shoppers visit the store less frequently and have more “lumpy” expenditures than high permanent income shoppers. Yet, the mean response is relatively flat across permanent income bins. Average payday responses are 4.8% in excess of average daily expenditures for the first quartile, 5.7% for the second, 4.9% for the third, and 3.4% for the fourth quartile.

[Figure 8](#) shows how payday responses vary continuously with our permanent income measure by plotting local smoothed means of the coefficients on *Payment* against average pay cycle expenditures. There is a clear negative relationship between spending on payday and permanent income for shoppers with below-median pay cycle expenditures on fresh groceries (25,000 JPY \approx \$250). However, this relationship becomes virtually flat once we look at shoppers with above-median average pay cycle expenditures.

We also estimate extensive margin regressions by permanent income decile where we replace the outcome variable in (17) with a dummy equal to one if the shopper visits the store on a given date. [Table 4](#) summarizes the payday responses by permanent income decile along the intensive, extensive, and total response (intensive \times extensive) margins. Along all three margins, there is a

¹⁸We emphasize that in our setting there is limited scope for measurement error from using expenditures as a proxy for permanent income since purchases are not self-reported by households.

¹⁹We provide more details on the contents of this category in [Appendix B](#).

clear negative relationship between payday responses and permanent income below the median, but a flat relationship above the median. The results here are consistent with a robust finding in the literature that excess spending responses to predictable payments are typically concentrated among low-income individuals (e.g. [Zeldes 1989](#); [Broda & Parker 2014](#)). To the extent that the ratio of pension payments to income is relatively constant across the distribution of total expenditures, our findings cast doubt on the near-rationality hypothesis advanced in [Kueng \(2018\)](#), under which we would expect a flat profile of spending responses with respect to total expenditures.²⁰

However, we caution against the interpretation of our results as evidence that liquidity constraints are driving payday expenditures in our setting for two reasons. One reason is that our setting features predictable, lump-sum payments, which most shoppers above the age of 65 have been receiving for many years; the average age among shoppers over 65 in our sample is 71. It is then difficult to argue that recipients lack the ability to build up a buffer to smooth variations over the pay cycle ([Fuchs-Schündeln & Hassan 2016](#)). Second, as noted by [Parker \(2017\)](#), behavioral factors such as present-bias and rule of thumb spending could contribute to both low income and excess spending responses to predictable payments. Our findings of heterogeneous responses by permanent income would therefore also be consistent with a higher incidence of behavioral traits such as present-bias and mental accounting among low-income shoppers.

5 DO RETAILERS CAPTURE THE INCIDENCE OF PAYMENTS?

So far we have demonstrated how pension receipt affects nominal expenditures. However, if retailers engage in price discrimination around scheduled payment dates in anticipation of an increase in demand, the payday coefficients β_j in equation (17) will capture changes in prices even when the real expenditure amount stays the same. In this section, we investigate how prices of individual goods respond to pension payment distribution dates. In particular, we construct a price index that isolates the retailer’s pricing response from consumer substitution across goods.

5.1 STORE-LEVEL PRICE INDICES

We first consider daily geometric average prices within a store:

$$\ln P_{s,t} = \frac{1}{n_{s,t}} \sum_k \ln p_{k,s,t} \tag{19}$$

²⁰The relationship between past earnings and pension amounts is linear for most recipients. There is some concavity in the benefits schedule for recipients who worked at larger firms who typically receive large retirement bonuses equal to several years of earnings. Under recent program rules, bonus amounts that can count towards pension benefit calculations are capped at 1.5 million yen per month or 5.73 million yen in any year. To help ensure that payments are a constant share of income throughout the distribution, we winsorized average pay cycle expenditures at the 1st and 99th percentile thresholds to limit the influence of very low and very high permanent income individuals and obtained qualitatively similar results.

where $p_{k,s,t}$ is the price for good k at store s on day t , and $n_{s,t}$ is the number of goods sold at store s at day t . We compute the price of individual good $p_{k,s,t}$ by dividing its total sales by total quantity purchased by shoppers over age 65 at store s at day t . Thus, this price index $\ln P_{s,t}$ represents the average (log) price of goods that consumers over age 65 buy on each day in each store.

To estimate the effect of pension receipt on the average price, we run the following regression:

$$\ln \widehat{P}_{s,t} = \sum_{j=-7}^{+7} \gamma_j \cdot \text{Payday}_{t+j} + \delta_{dow} + \phi_{wom} + \psi_{my} + \xi_h + \eta_s + \epsilon_{s,t} \quad (20)$$

where $\ln \widehat{P}_{s,t}$ is the daily average price index at store s , normalized by its mean value over our sample period. As in the expenditure regressions described in [Section 4.1](#), δ_{dow} are day-of-week fixed effects, ϕ_{wom} are week-of-the-month fixed effects, ψ_{my} are month-by-year fixed effects, and η_s are store fixed effects. ξ_h is a dummy equal to 1 if date t falls on a public holiday. Payday_{t+j} is a dummy variable that takes 1 if pension payments are scheduled at time $t+j$. The payday coefficients γ_j are interpreted as how the daily average price of goods purchased within a store deviates from its sample mean if consumers receive pension payments at time $t+j$.

[Figure 9.A](#) plots the coefficients on the payday dummies γ_j for a two-week interval around pension arrival dates. On pension payment dates, the average store-level price tends to be 1.6% higher than its sample mean. There are also small but precisely estimated positive effects around pension payment dates. A change in the geometric average price between time t and $t+1$ consists of a change in the number of goods purchased and a change in prices of goods purchased. [Figure 9.B](#) shows estimates of equation (20) where we replace prices on the left-hand side with the log number of goods sold $\ln \widehat{n}_{s,t}$ normalized by its mean value over our sample period. On payday, the number of goods increases by about 6%. This increase in the variety of goods puts downward pressure on the price index measured in equation (19).

These results indicate that the positive response of store-level prices to payday could be due to retailers raising their prices around pension payment dates, consumers substituting towards higher quality goods when they receive a payment, or some combination of the two responses. If consumers tend to buy higher quality goods around pension paydays, the average price paid by consumers could increase without any change in prices of individual goods set by retail stores to capture the incidence of pension payments. In the next subsection, we introduce a simple price index that isolates the change in within-store average prices due to consumer substitution across goods.

5.2 PRICE DISCRIMINATION OR SUBSTITUTION EFFECTS?

We consider a counterfactual where there is no retailer pricing response in even months in which pension payments are delivered. In each even month, we replace the prices of individual goods $p_{k,s,t}$ by their last observed daily value in the preceding odd month. We then fix these individual prices over time within the adjacent even month. [Figure 10](#) illustrates this with an example contrasting

actual and counterfactual price series. The black line depicts the actual price of an individual good while the red line represents its counterfactual price fixed at the last observed price in the preceding odd month. Thus, the counterfactual price is the price that consumers could have paid if there had been no retailer pricing response in even months. By fixing prices to be equal to prices in odd months, so that time variation in the price index in payment months only comes from a change in the set of goods purchased by consumers, we can isolate a change in the price index due to consumers' substitution across goods from a change due to a retailer response such as temporary sales. The difference between the black line and the red line in the figure represents the response of a good's price due to retailer intertemporal price discrimination.

Figure 11 plots the estimates of the coefficients on *Payday* obtained from equation (20) using observed prices minus the same coefficients from using the store-level counterfactual price indices as the dependent variable. The estimated differences are positive for each day within a two-week window around payday, implying that consumers would have paid lower prices if there were no retailer pricing response. However, these differences between the actual and counterfactual pricing responses around payday are economically insignificant. Actual prices are 1.58% higher on payday, while counterfactual prices are 1.40% higher. The gap of 0.18 p.p. between the two estimates captures the retailer's pricing response, which accounts for only 11.39% of the overall spike in payday prices. This demonstrates that while retailers capture a small portion of the incidence of pension payments, the observed increase in average prices around payday is almost entirely due to consumer substitution towards goods of higher value.²¹

6 MODEL CALIBRATION

In this section we describe how we calibrate our optimal payment frequency model using our estimates of the expenditure response to payments in the preceding sections and statistics on costs and benefit expenditures from JPS.

6.1 ESTIMATING ADMINISTRATIVE COSTS

Our identification of the slope of the administrative cost function in our model $\mu'(T)$ relies on a reform to the pension system in 1988 which reduced the length of the pay cycle from three months to two months. The provisions of this reform are discussed in detail in Stephens & Unayama (2011). In February 1988, poor elderly recipients of the Old Age Welfare Pension benefit and existing claimants who only qualified for the national pension but not the employee-based pension benefit

²¹We obtain qualitatively and quantitatively similar results to those reported in this subsection when we instead use the last observed price recorded at least one week before a payday. The implied retailer response from comparing the observed increase in prices on payday to the counterfactual increase in prices is 0.13 p.p. compared to 0.18 p.p. under our preferred counterfactual.

started receiving benefits in each even month. Hence, the bulk of payments from JPS were switched from a quarterly to bimonthly schedule starting in FY 1987.

Since the 1988 reform did not result in any new changes to the formula determining overall benefit amounts, this shift in policy represents a pure shock to T that did not shift the intercept of the administrative cost function $\mu(\cdot)$. While the policies governing the national pension system are set at a national level, many of the day-to-day functions of implementing the system are delegated to local town governments and local branch offices of the system which tend to be located in municipalities which lie in the center of a commuting zone. In particular, while claimants can apply to start receiving benefits at their local town hall office, the closest branch office is responsible for functions such as processing applications, reconciling benefits, confirming eligibility, and investigating fraud. This fiscal federalism built into the system implies that local governments with a branch office within their jurisdiction were more exposed to the shock to administrative costs, and thus would have seen expenditures allocated to the elderly welfare rise more relative to non-branch office governments following the reform.

Given this natural experiment, we run standard difference-in-differences regressions of the following form:

$$\log \mu_{j,t} = \beta \cdot Branch_j \times Post_t + \gamma_j + \delta_t + \epsilon_{j,t} \quad (21)$$

where $Branch_j$ is a dummy equal to 1 if municipality j contains a branch office of the Japanese National Pension System, $Post_t$ is a dummy equal to 1 if year t is FY 1987 or later, γ_j are municipality fixed effects, and δ_t are (fiscal) year fixed effects. The dependent variable is the log of costs per person over 65 associated with providing elderly welfare benefits, which include costs incurred from supporting the pension system. We consider this measure to be a proxy for administrative cost per claimant, since over 90% of individuals over age 65 receive benefits from JPS but a relatively small share of this population receive other forms of welfare benefits. Under the assumption of parallel trends in our cost per claimant measure, the coefficient β captures the mid-run increase in costs due to increasing the frequency of benefits from every three months to every two months without changing the formula for benefit amounts.

Figure 12 plots the average of this cost measure across all cities with a JPS branch office along with the average across all non-branch cities. The level and slope of costs across the two groups of cities is very similar up until 1987 when the government announced the reform.²² The cost series continue to diverge until the government froze increases in contribution rates for several years starting in December 1998. The government then proposed an overhaul of the system in 1999

²²Since the municipal budget data are for a fiscal year, which runs from April 1 to March 31 in Japan, the divergence in the cost series beginning in 1987 is due to overlap of the 1987 fiscal year with the start of the new payment schedule in February 1988, rather than an anticipation effect.

(passed in 2000) which aimed to reduce costs and aggregate spending and benefits.²³ To isolate variation in municipal administrative costs due to the 1988 reform which changed only the payment frequency, we restrict our sample to the years 1980-1996, which creates a symmetric window around the reform.

Table 5 presents results from estimating equation (21). In some specifications we include pre-reform quintiles of city size and per capita income interacted with year fixed effects to account for the fact that branch office and non-branch office cities are not balanced on these characteristics. Administrative costs increase by approximately 4.3% in branch office cities in the post-reform years, or 7.3% conditional on municipal population and income. The effect we uncover in columns (1) and (7) disappears once we exclude the 23 governments in metropolitan Tokyo and the five largest cities outside Tokyo from our sample. We speculate that JPS branch offices in central Tokyo were impacted more than other large cities with branch offices due to the comprehensive public transport system reducing the commute time to these offices. In short, we find limited evidence that the switch from a quarterly to bimonthly payment schedule had an economically significant effect on government costs of administering the public pension system.²⁴

In our optimal frequency model, the government’s administrative cost $\mu(T)$ is a function of the number of days T between payments. Taking our baseline estimate of a 4.3% increase in costs associated with a decrease in the pay cycle from 90 to 60 days implies a 0.14% increase in costs for each day the government shortens the pay cycle, assuming $\mu(T)$ is linear on the interval $T \in [60, 90]$. To match our aggregate statistics on the 2012 pension system and calibrate our model, we translate this estimate to real 2012 expenditures per claimant. Scaling reported total administrative costs for FY 2011 by our 4.3% estimate implies an increase in 8,111 JPY per claimant. The implied annual cost for each day the pay cycle length is shortened amounts to 270 JPY per claimant. For our upper bound estimate of a 7.3% increase in administrative costs, the implied cost increase is 8,344 JPY per claimant, or 278 JPY per claimant for each day the pay cycle length is shortened.²⁵

²³Major tenets of the 2000 reform included switching from wage-indexing to CPI-indexing of benefits, introducing an earnings test for claimants aged 65 to 69, reducing earnings-related benefits by 5%, and gradually increasing the normal retirement age from 60 to 65 (Takayama, 2001).

²⁴We acknowledge one drawback to this approach is external validity of the results to our sample time period of 2011-2014 for the retail scanner data. Prior to 2008, JPS did not offer internet services, so participants had to visit a local JPS office or a branch office to check their contributions and balances. The slope of the administrative cost function may not be a time-invariant policy parameter if the introduction of internet services reduced the sensitivity of administrative costs to payment frequency.

²⁵A simple non-parametric differences in means between the branch office and non-branch office cities pre- vs. post-reform yields an estimate of an increase of 14,464 JPY (in 2012 yen) per claimant due to moving to a bimonthly payment schedule, which implies an increase in 482 yen per claimant per pay cycle day. The absolute difference from our measure based on scaling up FY 2011 administrative costs is primarily due to the larger population of claimants in 2011-2012 relative to the 1980s and 1990s.

6.2 CALIBRATION PROCEDURES

We now compute the optimal payment frequency by combining our empirical results on the behavior of consumption expenditures within the pay cycle with our results on the shape of the government’s administrative cost function for providing pension benefits. This payment frequency is the solution to the government’s optimization problem in equation (6) and characterized by the first-order condition in (7). We perform computations for the two main versions of our model featuring quasi-hyperbolic discounters and payday liquid consumers. In the version with payday liquid consumers, we assume that the observed spike in consumption on payday is an increasing function of the pay cycle length T , as modeled in equation (16). This assumption is consistent with our empirical results in Section 4.3, where we showed that the spike in expenditures on payday increases by 0.23 percentage points for each additional day in the pay cycle.

Our modeling framework consists of two sets of parameters: preference parameters captured by ρ in the isoelastic utility case and the deviation function $f(t)$; and parameters associated with the government budget captured by p, \bar{B} . We analyze how the behavior of the welfare loss and optimal frequency varies with ρ and any latent parameters of $f(t)$ for each version of our model outlined in Section 2. To calibrate the government’s budget constraint, we rely on information on aggregate flows and participants in FY 2011 from the Japanese Pension System Annual Report 2012.²⁶

For the administrative cost function, we assume a convex cost function of the form $\mu(T) = \kappa_\ell/T^\ell$ and for each power ℓ calibrate κ_ℓ so that $\mu(60)$ equals the administrative service costs reported for FY 2011 in the annual JPS Business Report.²⁷ In FY 2011 there were 67.37 million contributors to the system, and 38.67 pension recipients, implying that pensioners made up $p = 0.377$ of all participants. Over the entire fiscal year total pension payments amounted to 48,867.5 billion yen, or roughly 10% of nominal 2012 GDP. This amount implies an average daily payment per claimant of $\bar{B} = 3,462$ yen (\approx \$32.50), or an average payment of 207,732 (\approx \$1,950) over the pay cycle. JPS reports administrative service costs of 300.722 billion yen in FY 2011.

Our results in Section 6.1 from exploiting differential exposure of local governments to the 1988 pension schedule reform suggest that $\mu(T) = \kappa_\ell/T^\ell$ with $\ell = 0.27$, calibrated to the aggregate system flows in 2012, is a good approximation for the administrative cost function.²⁸ While we view this version of the government’s cost function as empirically plausible, we also provide results for more or less convex cost functions by altering ℓ to obtain conservative upper bound estimates for the optimal frequency.

²⁶The official pension system report is available at http://www.nenkin.go.jp/files/13nenkinD_synthesis.pdf.

²⁷A copy of the FY 2011 JPS report can be found at <http://www.nenkin.go.jp/info/disclosure/jigyo.files/23-2.pdf>.

²⁸More concretely, we set $\mu(90) - \mu(60)$ equal to the cost per claimant increase implied by our regression estimates in Table 5. The formula used to obtain this value is $\ell^* = \frac{\log(1 - \hat{x} \cdot n/c)}{\log(2/3)}$, where \hat{x} is the cost per claimant increase implied by our regression estimates, n is the number of pension recipients, and c is total administrative costs reported by the government. Using our preferred estimate of a 4.3% cost increase from the 1988 reform, we obtain $\ell = 0.27$. For our upper bound estimate of a 7.3% cost increase from the 1988 reform, we obtain $\ell = 0.28$.

Lastly, we choose preference parameters that match the results from estimating expenditure regressions of the form in equation (17). We found that expenditures increase by roughly 10% for all goods or all raw foods expenditures upon payment receipt. In the payday liquidity model, this corresponds to a value of $x = 0.1$, which is one-third the estimate of $x = 0.3$ obtained by [Olafsson & Pagel \(2018\)](#) in their study of total expenditures on consumption and non-consumption goods. The estimates are consistent to the extent that grocery store items are less likely than other non-grocery categories of expenditures to be classified by consumers as splurge goods. We use our results on heterogeneity by pay cycle length in [Section 4.3](#) to set $x'(T) = 0.0023$, reflecting that the payday spike in expenditures tends to be 0.23 percentage points higher for each additional day in pay cycle.

Our estimate of a 10% spike in perishable grocery expenditures on payday also allows us to calibrate the degree of present-bias behavior in the version of the model with quasi-hyperbolic discounters. For quasi-hyperbolic discount factors β close to one, the decline in consumption over the pay cycle is approximately linear. Hence, to calibrate the overall daily discount rate in the model, we can set a linear rate of negative consumption growth between time $t = 0$ (payday) and time $t = T - 1$. The fact that raw foods expenditures increase by 10% between the day before the new pay cycle starts and the day a new pay cycle begins implies that this daily rate of negative consumption growth is $\nu = 0.10/60 \approx 0.002$, or 0.2% of payday consumption. This estimate is half of the $\nu = 0.004$ estimate found by [Shapiro \(2005\)](#) in his analysis of food stamp recipients' consumption between monthly pay cycles, which suggests that present-bias plays a less prominent role in our sample of pension recipients.

6.3 OPTIMAL PAYMENT FREQUENCY RESULTS

[Figure 13](#) presents the optimal payment frequency for each version of our model, with different values of the inverse IES $\rho \in [1, 50]$ and for three different versions of the government's convex administrative cost functions: $\mu(T) = \kappa_\ell/T^\ell$ for $\ell = 1, 2, 3$. The figure showcases two key results. First, in each model version, the optimal frequency is almost completely invariant to the assumed value for the inverse IES, and the optimal frequency is (weakly) decreasing in ρ . Intuitively, ρ determines the size of the welfare loss from consumption non-smoothing but not the increase in marginal costs to the government of decreasing T . However, the marginal welfare loss $-\lambda'(T)$ is largely invariant to the inverse IES. The optimal pay cycle length T^* is decreasing in ρ because the marginal benefit to the government of decreasing T is steeper for higher ρ . This is because a higher ρ means that consumption is less substitutable between periods, so individuals would be willing to pay more to get closer to the smooth consumption path in the *ex ante* sense in which we have defined the welfare loss.

Second, for each of the administrative cost functions across both model versions, the implied optimal frequency is less than one month. For $\ell = 1$, the optimal frequency is $T^* = 6.72$ days, whereas for $\ell = 2$ we obtain $T^* = 18.27$ days, and for $\ell = 3$ we obtain cycles of length $T^* = 27.33$ days. In [Figure 14](#), we show how T^* varies continuously as a function of the convexity of

administrative costs for a fixed value of the inverse IES ($\rho = 10$).²⁹ For very non-convex cost functions ($\ell \leq 0.45$) costs become negligible and the government finds it optimal to disburse daily payments. Even for extremely convex administrative costs the optimal interval length remains below two months. Thus, the government can improve welfare by distributing pension payments more frequently than the current bimonthly schedule.

7 CONCLUSION

We ask how governments should choose the frequency of public benefits. In our simple model the government chooses the length of the pay cycle while facing a tradeoff between minimizing the welfare loss associated with consumption non-smoothing and incurring increased administrative costs. In an empirical application, we combine results on consumer and retail responses to bimonthly Japanese old-age pension payments with evidence on how administrative costs vary with pay cycle length to calibrate two versions of the model: one in which consumers are quasi-hyperbolic discounters and another in which consumers deviate from a smooth path only on payday. Each version of the model results in an optimal pay cycle length of around one month, suggesting the government could improve welfare by moving to a monthly disbursement schedule.

Regular shoppers increase overall grocery expenditures by 10% within two days of the scheduled payment date, but most of this expenditure response is concentrated within categories of discretionary goods such as alcohol, desserts, and prepared foods. These findings accord with models of intertemporal decision-making that feature rule-of-thumb spending or mental accounting of payments. We find little evidence to support the notion that the government's choice of payment frequency could influence the degree of price discrimination around paydays. While within-store average prices increase by 1.6% on payday, we show using a simple counterfactual price index that the observed positive payday effect on average store-level prices is almost entirely due to consumers substituting towards a basket of more numerous, higher quality goods, rather than retailers raising prices to capture the incidence of pension payments.

Although our empirical application focuses on the Japanese pension system, our framework applies more generally to other public benefits programs in advanced economies. In line with the literature on the welfare loss from consumption non-smoothing, we demonstrate under empirically plausible models for individual consumption paths that this welfare loss is likely to be small even under extreme assumptions about the intertemporal elasticity of substitution. The main lesson here is that for public benefits programs where administrative costs vary minimally with respect to the frequency of disbursement (e.g. due to economies of scale or limited scope for benefits fraud), the government can achieve a welfare improvement by splitting up entitlements into smaller, more frequent payments.

²⁹Again, the differences across the two versions of the optimal frequency model are only present after the sixth decimal point. Intuitively, this is because both models produce marginal welfare losses that are very flat with respect to payment frequency.

REFERENCES

- Aguiar, M. & E. Hurst** (2005): "Consumption versus Expenditure," *Journal of Political Economy*, 113(5): 919-948.
- Baker, S.** (2018): "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data," *Journal of Political Economy*, 126(4): 1504-1557.
- Battistin, E., A. Brugiavini, E. Rettore, & G. Weber** (2009): "The Retirement Consumption Puzzle: Evidence from a Regression Discontinuity Approach," *American Economic Review*, 99(5): 2209-2226.
- Baugh, B. & J. Wang** (2018): "When Is it Hard to Make Ends Meet?," *mimeo*, University of Illinois at Urbana-Champaign.
- Bernheim, B.D., J. Skinner, & S. Weinberg** (2001): "What Accounts for the Variation in Retirement Wealth Among U.S. Households," *American Economic Review*, 832-857.
- Best, M., J. Cloyne, E. Ilzetzki, & H. Kleven** (2018): "Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches," forthcoming *Review of Economic Studies*.
- Broda, C. & J.A. Parker** (2014): "The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption," *Journal of Monetary Economics*, 68: S20-S36.
- Browning, M. & T.F. Crossley** (2001): "The Life-Cycle Model of Consumption and Saving," *Journal of Economic Perspectives*, 15(3): 3-22.
- Chevalier, J.A., A.K. Kashyap, & P.E. Rossi** (2003): "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," *American Economic Review*, 93(1): 15-37.
- Chevalier, J.A., A.K. Kashyap** (2019): "Best Prices: Price Discrimination and Consumer Substitution," *American Economic Journal: Economic Policy*, 11(1): 126-159.
- Cochrane, J.H.** (1989): "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *American Economic Review*, 79(3): 319-337.
- Dobkin, C. & S.L. Puller** (2007): "The Effects of Government Transfers on Monthly Cycles in Drug Abuse, Hospitalization and Mortality," *Journal of Public Economics*, 91(11-12): 2137-2157.
- Farhi, E. & X. Gabaix** (2015): "Optimal Taxation with Behavioral Agents," NBER Working Paper, No. 21524.
- Foley, C.F.** (2011): "Welfare Payments and Crime," *Review of Economics and Statistics*, 93(1): 97-112.

- Fuchs-Schündeln, N. & T.A. Hassan** (2016): “Natural Experiments in Macroeconomics,” *Handbook of Macroeconomics*, 2: 923-1012.
- Gelman, M., S. Kariv, M.D. Shapiro, D. Silverman, & S. Tadelis** (2014): “Harnessing Naturally Occurring Data to Measure the Response of Spending to Income,” *Science*, 345(6193): 212-215.
- Goldin, J., T. Homonoff, & K. Meckel** (2019): “Issuance and Incidence: SNAP Benefit Cycles and Grocery Prices,” *mimeo*, NYU Wagner.
- Hastings, J. & E. Washington** (2010): “The First of the Month Effect: Consumer Behavior and Store Responses,” *American Economic Journal: Economic Policy*, 2(2): 142-162.
- Hastings, J. & J.M. Shapiro** (2018): “How Are SNAP Benefits Spent? Evidence from a Retail Panel,” *American Economic Review*, 108(12): 3493-3540.
- Kueng, L.** (2015): “Revisiting the Response of Household Spending to the Alaska Permanent Fund Dividend using CE Data,” *mimeo*, Northwestern Kellogg.
- Kueng, L.** (2018): “Excess Sensitivity of High-Income Consumers,” *Quarterly Journal of Economics*, 133(4): 1693-1751.
- MacDonald, J.M.** (2000): “Demand, Information, and Competition: Why Do Food Prices Fall at Seasonal Demand Peaks?” *The Journal of Industrial Economics*, 48(1): 27-45.
- Mastrobuoni, G. & M. Weinberg** (2009): “Heterogeneity in Intra-Monthly Consumption Patterns, Self-Control, and Savings at Retirement,” *American Economic Journal: Economic Policy*, 1(2): 163-189.
- Nevo, A. & K. Hatzitaskos** (2006): “Why Does the Average Price Paid Fall during High Demand Periods,” *mimeo*, Northwestern University.
- Olafsson, A. & M. Pagel** (2018): “The Liquid Hand-to-Mouth: Evidence from Personal Financial Management Software,” *Review of Financial Studies*, 31(11): 4398-4446.
- Parker, J.A.** (2017): “Why Don’t Households Smooth Consumption? Evidence from a 25 Million Dollar Experiment,” *American Economic Journal: Macroeconomics*, 9(4): 153-183.
- Parsons, C.A. & E.D. Van Wesep** (2013): “The Timing of Pay,” *Journal of Financial Economics*, 109: 373-397.
- Shapiro, J.M.** (2005): “Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle,” *Journal of Public Economics*, 89(2-3): 303-325.
- Shimizutani, S. & T. Oshio** (2016): “Public Pension Benefits Claiming Behaviour: New Evidence from the Japanese Study on Ageing and Retirement,” *The Japanese Economic Review*, 67(3): 235-256.

- Shoji, T.** (2018): “Liquidity Constraints, Storage Costs, and Consumer Stockpiling,” *mimeo*, University of Tokyo.
- Stephens Jr., M.** (2003): “‘3rd of the Month’: Do social Security Recipients Smooth Consumption between Checks,” *American Economic Review*, 93(1): 406-422.
- Stephens Jr., M.** (2006): “Paycheque Receipt and the Timing of Consumption,” *The Economic Journal*, 116: 680-701.
- Stephens Jr., M. & T. Unayama** (2011): “The Consumption Response to Seasonal Income: Evidence from Japanese Public Pension Benefits,” *American Economic Journal: Applied Economics*, 3(4): 86-118.
- Stephens Jr., M. & T. Unayama** (2012): “The Impact of Retirement on Household Consumption in Japan,” *Journal of the Japanese and International Economies*, 26(1): 62-83.
- Takayama, N.** (2001): “Pension Reform in Japan at the Turn of the Century,” *The Geneva Papers on Risk and Insurance*, 26(4): 565-574.
- Thaler, R.H.** (1999): “Mental Accounting Matters,” *Journal of Behavioral Decision Making*, 12(3): 183-206.
- Warner, E.J. & R.B. Barsky** (1995): “The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays,” *Quarterly Journal of Economics*, 110(2): 322-352.
- Zeldes, S.P.** (1989): “Consumption and Liquidity Constraints: An Empirical Investigation,” *Journal of Political Economy*, 97(2): 305-346.

FIGURE 1.A. QH Discounting: Welfare Loss as Function of Interval




FIGURE 1.B. QH Discounting: Welfare Loss as Function of Inverse IES




FIGURE 2.A. Payday Liquidity: Welfare Loss as Function of Interval




FIGURE 2.B. Payday Liquidity: Welfare Loss as Function of Inverse IES




FIGURE 3.A. Payday Liquidity $x(T)$: Welfare Loss as Function of Interval




FIGURE 3.B. Payday Liquidity $x(T)$: Welfare Loss as Function of Inverse IES




FIGURE 4: Shoppers who Visit Store k Times per Month

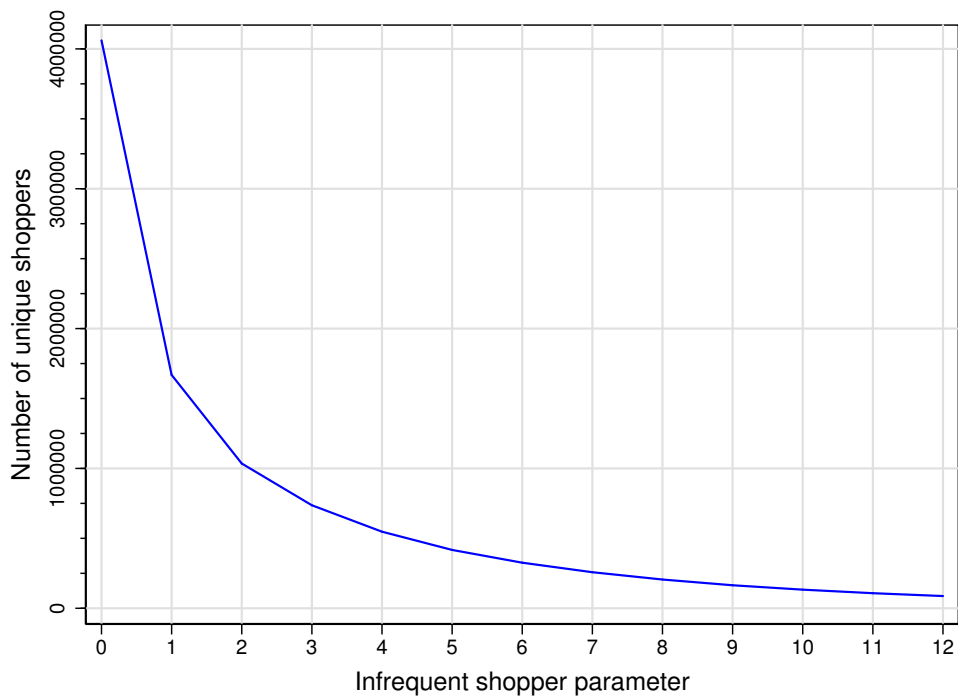


FIGURE 5: Response of Total Expenditures to Payday




FIGURE 6: Response of Major Subcategory Expenditures to Payday



Figure 7: Payday Responses by Quartile of Pay Cycle Expenditures

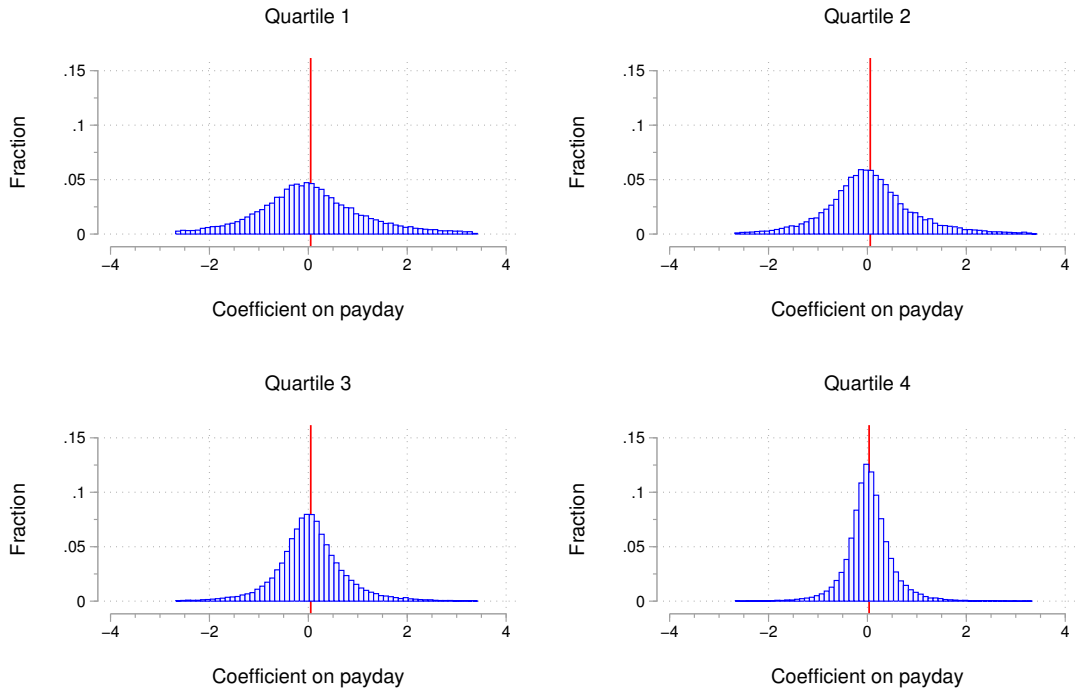


Figure 8: Payday Responses as a Function of Average Pay Cycle Expenditures

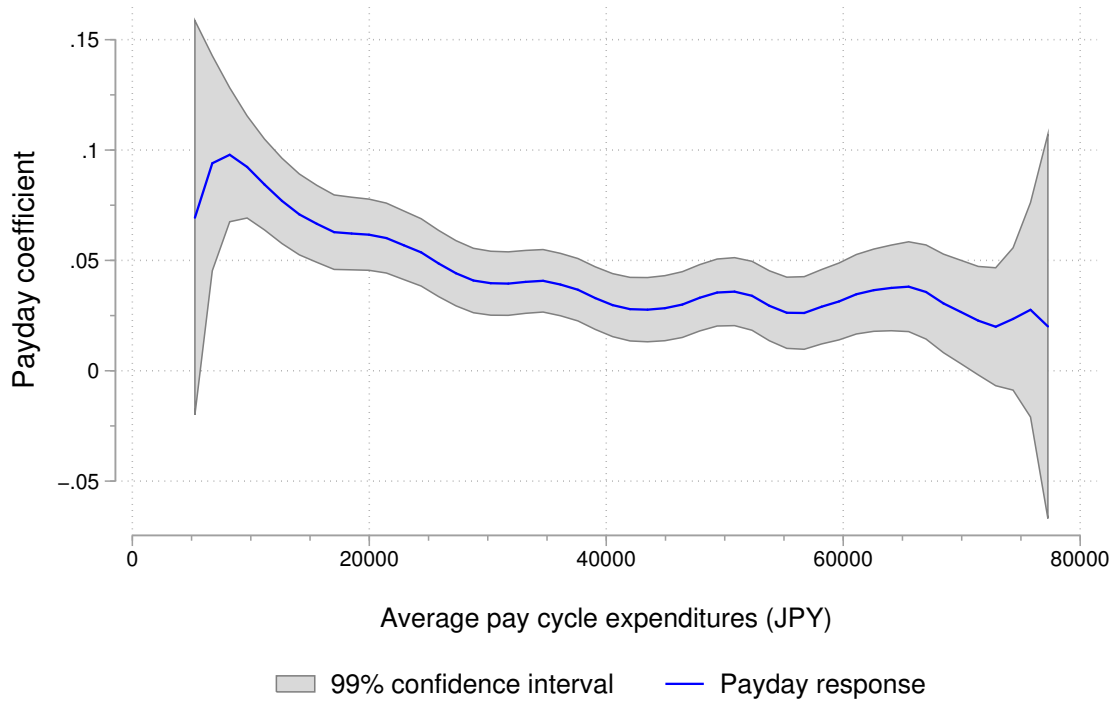


FIGURE 9.A: Store-level Average Prices around Payday

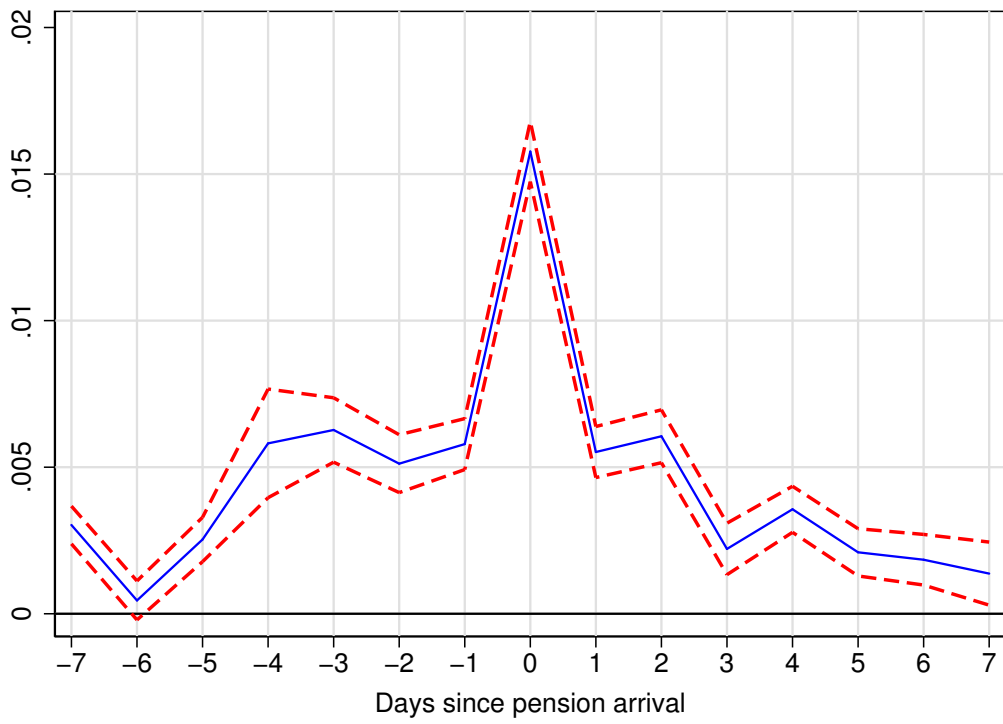


FIGURE 9.B: Number of Unique Goods Purchased around Payday




FIGURE 10: Illustration of Actual and Counterfactual Prices for a Single Good

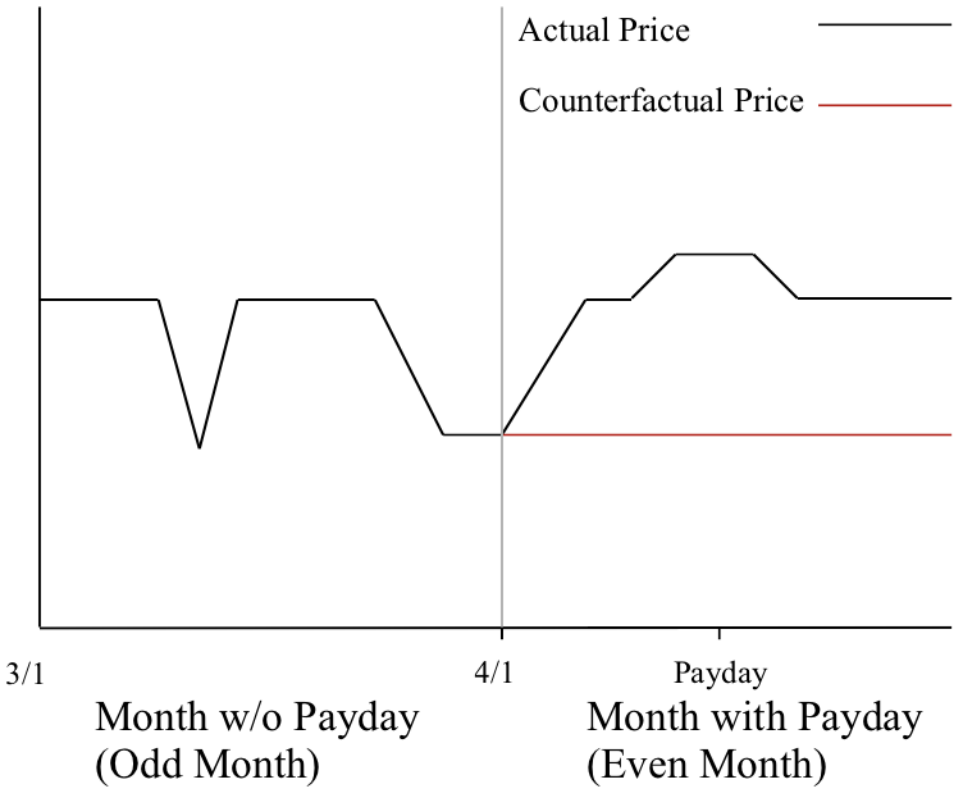


FIGURE 11: Retailer Pricing Response around Payday




FIGURE 12: Municipal Government Expenditures on the Elderly

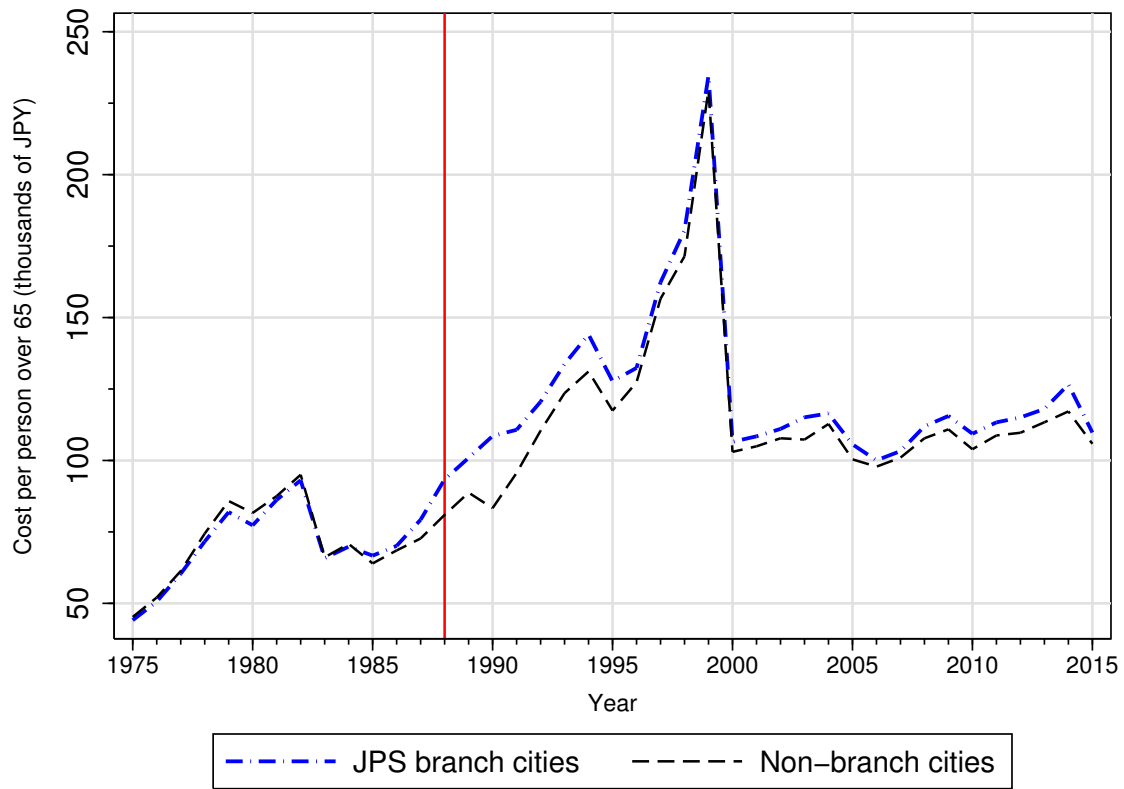


FIGURE 13A. QH Discounting: Optimal Frequency by Inverse IES




FIGURE 13B. Payday Liquidity: Optimal Frequency by Inverse IES




FIGURE 14A. QH Discounting: Optimal Frequency by Convexity of Admin Costs

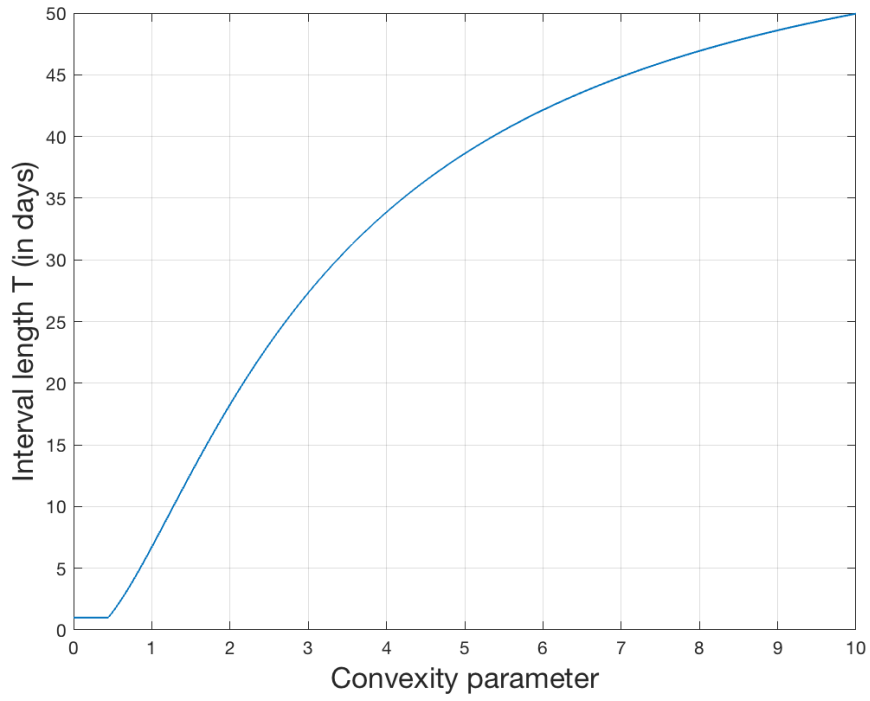


FIGURE 14B. Payday Liquidity: Optimal Frequency by Convexity of Admin Costs




TABLE 1. Retail Expenditures Summary Statistics

	All Goods	Raw Foods
Avg. monthly expenditures	23,102 (15,752)	10,127 (7,260)
Avg. expenditures per trip	2,320 (1,180)	1,086 (579)
Avg. number of monthly trips	9.1 (6.2)	8.5 (5.7)
Avg. periodicity	2.7 (2.5)	2.9 (2.7)
% Early retirement age	40.6%	42.7%
% Normal retirement age	27.9%	29.3%
# Stores	511	511
# Shoppers	1,035,431	831,662

Notes: The table reports separate summary statistics for purchases on all goods in our retail panel and purchases on raw foods, which are defined as all perishable food items (fruits and vegetables, fish, meat and poultry). Monthly and per-trip expenditures are reported in nominal Japanese yen (100 JPY \approx 1 USD). Periodicity refers to the number of days between consecutive shopping trips. % Early retirement age refers to the percentage of shoppers who had reached the earliest possible age (60 years old) for claiming benefits as of the beginning of the panel, and % Normal retirement age refers to the percentage of shoppers who had reached the claimant age at which benefit amounts are no longer taxed for early withdrawal (65 years old). All statistics were computed from a sample of shoppers aged 20 to 90 years old who visited a store in our sample at least twice per month.

TABLE 2. Intensive and Extensive Margin Spending Responses

Category	Intensive	Extensive
All goods	0.134*** (0.002)	0.008*** (0.000)
Prepared foods	0.154*** (0.004)	0.003*** (0.000)
Sweets/desserts	0.136*** (0.005)	0.008*** (0.000)
Alcohol	0.125*** (0.011)	0.005*** (0.001)
Fresh produce	0.081*** (0.003)	0.004*** (0.000)
Fresh fish	0.129*** (0.005)	-0.001* (0.001)
Meat & poultry	0.127*** (0.004)	0.004*** (0.000)
Processed fruits/vegetables	0.007 (0.026)	0.000 (0.002)
Preserved fish	0.042*** (0.007)	0.002*** (0.001)
Grains	0.070*** (0.006)	0.002*** (0.000)

Notes: The table reports separate results for intensive and extensive margin versions of the baseline regression model. The dependent variable in the intensive margin regressions is expenditures on a store visit relative to average daily expenditures. The dependent variable in the extensive margin regressions is a dummy for whether the shopper makes a purchase on a given date. Robust standard errors are clustered at the individual shopper level in parentheses. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

TABLE 3. Evidence for Pent-up Demand: $x'(T) > 0$

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Payment</i> \times <i>Length</i>	0.00151***	0.00176***	0.00226***	0.00253***	-0.733***	-0.736***
<i>Payment</i> \times <i>Length</i> ²			-0.0000126	-0.0000128	0.0246***	0.0247***
<i>Payment</i> \times <i>Length</i> ³					-0.000205***	-0.000206***
<i>Length</i>		-0.000265***		-0.000265**		-0.000267***
Total response ($T = 58$)	0.088	0.087	0.131	0.131	0.242	0.194
Total response ($T = 60$)	0.091	0.090	0.136	0.136	0.300	0.248
Total response ($T = 62$)	0.094	0.093	0.140	0.140	0.259	0.203
Joint F-test (p-value)	–	–	0.000	0.000	0.000	0.000
N	956,844,278	956,844,278	956,844,278	956,844,278	956,844,278	956,844,278
Adj. R ²	0.013	0.013	0.013	0.013	0.013	0.013

Notes: The table reports results of regressions interacting a dummy for pension receipt with the number of days since the last pension payment, with either quadratic or cubic interaction terms. The dependent variable is expenditures on raw foods relative to average daily expenditures on raw foods. We compute the implied total response of expenditures for different interval lengths using the fitted values from each model, conditional on calendar and individual fixed effects. The table also reports the p-values from a joint F-test for significance of coefficients on the interaction terms. In some specifications we include the interval length as a separate regressor to allow for a spillover effect of interval length to shoppers who themselves are not eligible for pension receipt. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

TABLE 4. Expenditure Responses by Permanent Income Decile

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Intensive margin	0.106*** (0.011)	0.097*** (0.009)	0.088*** (0.008)	0.074*** (0.007)	0.069*** (0.006)	0.054*** (0.006)	0.052*** (0.005)	0.048*** (0.004)	0.035*** (0.003)	0.031*** (0.003)
Intensive margin w/controls	0.0104*** (0.011)	0.094*** (0.009)	0.085*** (0.008)	0.073*** (0.007)	0.067*** (0.006)	0.051*** (0.006)	0.050*** (0.005)	0.046*** (0.004)	0.033*** (0.003)	0.029*** (0.003)
Total response	0.069*** (0.005)	0.062*** (0.005)	0.047*** (0.005)	0.045*** (0.005)	0.036*** (0.004)	0.035*** (0.004)	0.026*** (0.004)	0.029*** (0.004)	0.022*** (0.003)	0.003 (0.003)
Extensive margin	0.008*** (0.001)	0.007*** (0.001)	0.003* (0.001)	0.003** (0.001)	0.002 (0.001)	0.001 (0.001)	-0.002 (0.001)	-0.002 (0.001)	-0.002 (0.001)	-0.012*** (0.001)

Notes: The table reports expenditure responses of regular shoppers of age 65 or older to payday, sorted by decile of a proxy for permanent income based on expenditures. Each column number refers to results for a decile of average total expenditures over bimonthly pay cycles. Each row refers to a version of our baseline specification in equation (17) with a different measure of expenditures as the outcome variable. Intensive margin regressions in row 1 condition on the individual visiting the store on a given date. Regressions in row 2 record intensive margin responses but control for the number of periods in between shopping trips. For total response regressions, the dependent variable is expenditures on raw foods relative to average daily expenditures on raw foods. Extensive margin refers to regressions where the outcome variable is a dummy equal to one if the shopper visits a store in our sample on a given date. Hence, coefficients reported for the extensive margin regressions represent changes in probability that a shopper visits a store as a result of payday. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

TABLE 5. Effect of Pension Frequency Reform on Municipal Admin Costs

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Branch</i> × <i>Post</i>	0.043** (0.017)	-0.012 (0.014)	0.034 (0.023)	-0.003 (0.016)	0.073*** (0.026)	0.021 (0.017)
City & year FEs	✓	✓	✓	✓	✓	✓
Incl. Tokyo	✓		✓		✓	
Incl. major cities	✓		✓		✓	
1985 population bin X year FEs			✓	✓	✓	✓
1985 per capita income bin X year FEs					✓	✓
N	11,111	10,635	11,111	10,635	11,111	10,635
# Municipalities	663	635	663	635	663	635
Adj. R ²	0.517	0.554	0.856	0.863	0.863	0.866

Notes: The dependent variable in each regression is log expenses on elderly welfare per resident at or above age 65. $Branch_j = 1$ if municipality j contains a Japan Pension System branch office. $Post_t = 1$ for years 1988–1996. All regressions include observations for years 1980 – 1996 and a full set of year fixed effects. Robust standard errors clustered at the municipality level. Tokyo consists of the 23 central wards for which separate expenditure time series are available. Major cities consist of the historically five most populous cities outside of Tokyo: Yokohama, Nagoya, Kyoto, Osaka, and Kobe. 1985 population bin refers to quintiles of 1985 Census population. 1985 per capita income bin refers to quintiles of per taxpayer taxable income in 1985. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

APPENDIX A: MODEL DERIVATIONS

A.1 GENERAL EXPRESSION FOR THE WELFARE LOSS

We can rearrange the compensating variation equation in (5) to obtain an intuitive expression for the welfare loss:

$$1 - \lambda(T) = 1 - \frac{T \cdot u^{-1} \left\{ \frac{1}{T} \int_0^T u \left(c_0 \cdot \exp(-f(t)) \right) dt \right\}}{\int_0^T c_0 \cdot \exp(-f(t)) dt} \quad (22)$$

where $c_0 \equiv \exp(\theta)$ is the value of consumption on payday, as defined in (4). Note that in general c_0 is a function of T , but we suppress this dependency for ease of exposition. The numerator in this expression represents the amount of total consumption in a counterfactual scenario where daily consumption is such that the pensioner receives the average daily utility over the actual consumption path. The denominator is total consumption over the pay cycle. Let us denote empirically observed total consumption by C^{tot} . Then we can compactly write the welfare loss as:

$$1 - \lambda(T) = 1 - \frac{T \times \bar{C}}{C^{tot}} = \frac{C^{tot} - T \times \bar{C}}{C^{tot}} \quad (23)$$

We can now examine how the welfare loss varies with the government's choice of the interval T between payments. By the quotient rule we can write the marginal compensating variation as

$$\lambda'(T) = \frac{(T \times \bar{C})' \cdot C^{tot} - (T \times \bar{C}) \cdot (C^{tot})'}{(C^{tot})^2} \quad (24)$$

where the derivatives of \bar{C} and C^{tot} with respect to T are

$$\begin{aligned} (\bar{C})' &= \frac{\partial}{\partial T} u^{-1} \left\{ \underbrace{\frac{1}{T} \int_0^T u \left(c_0 \cdot \exp(-f(t)) \right) dt}_{\equiv \bar{U}(T)} \right\} \\ &= \frac{\partial}{\partial T} u^{-1} \left(\bar{U}(T) \right) \cdot \frac{\partial \bar{U}(T)}{\partial T} = u^{-1} \left(\bar{U}(T) \right) (u^{-1})' \left(u \left(\bar{U}(T) \right) \right) \cdot \frac{\partial \bar{U}(T)}{\partial T} \\ &= u^{-1} \left(\bar{U}(T) \right) (u^{-1})' \left(u \left(\bar{U}(T) \right) \right) \cdot \frac{1}{T} \left[u \left(c_0 \cdot \exp(-f(T)) \right) - \frac{1}{T} \int_0^T u \left(c_0 \cdot \exp(-f(t)) \right) dt \right] \end{aligned} \quad (25)$$

$$(C^{tot})' = \frac{\partial}{\partial T} \int_0^T c_0 \cdot \exp(-f(t)) = c_0 \cdot \exp(-f(T)) = \bar{B} \quad (26)$$

Thus combining (19)-(21) we can write the FOC to the government's problem in (7) as

$$\frac{p(\bar{C} \times C^{tot}) + p(T \times (\bar{C})' \cdot C^{tot}) - p(T \times \bar{C}) \cdot c_0 \cdot \exp(-f(T))}{(C^{tot})^2} = \mu'(T) + p \cdot \bar{B} \quad (27)$$

We have now obtained an expression that defines the optimal payment frequency T^* as an implicit function of model parameters and the administrative cost function $\mu(T)$. For each case of our model, we solve for (22) using standard algorithms for solving nonlinear equations.

A.2 DETAILS FOR QUASI-HYPERBOLIC DISCOUNTING CASE

For QH discounters with log utility, it is straightforward to compute a closed-form expression for the welfare loss:

$$1 - \lambda(T) = \frac{c_0}{\bar{B}} \cdot \exp\left(-\nu \cdot T/2\right)$$

For isoelastic utility functions with $\rho \neq 1$, we can define the welfare loss in terms of \bar{C} and C^{tot} :

$$\begin{aligned} \bar{C} &= u^{-1} \left\{ \left(\frac{c_0^{1-\rho}}{\nu \cdot T(1-\rho)^2} \right) \left[1 - \exp\left((\rho-1)\nu \cdot T\right) \right] \right\} \\ C^{tot} &= \frac{c_0}{\nu} \cdot \left(1 - \exp(-\nu \cdot T) \right) = \bar{B} \cdot T \end{aligned}$$

Putting everything together, we get the welfare loss for QH discounters with isoelastic utility and inverse IES ρ :

$$1 - \lambda(T) = \begin{cases} 1 - \frac{1}{\bar{B}} \cdot \exp\left(c_0 - \nu T/2\right) & \text{if } \rho = 1 \\ 1 - \frac{c_0}{\bar{B}} \cdot \left[\frac{1 - \exp\left((\rho-1)\nu T\right)}{\nu T(1-\rho)} \right]^{\frac{1}{1-\rho}} & \text{if } \rho \neq 1 \end{cases}$$

The marginal welfare loss is therefore:

$$-\lambda'(T) = \begin{cases} \frac{c_0 \nu}{2\bar{B}} \cdot \exp(-\nu T/2) & \text{if } \rho = 1 \\ \frac{c_0}{\bar{B}\nu(1-\rho)^2 T^2} \left[\frac{\exp(\nu(\rho-1)T) - 1}{\nu(\rho-1)T} \right]^{\frac{\rho}{1-\rho}} \frac{\left(\exp(\nu(\rho-1)T)(\nu(\rho-1)T - 1) + 1 \right)}{1 - \frac{c_0}{\bar{B}} \exp(-\nu T/2)} & \text{if } \rho \neq 1 \end{cases}$$

A.3 DETAILS FOR PAYDAY LIQUIDITY CASE

Expression (15) shows the welfare loss when consumers exhibit payday liquidity with a spike on payday that is independent of the interval length T . When we allow for the spike to depend on interval length, the welfare loss expression is the same, except we write $x(T)$. For the log utility case, it is straightforward to compute the marginal welfare loss:

$$-\lambda'(T) = \frac{(\bar{c}/\bar{B})(1+x(T))^{1/T-1}}{T^2} \left[(1+x(T)) \cdot \log(1+x(T)) - T x'(T) \right]$$

From this equation, we see that irrespective of administrative costs, for this type of payday liquid consumer the government faces a tradeoff between lowering the spike magnitude and decreasing subdivision of the welfare loss when deciding to shorten the pay cycle.

For the case where $\rho \neq 1$, we can use the general formula for the welfare loss in (18) to compute the marginal welfare loss. In this context \bar{C} and C^{tot} are given by:

$$\bar{C} = \frac{\bar{c}}{T^{1/(1-\rho)}} \left[\left(1 + x(T)\right)^{1-\rho} + (T - 1) \right]^{\frac{1}{1-\rho}}$$

$$C^{tot} = \bar{B} + \frac{\bar{B} \cdot T}{T + x(T)} = \bar{B} \cdot T$$

Using these expressions, we obtain the marginal welfare loss when $\rho \neq 1$:

$$-\lambda'(T) = \frac{\bar{c}}{\bar{B}} \cdot \frac{\left((1+x)^\rho - (1+x) \right) T^{\frac{2-\rho}{\rho-1}} \left((1+x)^{1-\rho} + x - 1 \right)^{\frac{1}{1-\rho}}}{(\rho-1) \left(T(1+x)^\rho - (1+x)^\rho + (1+x) \right)}$$

APPENDIX B: CONSTRUCTING GOODS SUBCATEGORIES

Our data contain 968 major goods categories provided by Magee Co., with some categories containing thousands of barcodes each corresponding to a unique product. The majority of these categories consist of perishable food items. However, while most of the goods in our data would be considered non-durable by standard definitions proposed in the empirical literature on consumption patterns, we find it useful to further aggregate these categories based on whether they consist of fresh foods or packaged foods, alcohol, tobacco, or other non-discretionary spending. Relative to total expenditures, these subcategories provide information on how diet composition varies over the pay cycle, and information on which goods are more likely to be purchased on paydays.

In the end, we consider 13 subcategories, plus a category called “raw foods” which contains all fresh, non-packaged food items in our sample. We summarize the contents of these subcategories in Table B1. Tables B2 through B4 provide summary statistics for each subcategory. Further summary of the goods included in these subcategories can be found in the Online Data Appendix.

Table B.1: One-digit Goods Subcategory Contents

One-digit Category	Two-digit Category	Four-digit Categories
Fresh fruits & vegetables	Fresh fruits	seasonal fruits, imported fruits, assorted fruits, fruit-related products,
	Fresh vegetables	leafy veg., stalk veg., root crops, edible plants, edible seeds, mushrooms, germinated veg., assorted veg.
Processed fruits & vegetables	Processed fruits	frozen fruits, cut fruits,
	Processed vegetables	boiled veg., frozen veg., cut veg.
Fresh fish	Fresh fish	round items, filet, shellfish, assorted fish
	Sashimi	brick form, sashimi, tataki, raw fish, assorted fresh fish
Preserved fish products	Salted & dried fish	boiled fish, frozen fish, seasoned fish, pickled fish, salted fish, dried fish, fish eggs, seaweed
Raw meat & poultry	Beef	wagyu, domestic beef, imported beef,
	Pork	domestic pork, imported pork
	Chicken	domestic chicken, imported chicken, brand name chicken, duck meat
	Meat varieties	lamb, horse meat, minced meat, offal, raw meat, eggs, dairy products
Grains	Cereals	powder, rice, mochi, raw noodles, dough, bread, cereal
Other processed foods	Seasonings	cooking oil, spices, condiments, spread/dips, toppings, rice seasoning
	Dry produce	dried fish, dried fruits
	Processed food	pickled items, processed fish, pastes, cooked beans, processed meats
	Instant foods	cup noodle, instant soup, frozen foods, sealed rice pouch
Prepared foods	Semi-prepared dishes	fried, simmered, grilled, Japanese, Western, Chinese
	Side dishes	fried, grilled, grilled eel, Japanese, Western, Chinese
	Bento	cooked rice, sushi, bread dishes, noodle dishes
Sweets and desserts	Confectionary	toppings, jelly/pudding, ice cream, frozen confections, candies/cookies, rice crackers
Non-alcoholic beverages	Beverages	coffee/tea, milk-based drinks, vegetable/fruit drinks, soft drinks
Alcohol	Alcohol	beer, liquers, wine liquor, sake
Tobacco	Tobacco	tobacco
Other discretionary	Other	flowers, gifts/confections, kiosk goods, service counter goods

Table B.2: Summary Statistics for Fresh Foods Expenditures

	Fruits/Vegetables	Fish	Meat/Poultry	Prepared Foods	Raw Foods
Avg. monthly expenditures	4,476	3,746	4,554	4,567	10,127
% Normal retirement age	30.2%	37.2%	26.4%	29.6%	29.3%
Avg. number of monthly trips	7.9	6.1	7.0	7.2	8.5
Avg. periodicity	3.1	3.8	3.4	3.1	2.9
# Stores	510	495	508	509	511
# Shoppers	659,402	165,408	527,578	420,497	831,662

Table B.3: Summary Statistics for Processed Foods Expenditures

	Fruits/Vegetables	Fish	Grains	Other Processed
Avg. monthly expenditures	981	2,203	2,734	5,661
% Normal retirement age	25.0%	36.0%	26.8%	28.8%
Avg. number of monthly trips	4.2	5.5	6.8	8.0
Avg. periodicity	4.7	4.3	3.5	3.0
# Stores	503	505	509	510
# Shoppers	30,952	127,599	417,280	782,680

Table B.4: Summary Statistics for Discretionary Expenditures

	Sweets/Desserts	Non-alcoholic Beverages	Alcohol	Tobacco	Other
Avg. monthly expenditures	2,689	2,745	5,745	6,510	2,055
% Normal retirement age	24.5%	23.9%	20.0%	24.0%	32.4%
Avg. number of monthly trips	6.8	6.9	6.7	5.4	5.1
Avg. periodicity	3.5	3.4	3.0	3.1	3.9
# Stores	510	510	508	391	475
# Shoppers	379,579	430,856	85,084	8,649	35,474

APPENDIX C: USING GOOGLE DATA TO TEST FOR INATTENTION

This appendix outlines how we constructed daily time series of Google searches related to payments from the Japanese National Pension System. We use these series to gauge the extent to which the timing of upticks in search activity line up with scheduled payment dates. This exercise is useful for two reasons: (i) it allows us to assess whether shoppers anticipate pension payments, and (ii) it allows us to discern whether shoppers are inattentive to changes in the payment schedule due to holidays and weekends. In a recent paper, [Baugh & Wang \(2018\)](#) argue that consumers may be unaware of or ignore rules governing regular Social Security payment schedules, even when such rules are public information. We find limited evidence in the Google data to support this hypothesis in the context of the Japanese public pension system.

C.1 CONSTRUCTING DAILY GOOGLE SEARCH DATA

When querying Google Trends, the user provides search terms, a region code (e.g. “JP” for Japan), and a time range. As of writing, Google Trends data are available from January 1, 2004.³⁰ It is possible to add punctuation to queries to filter the results. For instance, one can wrap double quotation marks around keywords to restrict to searches containing the exact phrase. One can also use the addition operator to generate intersected queries. In our application, we use the minus operator to search for “public pension payments - fees,” which excludes searches related to contributions from participants who have not yet begun to claim benefits.³¹

Instead of reporting the total amount of searches for the user-provided keywords, Trends reports a search volume index (SVI). For example, if one typed “pension” into the Google Trends search bar, for each t in time range τ and region r one would then obtain:

$$\text{pension search index}_{r,t} = 100 \times \frac{\text{SP}_{r,t}}{\max_r \text{SP}_{r,t}}$$

where the search propensity (SP) is computed as:

$$\text{SP}_{r,t} = \frac{\left[\text{Google searches containing “pension”} \right]_{r,t}}{\left[\text{total Google searches} \right]_{r,t}}$$

The SVI is a function of the aggregate search volume Google receives. This raises the issue that some of what the SVI measure captures is trends in total Google searches over time. We check that this denominator effect is not confounding our results by obtaining total page views from the

³⁰Note that Google frequently changes the Trends interface, so some procedures discussed here may not be applicable in the future. The current web browser version of the tool can be accessed at <http://trends.google.com/trends/explore>

³¹Our results are qualitatively similar for non-interacted keyword searches of “public pension payments.”

Wikipedia page for the Japanese public pension system, rescaling the page views series as Google does, and then comparing results using the scaled vs. unscaled page views.

Google Trends, by default, returns a monthly series for time ranges longer than 5 years, a weekly series for time ranges shorter than 5 years but longer than 8 months, and a daily series for time ranges spanning less than 8 months. Hence, to obtain SVI data at the daily frequency for a long time period, the user must generate separate queries for several subperiods. Since SVI observations in each subperiod are constructed relative to the week with the greatest search volume, the user must create a common index across the subperiods to maintain the interpretation of the SVI as a measure of relative search propensity over the entire sample period.

For instance, to create a weekly series we perform the following steps:

1. For each keyword search, download the full monthly series for 2004-2018.
2. Divide the period into subperiods τ_k of equal length and download the weekly data for each subperiod. Merge the subperiods together into a single data frame.
3. Create a weekly data frame that assigns to each week the SVI observation for the corresponding month.
4. Create another weekly data frame that assigns to each week j the SVI observation in the first week j' of the corresponding month i .
5. Rescale the downloaded weekly SVI series via:

$$Z_{i,j}^* \equiv \underbrace{\frac{\text{SVI}_{i,\tau}}{\text{SVI}_{j',\tau_k}}}_{\text{correction factor}} \times \underbrace{\text{SVI}_{i,j,\tau_k}}_{\text{downloaded data}}$$

For cases where the first week of month i has an SVI of zero, normalize $Z_{i,j}^* = 0$.

6. Create the new series $SVI_{i,j,\tau}^*$ from the data frame of scale factors Z^* :


$$SVI_{i,j,\tau}^* \equiv \text{round} \left(\frac{Z_{i,j}^*}{\max_{(i,j)} \{Z^*\}} \times 100 \right)$$

The procedure is similar for obtaining daily-frequency data over the entire available period. We first create a rescaled weekly SVI series using the steps outlined above, and then use this rescaled weekly series to construct a correction factor that creates a common index across all of the subperiods of daily observations. Our daily time series of Google search activity is therefore an “index within an index.”

Figure C.1 plots the resulting raw, daily SVI series for searches of “public pension payment,” excluding searches containing “fees,” for our sample time period of April 2011 to October 2014. The dashed red lines indicate scheduled payment dates. Even in the raw data, there is a clear uptick

in search activity on payment dates, even when a date is rescheduled from the default 15th of the month. There are also clear spikes in search activity around announcements of reforms to JPS that would affect the determination of benefits. We control for policy announcement effects to isolate the response of searches to the payment schedule in the next subsection.

Figure C.1: Raw Daily Google Searches for “Public Pension Payments”



C.2 ESTIMATION PROCEDURE

We run time series regressions of the following form using Google SVI for “public pension payments” relative to average daily SVI as the outcome variable:

$$y_t = \sum_{j=-7}^{+7} \beta_j \cdot \text{Payday}_{t+j} + \gamma \cdot t + \delta_{dow} + \phi_{wom} + \psi_{my} + \xi_h + \alpha_p + \epsilon_t$$

where we control for the full set of time fixed effects capturing cyclicity in searches within the month (ϕ_{wom}), within the week (δ_{dow}), and within the year (ψ_{my}), as well as holiday effects ξ_h . The linear time trend $\gamma \cdot t$ accounts for the secular increase in internet search activity over the time period. The dummy α_p equals unity on dates when the government first announced policy changes to the pension system, such as changes to the formula determining benefits and the normal retirement

age that would affect certain cohorts. Such policy changes are not directly related to the frequency or delivery schedule for payments. However, failing to account for policy announcements could lead to a spurious correlation between payment dates and search activity, as several announcements are timed close to scheduled delivery dates.³²

Figure C.2 plots the estimated coefficients $\hat{\beta}_j$ which capture search behavior around scheduled pension paydays. Panel A does this for the time series (April 2011 to October 2014) that overlaps with our retail panel, and Panel B extends the time series to the full available range of Google data (January 2004 to December 2018). In both cases, there is a clear spike in search behavior directly prior to payday, and search activity declines thereafter. This response is equal to 21% above the average daily level of searches on the day before payday in the shorter time series, and 5% above average in the full time series. In the full sample, the decline in searches is sharper, with search volume returning to the average level within two days after a payday. These findings support the notion that individuals are highly attentive to even small deviations of delivery dates from the benchmark 15th of the month.

³²We obtained a list of pension program announcements from the official JPS website: <http://www.nenkin.go.jp/oshirase/taisetu/index.html>. From this list of 58 announcements during our sample time period, we also tried excluding announcements that were purely clerical in nature (e.g. new application form for dependent beneficiaries), which left us with only three substantive policy change dates. In either case, the estimated coefficient on the dummy α_p was insignificant.

FIGURE C.2A: Sample Time Period: Payday Google Searches

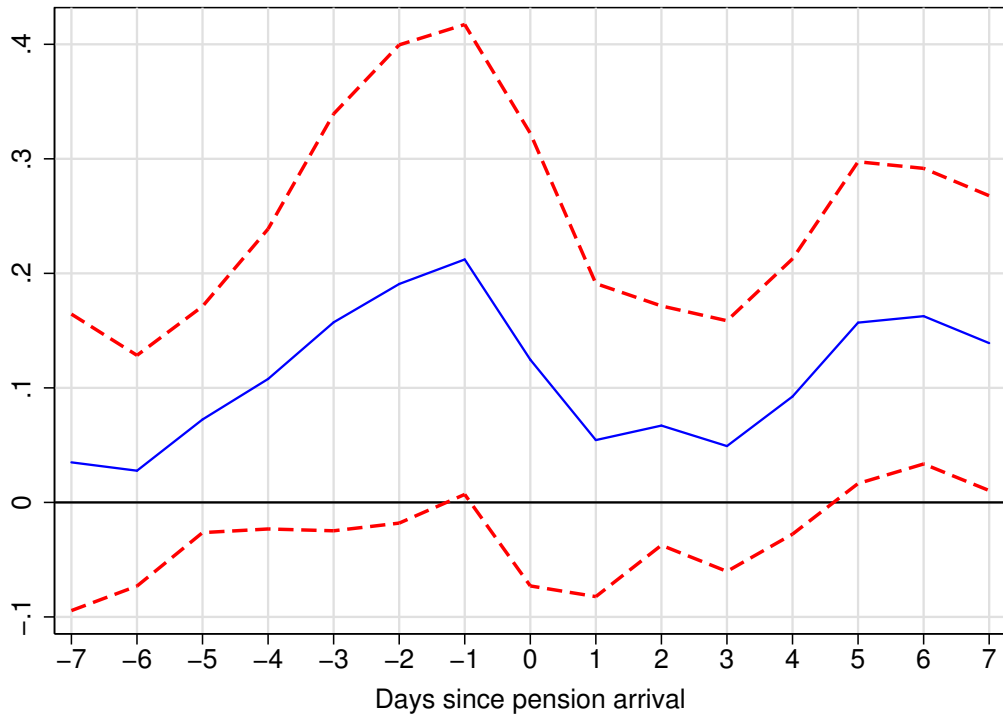


FIGURE C.2B: Full Time Period: Payday Google Searches

