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

Cameron LaPoint[†] Shogo Sakabe[‡]

Yale SOM Columbia

October 2022

Abstract

How governments should choose the frequency of payments has received little attention in the literature on the optimal design of benefits programs. We propose a simple model in which the government chooses the interval length between payments, subject to a tradeoff between costs of providing more frequent benefits and welfare gains from mitigating consumption non-smoothing. Using a high-frequency retail dataset that links consumers to their purchase history, we apply the model to the Japanese National Pension System. Our evidence suggests suboptimal intra-cycle consumption patterns, with negligible retailer price discrimination. Model calibrations support the worldwide prevalence of monthly payment systems.

Keywords: optimal payment frequency, splurge goods, consumption smoothing, present bias, incidence, retail scanner data, store-level prices, pension payments

JEL classifications: D12, D91, G51, H21, H55, I38

*We would like to thank Michael Best, Judy Chevalier, Wojciech Kopczuk, Kazuki Onji, Michaela Pagel, Daniel Reck, Jón Steinsson, Neil Thakral, Takashi Unayama, Tsutomu Watanabe, David Weinstein, and participants in the Applied Micro Methods Colloquium at Columbia and Japan Empirical Economics Seminar at Osaka University for helpful advice and comments. We thank Louisa Liu for excellent research assistance. We are grateful to Magee Co., Ltd. (formerly, IDs 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 supporting this project through the provision of dissertation fellowships to both authors. Additionally, Sakabe thanks the Nakajima Foundation for generous financial support. An earlier version of this paper circulated as CJEB Discussion Paper, No. 370.

[†]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; Web: http://shogosakabe.github.io

1 Introduction

Most 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 transfer programs is how the government should choose the frequency of disbursal to eligible households. Several studies have found evidence that infrequently distributing benefits to people who live "paycheck to paycheck" 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 recipients without savings consume fewer calories and switch to unhealthy foods near the end of the pay cycle.

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 welfare gain from reducing deviations from full consumption-smoothing behavior.

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 1,000 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 shopper ID and the birth month/year of the registered shopper. 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 administrative burden imposed by 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.

We use our 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. Moreover, even for extreme assumptions about consumer preferences, such as high intertemporal elasticity of substitution, or programs featuring highly convex administrative cost functions with respect to interval length, our model rarely generates optimal frequencies exceeding one month. Hence, the model supports the worldwide prevalence of monthly public payment schedules.[1]

Using the Japanese old-age pension system as our 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

[1]We hand collected information related to the administration of public benefits for 36 OECD countries. Of these, 25 disburse universal public pensions on monthly schedules, and eight distribute payments at even shorter intervals. Other forms of public benefits such as means-tested welfare programs are also commonly distributed on a monthly basis.

reduces the need of retirees to save their pension income for uncertain medical expenses. Third, as Stephens & Unayama (2011) document, 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 posited that policymakers might improve welfare by increasing the frequency of benefit payments. Parsons & Van Wesep (2013) and Baugh & Correia (2022) build on this notion in studying how pay cycle duration influences consumption and borrowing patterns. We complement these papers by offering a sufficient statistics approach to computing optimal pay frequency from a regulator's point of view. Dobkin & Puller (2007) argue that recipients of SSI drive spikes in drug-related hospital admissions at the beginning of the month, while others show that crime increases (Foley 2011) and mortality falls (Evans & Moore 2012) towards the end of monthly welfare cycles. The link between physical harm and non-smooth consumption paths at higher frequencies suggests estimates of the utility loss from permanent income hypothesis (PIH) deviations, as in Cochrane (1989) and Browning & Crossley (2001), may be severely underestimated.

Our paper speaks to the vast literature testing the permanent income hypothesis in the context of the consumption response to *regular* income sources – that is, recurring payments, such as pension payments, that are delivered more than once per year.[2] Stephens & Unayama (2011) look at the Japanese pension system during an earlier period in the 1980s when pension payments were distributed quarterly. Consistent with our results, they provide 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 empirical research design to ours, Baugh & Wang (2021)

[2]The vast majority of these studies use quarterly or monthly panel data and examine payments that occur at the cutoff separating two discrete time observations (e.g. Browning & Collado 2001; Stephens 2003,2006). 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.

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 those authors argue that inattention to changes in the length of the pay cycle can explain their results, we provide evidence from Google search data that consumers in our setting are highly attentive to even small changes to usual disbursement dates.

Our analysis of expenditure patterns of the elderly also contributes to the literature on the retirement consumption puzzle. Bernheim, Skinner, & Weinberg (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 eligibility in Italy to estimate a 9.8% drop in consumption due to male retirement. 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. Shoppers who cross retirement age thresholds in our panel do not reduce spending on perishables, indicating savings stocks are likely to be high among benefit recipients.

While we cannot directly test for liquidity constraints among elderly consumers in our sample, we document similar payday responses for shoppers who visit stores at different average shopping frequencies. Coibion, Gorodnichenko, & Koustas (2021) show that the ability of upper-income households to buy in bulk helps explain increasing U.S. expenditure inequality. Based on the insight of Bils & Klenow (2001) that quality Engel curves slope upward, we also use prices paid within a goods category as a tag for income and find similar spending patterns when we split shoppers along that dimension. Our results therefore imply at least part of the response of expenditures to payday is due to factors other than short-term liquidity constraints.

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 inhibit 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, Kashyap, & Rossi (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, Homonoff, & Meckel (2022) 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 capture the incidence of pension payments by price discriminating against elderly customers on or around scheduled delivery dates. Average prices paid within a store increase by 1.6% on payday, but the number of unique goods purchased increases by 6%. Using a price index that isolates a change in average prices due to consumer substitution from the retailer's pricing response, we demonstrate that the observed spike in 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 splurge behavior on payday reinforces the theory, introduced in Chevalier & Kashyap (2019), of consumer "type switching" from bargain hunters to brand loyalists during periods of peak demand. The substitution and variety effects we uncover are also consistent with evidence on the non-fungibility of SNAP benefits (Hastings & Shapiro 2018) and tax refunds (Baugh, Ben-David, & Parker 2021), and with the literature on mental accounting (Thaler 1999; Farhi & Gabaix 2020), which argues individuals earmark income sources for specific spending categories. Our finding of within-retailer substitution and variety effects complements the analysis in Baker, Baugh, & Kueng (2021), who show households switch to higher quality retailers when their income increases. Overall, even though our model can easily incorporate strategic retailer responses, 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 our optimal payment frequency model. Section 3 provides background on the Japanese pension system and describes our data. Section 4 presents our main empirical results on the expenditure response to payment receipt. Section 5 assesses possible retailer pricing responses and shows consumer substitution towards higher quality goods. Section 6 discusses our calibration of 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 policy

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 $b(T) = \overline{B} \cdot T$. 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.[3]

In addition to the direct cost of $p \cdot b(T)$ of delivering benefits to retirees, the government faces an administrative cost function $\mu(T)$.[4] 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:

$$(1-p) \cdot \tau(b) = p \cdot b(T) + \mu(T) \implies \tau(b) = \frac{p \cdot \overline{B} \cdot T + \mu(T)}{1-p} \tag{2.1}$$

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

$$\max_{\{C(t)\}} \int_0^T u(C(t))dt \text{ s.t. } C(t) = S(t) + w(t) - \frac{\tau(b)}{T} \tag{2.2}$$

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 interval $[0, T]$. However, we suppose that, following Shapiro (2005), the pensioner

[3] The 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).

[4] These 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.

instead chooses a consumption path given by

$$C(t) = \exp\Big(\theta - f(t)\Big) \tag{2.3}$$

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 as to 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:

$$\int_0^T \exp\Big(\theta - f(t)\Big) = b(T) \implies \exp(\theta) = \frac{b(T) \cdot f'(0) f'(T)}{f'(T) - f'(0) \cdot \exp(-f(T))} \tag{2.4}$$

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 can be obtained as the fraction λ that solves the following equation:

$$\int_0^T u^r\Big(\exp\Big(\theta - f(t)\Big)\Big)dt = \int_0^T u^r(\lambda \overline{B})dt \tag{2.5}$$

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, $\gamma = 1$.[5] Given a constant MCF, the optimal frequency problem a utilitarian government faces is given by

$$\min_T \Big\{ -p \cdot \lambda(T) + \gamma \cdot \Big(p \cdot b(T) + \mu(T)\Big)\Big\} \tag{2.6}$$

where $\lambda(T)$ is the solution to (2.5). For any strictly concave $u^r(\cdot)$, the welfare loss $1 - \lambda$ will

[5] More 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.

be strictly convex in T. Therefore, in such a case the FOC of this problem is necessary and sufficient for a solution:

$$\frac{p \cdot \lambda'(T^*)}{\gamma} = \mu'(T^*) + p \cdot \overline{B} \tag{2.7}$$

We have now obtained T^* as an implicit function of observables. These include the daily average benefit amount among claimants over the pay period $\overline{B}$ 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 maintain an otherwise smooth consumption profile.

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:

$$u(c_0) + \beta \sum_{t=1}^{T} \delta^t u(c_t) \tag{2.8}$$

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.[6] Under these conditions, one can show that the decrease in log consumption over time is:

$$\frac{\partial \log(c_t)}{\partial t} = \frac{1}{\rho} \cdot \log \beta - \frac{1}{T - t + 1} + \frac{1}{T - t + \beta^{-1/\rho}} < 0 \tag{2.9}$$

where ρ is the inverse intertemporal elasticity of substitution (IES). For values of β close to one, the time path for consumption given by (2.9) is approximately linear over the pay cycle.[7] 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

[6] The 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.

[7] Exponential 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 $\log(\delta)/\rho$.

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.

To illustrate, for log utility ($\rho = 1$), the compensating variation term from (2.5) is

$$\lambda(T) = \frac{1}{\overline{B}} \cdot \exp\Big(\theta - \nu \cdot T/2\Big) \tag{2.10}$$

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

$$-\frac{p \cdot \nu}{2\overline{B}} \cdot c_0 \cdot \exp\Big(-\nu \cdot T/2\Big) = \mu'(T) + p \cdot \overline{B} \tag{2.11}$$

We provide a more general expression for the welfare loss under isoelastic utility 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 for Japanese pension payments. 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.[8]

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.

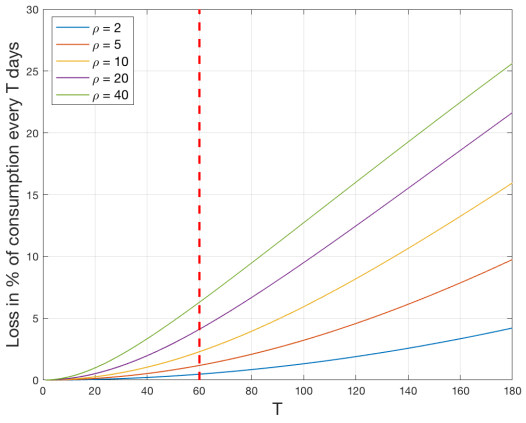
2.3 Payday liquid consumers

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

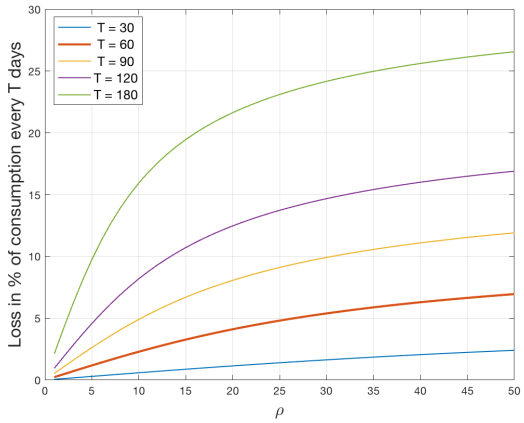
[8]Most non-experimental estimates for the inverse IES fall in the range $\rho \in [1, 3]$. However, Best at al. (2020) use notches in mortgage interest rate schedules in the UK to argue that $\rho = 10$ is plausible.

FIGURE 1. Welfare Loss under Quasi-Hyperbolic Discounting

By Interval Length



By Inverse IES



Notes: The figure plots the welfare loss implied by quasi-hyperbolic discounters as a function of the government's choice of interval length (left), or as a function of the inverse intertemporal elasticity of substitution (right), under isoelastic utility. Full derivations and closed-form expressions for the welfare loss can be found in Appendix A.2.

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 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. 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.

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.[9]

[9]It 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.

We assume pensioners consume over the time interval $[0, T-1]$ as follows:

$$C_t = \begin{cases} (1+x)\cdot \bar{c} & \text{if } t = 0 \\ \bar{c} & \text{if } t \in [1, T-1] \end{cases} \tag{2.12}$$

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

$$\sum_{t=0}^{T-1} C_t = (1+x)\cdot \bar{c} + \sum_{t=1}^{T-1} \bar{c} = (T+x)\cdot \bar{c} = \overline{B} \cdot T$$

$$\implies \bar{c} = \begin{cases} \frac{\overline{B}\cdot T}{T+x} & \text{if } T > 1 \\ \overline{B} & \text{if } T = 1 \end{cases} \tag{2.13}$$

$\bar{c}$ 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.[10]

Many utility functions can produce the consumption path in (2.13) 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 $\bar{c}$.[11] 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 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 (2.5) is:

$$\sum_{t=0}^{T-1} u(C_t) = T \cdot u(\lambda \overline{B}) \tag{2.14}$$

To facilitate comparison to other cases of our model, we continue to assume an isoelastic

[10] While 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.

[11] More 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 $\bar{c}$.

felicity function with inverse IES ρ. The expression for the welfare loss is then:

$$1-\lambda(T)=\begin{cases}1-\frac{\bar{c}}{\bar{B}}\cdot(1+x)^{1/T} & \text{if } \rho=1\\ 1-\frac{\bar{c}}{\bar{B}\cdot T^{1/(1-\rho)}}\Big[(1+x)^{1-\rho}+(T-1)\Big]^{\frac{1}{1-\rho}} & \text{if } \rho\neq 1\end{cases} \tag{2.15}$$

Figure 2 plots the welfare loss formula in (2.15) 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 (Panel A). 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:

$$\underbrace{\Big(1+x(T)\Big)\cdot\log\Big(1+x(T)\Big)}_{\text{loss from spike magnitude as } T\uparrow} \quad > \quad \underbrace{T\cdot x'(T)}_{\text{gain from subdivision as } T\uparrow} \tag{2.16}$$

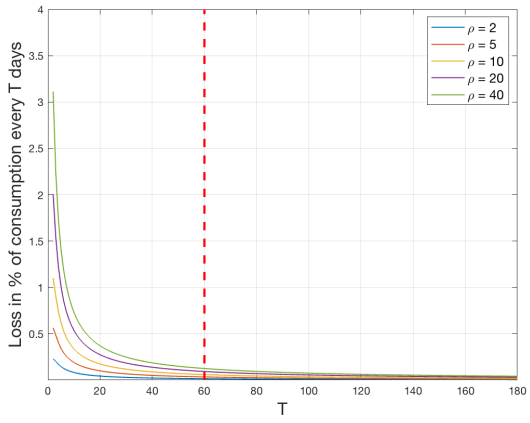
In our empirical setting, we estimate a slope of $x'(T) = 0.0013$ by exploiting exogenous calendar variation in pension delivery dates, so the spike on payday increases by 0.13 percentage points for each additional day in the pay cycle. For this value of $x'(T)$ the condition in equation (2.16) holds for any $\rho > 0$. Panel B of Figure 2 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.

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 $x = 0.1$

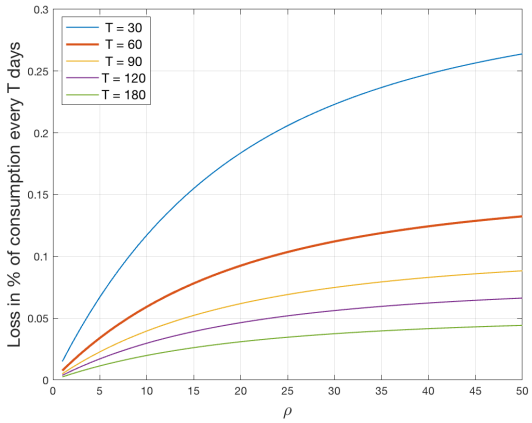
FIGURE 2. Welfare Loss under Payday Liquidity

A. Standard Payday Liquidity (Constant x)

By Interval Length

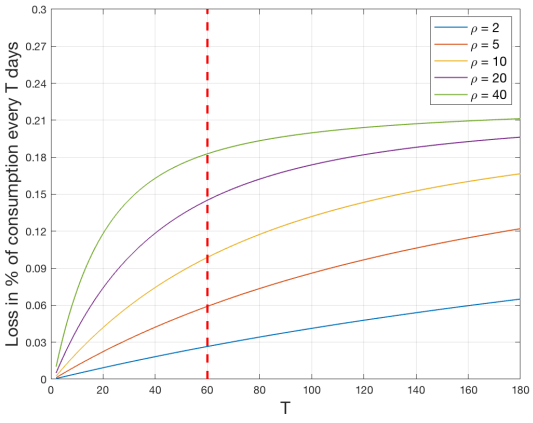


By Inverse IES

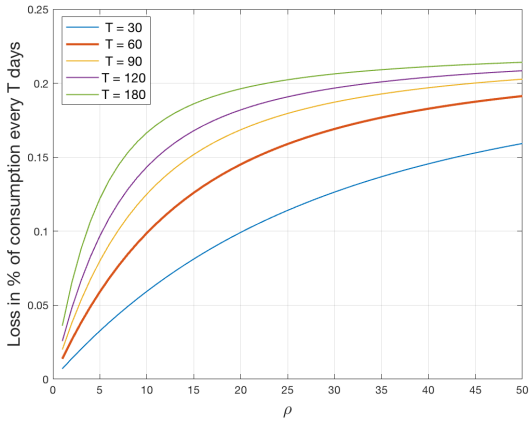


B. Payday Liquidity with Pent-up Demand (x linear in T)

By Interval Length



By Inverse IES



Notes: Panel A of the figure plots the welfare loss in equation (2.15) implied by payday liquidity behavior as a function of the government's choice of interval length, or as a function of the inverse intertemporal elasticity of substitution, under isoelastic utility. Panel B plots the same object when we allow the spike magnitude $x(T)$ to vary with interval length, which generates pent-up demand at the beginning of pay cycles. See Appendix A.3 for detailed derivations.

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 full derivations and model extensions 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 and municipal-level data used 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).[12] 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, a consumer-identifier code, a store-identifier code, and prices and quantities of barcode-level purchases. An unique characteristic of this dataset is that it includes consumer identifiers assigned at the time consumers sign up for a shopper loyalty program. This shopper ID has information about each consumer's birth year/month and gender. If a consumer uses their membership card when shopping, their identifier code is recorded on the sales receipt, which enables us to track their 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

[12]To our knowledge, the only other paper using these data is Shoji (2020), who investigates consumer stockpiling behavior in advance of a consumption tax hike in 2014.

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. The number of shoppers stabilizes at around 500,000 for shoppers who go to the store on at least a weekly basis, on average ($k \geq 4$).[13]

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. Such stores might exit either because the store closes for business or because management chooses to stop providing data to the marketing firm. Second, we restrict to stores that offer loyalty point card programs. Our results are unaffected when we restrict to a balanced panel of stores which record transactions during all non-holiday dates.

To facilitate comparison of our analysis to other papers in the literature and classify products by durability, we aggregate barcodes into 13 subcategories, plus a category called "raw foods" which contains all fresh, non-packaged food items in our sample.[14] Overall, 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, 2011, and 2013 waves of the survey, each of which includes data from interviews with approximately

[13]We obtain qualitatively and quantitatively similar results when we deviate from our baseline sample restriction of $k = 2$. See Appendix B.2 for the full relationship between the number of unique shopper IDs and trip frequency. We also show in Appendix F.1 that our spending response estimates are robust to using frequent shoppers *within* each goods subcategory, rather than fixing a balanced panel across all goods.

[14]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. We describe the contents of these categories in more detail and offer summary statistics in Appendix B.

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, we use questions on typical elderly expenditure patterns in an attempt to quantify the share of total expenditures we capture in our grocery retail data and determine the extent to which we may underestimate the true response of expenditures to pension income receipt. Third, we use detailed questions related to wealth and income to assess whether elderly households are liquidity constrained. We report summary statistics from the JSTAR data in Appendix B.3.

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 distributing national pension benefits. 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.[15]

We merge these data with a list of locations for 312 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 some 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.

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

3.2 Background on the Japanese Old-Age Pension System

The Japanese National Pension System (JPS) is an 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 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

$$EPI = (A - NP) + B + C \tag{3.1}$$

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).[16] 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 in Appendix B.3 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

[16]The 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.

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 strategy for identifying the effect of pension payment receipt on nominal expenditures.

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:

$$\frac{X_{i,c,t}}{\overline{X}_{i,c}} = \sum_{j=-7}^{+7} \beta_j \cdot Payment_{i,t+j} + \delta_{dow} + \phi_{wom} + \psi_{my} + \xi_h + \eta_i + \epsilon_{i,c,t} \tag{4.1}$$

where $X_{i,c,t}$ is expenditures of shopper i on goods within category c on date t, and $\overline{X}_{i,c}$ 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 $Payment_{i,t+j}$ 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.[17] 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 $Payment_{i,t+j}$ equal to 1 if the shopper is older than the normal eligibility age of 65 for the national pension

[17] These 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.

and $t + j$ is a scheduled pension delivery day.[18] 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. Because we do not directly observe pension receipt, our estimates $\widehat{\beta}_j$ are intent-to-treat (ITT) estimates of the effect of pension receipt on expenditures.[19]

The model in (4.1) 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 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 (4.1) in two ways, with results presented in Appendix F.1. 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. We then create a set of dummies for each quartile of $Period$ and interact these dummies with individual fixed effects.[20] The interpretation of $Period$ depends on the category of expenditures under consideration. For highly aggregated categories 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

[18] We also report results using earlier and later age cutoffs for treatment, spanning the earliest (60) and the latest 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.

[19] 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.

[20] 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 $Period_{i,c,t+1} = 1$. Since 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.

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.

Differences between treatment and control groups in store visits could also be driven by retailer-specific pricing responses. For instance, if pension recipients are more likely to favor stores in a particular retailer chain, those stores might offer temporary sales promotions on products in elderly shoppers' expenditure baskets to attract pensioners around paydays. To account for retailer strategies, we compare our point estimates obtained with and without including fixed effects for the 20 grocery chains in our sample. In Appendix D, we examine retailers' temporary sale responses using standard filtering methods for identifying sales (e.g. Nakamura & Steinsson 2008). While we uncover evidence consistent with stores offering fewer and less generous discounts on high-quality goods (measured by a modal price) on paydays, when we consider all goods in the expenditure basket, stores are only slightly more likely to offer sales on paydays vs. non-paydays. We present results in Appendix F.1 using expenditures deflated by store-level price indices as the outcome variable, which is a closer analog to the notion of consumption in the model environment of Section 2.1.

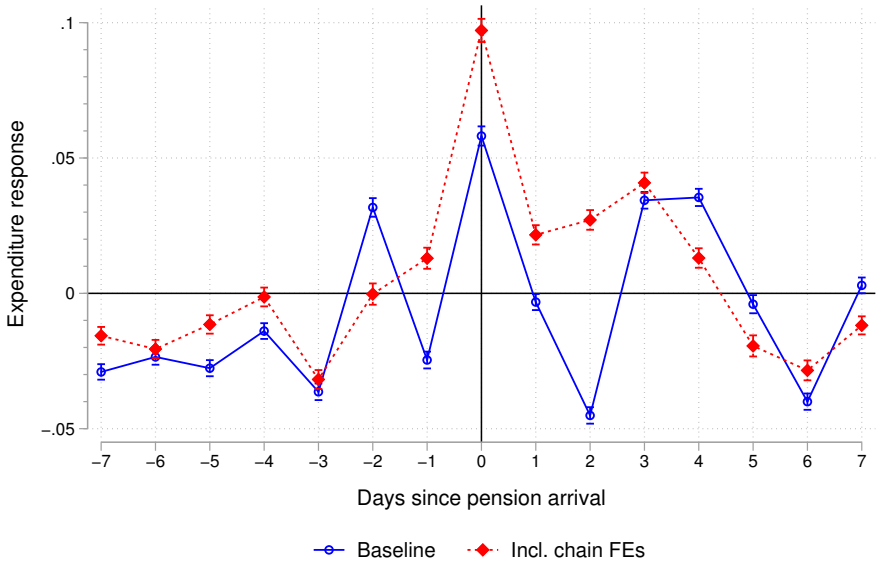
Finally, we emphasize that because payment dates are rescheduled when a payday falls on a weekend or holiday, the presence of weekends and holidays generates exogenous variation in the day of the month that income arrives. Our specification isolates this exogenous variation even without interacting individual 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 3 plots the event study coefficients β_j and 99% confidence intervals obtained from estimating equation (4.1) 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 largely spend more on payday on certain discretionary categories of goods such as alcohol, pre-packaged or prepared meals, and desserts. Table 1 shows heterogeneity in

FIGURE 3. Response of Total Expenditures to Payday



Notes: The figure plots in blue the estimated event study coefficients $\widehat{\beta}_j$ obtained from estimating regression equation (4.1) for all goods expenditures over a fixed panel of weekly shoppers. Point estimates in red obtained from a version of equation (4.1) which includes store chain fixed effects to account for variation in spending patterns due to retailer-specific pricing decisions. The y-axis records the percent increase in expenditures relative to average daily expenditures over the panel. We winsorize the top 1% of total daily expenditures. Bars indicate 99% confidence intervals, with standard errors obtained from clustering by shopper ID.

the response of expenditures when we estimate the event study equation in (4.1) using expenditures on different categories of goods to construct the dependent variable. We observe baseline spending responses ranging from 2.8% (processed fish) to 13.7% (alcohol and tobacco) of average daily expenditures on goods within a 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.

In even months when payments are delivered, shoppers who receive a payment spend more at the store around the payment date but are not much more likely to make a trip to a store compared to the same date in odd months when no payment is received. In Table 1, to isolate the intensive margin responses to pension receipt, we estimate a version of equation (4.1) defined at the individual-trip level, excluding observations within the panel where we observed zero spending. For the extensive margin, we instead replace the dependent variable in equation (4.1) with an indicator equal to unity if the shopper's expenditures are strictly positive on a given date. For all goods expenditures, the intensive margin estimates indicate that shoppers spent 9.6% more on payday visits relative to average expenditures

TABLE 1. Payday Spending Responses by Margin and Goods Category

Category	Overall	Incl. Chain FEs	Intensive	Extensive
All goods	0.059***	0.099***	0.096***	0.001***
	(0.001)	(0.002)	(0.002)	(0.000)
Raw foods	0.053***	0.093***	0.093***	0.001
	(0.001)	(0.002)	(0.002)	(0.000)
Prepared foods	0.079***	0.212***	0.219***	0.002***
	(0.003)	(0.007)	(0.006)	(0.000)
Sweets/desserts	0.069***	0.166***	0.167***	0.006***
	(0.003)	(0.007)	(0.007)	(0.000)
Alcohol	0.137***	0.275***	0.281***	0.004***
	(0.008)	(0.053)	(0.051)	(0.000)
Fresh produce	0.044***	0.077***	0.076***	0.001*
	(0.002)	(0.003)	(0.003)	(0.000)
Fresh fish	0.060***	0.226***	0.225***	0.001**
	(0.003)	(0.009)	(0.009)	(0.000)
Meat & poultry	0.049***	0.141***	0.132***	0.002***
	(0.002)	(0.005)	(0.004)	(0.000)
Grains	0.024***	0.092***	0.073***	0.001+
	(0.003)	(0.009)	(0.008)	(0.000)
Non-alcoholic beverages	0.048***	0.110***	0.101***	0.005***
	(0.002)	(0.006)	(0.006)	(0.000)
Tobacco	0.137***	0.135	0.140+	0.001
	(0.026)	(0.086)	(0.079)	(0.001)
Processed fruits/vegetables	0.051***	0.180***	0.120**	0.002***
	(0.006)	(0.039)	(0.037)	(0.000)
Preserved fish	0.028***	0.064***	0.060***	0.002***
	(0.003)	(0.011)	(0.010)	(0.000)
Other processed foods	0.056***	0.107***	0.102***	0.003***
	(0.002)	(0.003)	(0.003)	(0.000)

Notes: The table reports results for total spending and intensive and extensive margin versions of regression (4.1) using a panel of weekly shoppers and excluding leads and lags of *Payment*. Each cell in the table is the coefficient on *Payment* from a separate regression within a particular expenditure subcategory. Overall refers to the spending response estimated according to (4.1) and including shopper-day observations of zero expenditures. The second column indicates how our point estimates of the overall spending response changes when we include store chain fixed effects. 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. In each regression, we winsorize the top 1% of total daily expenditures. Robust standard errors clustered by shopper ID in parentheses. $***p < 0.001, **p < 0.01, *p < 0.05, +p < 0.1$

on other days when they visited a store. The point estimates along the intensive margin are broadly similar across spending categories to specifications in which we include store chain fixed effects. On the extensive margin, the probability that a consumer eligible for pension payments visits a store is roughly 0.1% 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.[21]

4.3 Heterogeneity by Pay Cycle Length

The results in the previous subsection on pensioners' expenditure responses 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 by augmenting equation (4.1) with terms that interact the payment dummy with a polynomial function of the pay cycle length:

$$\frac{X_{i,c,t}}{\overline{X}_{i,c}} = \beta_1 \cdot Payment_{i,t} \times Length_p + \beta_2 \cdot Payment_{i,t} \times \left(Length_p\right)^2$$
$$+\beta_3 \cdot Payment_{i,t} \times \left(Length_p\right)^3 + \delta_{dow} + \phi_{wom} + \psi_{my} + \xi_h + \eta_i + \epsilon_{i,c,t} \tag{4.2}$$

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 (4.2) 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 2 provides evidence from estimating equation (4.2) that shoppers tend to spend more

[21] We plot the event study coefficients by expenditure subcategory in Appendix F.1.

TABLE 2. Testing the Model of Payday Liquidity with Pent-up Demand

	(1)	(2)	(3)	(4)	(5)	(6)
$Payment \times Length$	0.0013^{***}	0.0015^{***}	0.0052^{***}	0.0055^{***}	-0.7360^{***}	-0.7390^{***}
	(0.0000)	(0.0000)	(0.0007)	(0.0007)	(0.0330)	(0.0330)
$Payment \times Length^2$			-0.00001^{*}	-0.00001^{*}	0.0247^{***}	0.0248^{***}
			(0.0000)	(0.0000)	(0.0011)	(0.0011)
$Payment \times Length^3$					-0.0002^{***}	-0.0002^{***}
					(0.0000)	(0.0000)
$Length$		-0.0002^{**}		-0.0002^{**}		-0.0002^{***}
		(0.0000)		(0.0000)		(0.0000)
Joint F-test (p-value)	–	–	0.000	0.000	0.000	0.000
N	660,895,995	660,895,995	660,895,995	660,895,995	660,895,995	660,895,995
Adj. R^2	0.017	0.017	0.017	0.017	0.017	0.017

Notes: The table reports results from estimating versions of equation (4.2), which interacts a dummy for pension receipt with the number of days since the last pension payment, for different polynomial orders. The dependent variable is expenditures on raw foods relative to average daily expenditures on raw foods. We use raw foods as the spending category for this exercise because it encompasses perishables which closely approximate consumption. We report p-values from an F-test for joint significance of coefficients on 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. Standard errors clustered by shopper ID. $^{***}p < 0.001, ^{**}p < 0.01, ^{*}p < 0.05$

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 0.13 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.[22] However, allowing for this direct effect of pay cycle length has no effect on the fitted values for $x(T)$.

Although the results in column (5) and column (6) of Table 2 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, our polynomial specifications are not globally well-defined. From this exercise, we take the coefficient on the linear interaction $Payment \times Length$ in column (1) to obtain our

[22]For 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.

preferred estimate of the expenditure spike function: $x(T) = 0.0013 \cdot T$. 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 consumers exhibit a spike in spending on or around the receipt of predictable and regular pension payments. What are the underlying economic mechanisms driving these responses? Our findings are consistent with features of the payday liquidity and quasi-hyperbolic discounting versions of the theoretical framework in Section 2, but other explanations, such as near-rationality or liquidity constraints could play a role.

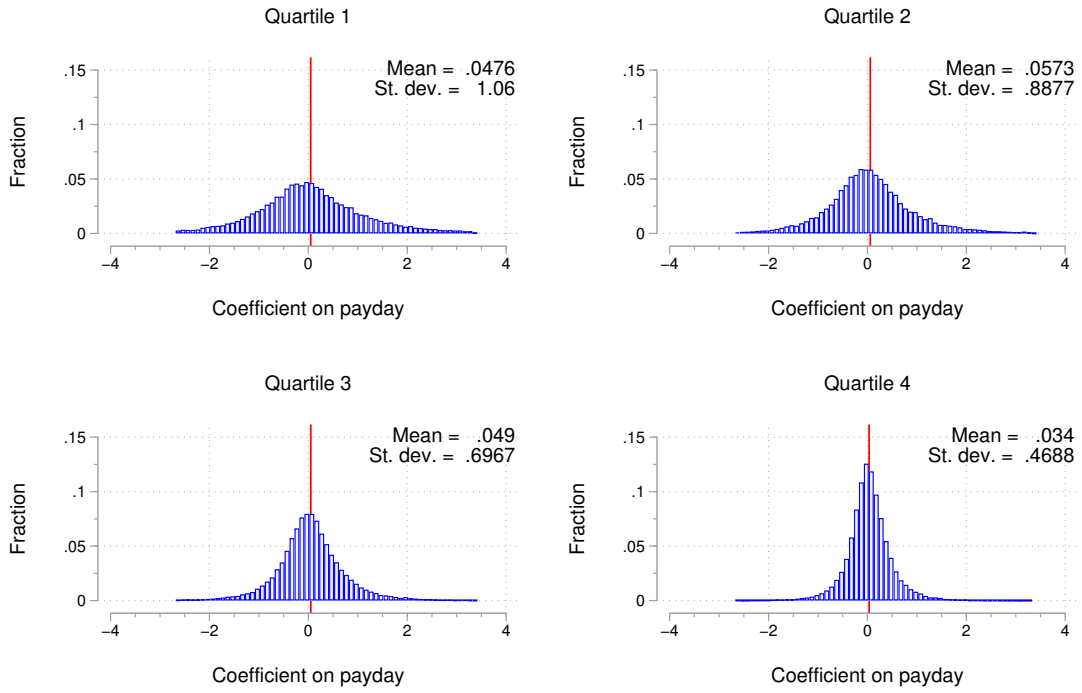
Distinguishing between these mechanisms is important for the welfare analysis we conduct in this paper. Near-rational consumers induce a welfare loss from excess spending that is proportional to the size of the payment relative to their permanent income (Kueng 2018). Recipients for whom the payment is 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 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.[23] In the following exercise, we restrict to shoppers aged 65 or over who are eligible for pension payments and who are

[23] We 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.

FIGURE 4. Payday Responses by Quartile of Pay Cycle Expenditures



Notes: We sort shoppers by quartile of average total pay cycle expenditures, which we consider a measure of permanent income. We then estimate the time series regression $X^i_{c,t}/\overline{X}_{i,c} = \beta^i \cdot Payment_{i,t} + \delta_{dow} + \phi_{wom} + \psi_{my} + \xi_h + \epsilon^i_t$ for each individual shopper ID using raw foods spending, and plot the distribution of the $\widehat{\beta^i}$ by each permanent income quartile. Vertical red lines indicate the within-quartile average payday response. We winsorize the top 1% of daily expenditures.

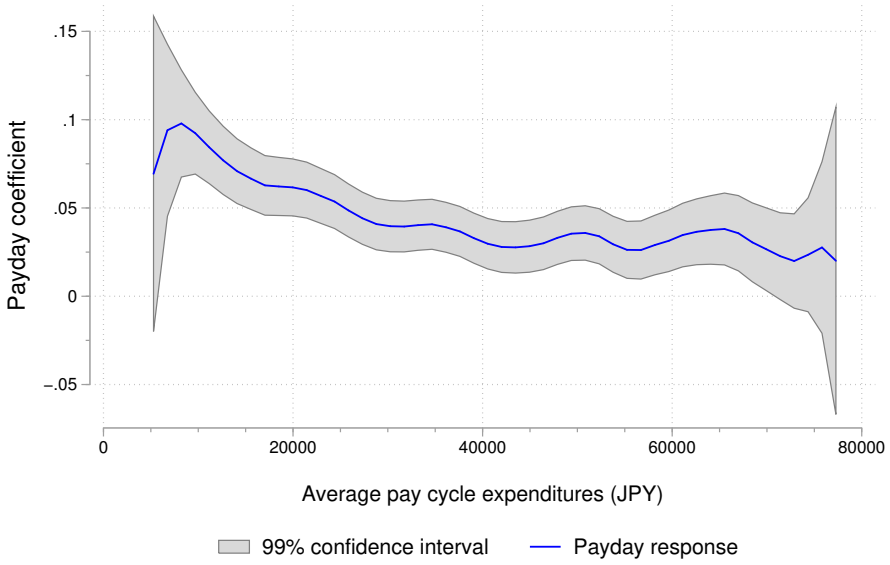
regular shoppers, defined as those who visit a store at least four times per month, or one week on average. These restrictions leave us with 116,533 shoppers.

We run our main specification in equation (4.1) 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.[24] Figure 4 plots the distribution of these individual payday responses binned by quartiles of average total expenditures over the pay cycle.[25] The figure shows that spending responses are more variable among low permanent income shoppers; such shoppers visit the store less

[24]We provide more details on the contents of this category in Appendix B.

[25]To limit variance in permanent income due to measurement error, we subset to shoppers who visit a store at least four times each month and appear in the sample for at least six full pay cycles. We also winsorize average daily expenditures at the 5th and 95th percentiles to help guarantee payments are a constant fraction of permanent income across different levels of permanent income.

FIGURE 5. Payday Responses as a Function of Average Pay Cycle Expenditures



Notes: We estimate the time series regression $X^i_{c,t}/\overline{X}_{i,c} = \beta^i \cdot Payment_{i,t} + \delta_{dow} + \phi_{wom} + \psi_{my} + \xi_h + \epsilon^i_t$ for each individual shopper ID using all goods expenditures. The figure fits a second-order local polynomial to the relationship between payday responses $\widehat{\beta^i}$ and average expenditures over the two-month pay cycle. We winsorize the top 1% of daily expenditures. 99% confidence intervals represented by the gray shaded area.

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 5 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 ≈ $250). However, this relationship becomes virtually flat once we look at shoppers with above-median average pay cycle expenditures.[26]

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

[26]In Appendix F.1, we separately re-estimate specification (4.1) for each permanent income decile. Along both the intensive and extensive margin, there is a clear negative relationship between payday responses and permanent income below the median, but a flat relationship above the median.

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.[27]

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?

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 (4.1) will capture changes in prices even when the real expenditure amount stays the same. In this section, we investigate how store-level prices respond to pension payment distribution dates. We construct a price index that isolates the retailer's pricing response from consumer substitution across goods but uncover limited evidence of retailer price discrimination.

5.1 Store-Level Price Indices

We first consider daily geometric average prices within a store:

$$\log P_{s,t} = \frac{1}{n_{s,t}} \sum_{k} \log p_{k,s,t} \tag{5.1}$$

[27]The relationship between past earnings and pension amounts is linear for most recipients. There is some concavity in the benefits schedule for recipients who receive large corporate retirement bonuses equal to several years of earnings. Under recent rules, the portion of bonus amounts that counts towards pension benefit calculations is capped at 1.5 million yen per month or 5.73 million yen in any year. Winsorizing average pay cycle expenditures helps ensure that payments are a constant share of income in our sample.

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.[28] Thus, this price index $\log P_{s,t}$ represents the average (log) price of goods that consumers over age 65, and who are thus eligible to receive pension payments, buy on each day in each store.

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

$$\widetilde{P}_{s,t} = \sum_{j=-7}^{+7} \gamma_j \cdot Payday_{t+j} + \delta_{dow} + \phi_{wom} + \psi_{my} + \xi_h + \eta_s + \epsilon_{s,t} \tag{5.2}$$

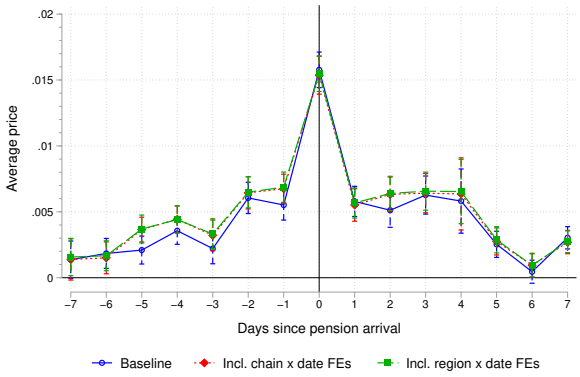
where $\widetilde{P}_{s,t}$ is the daily geometric average price index at store s, normalized by its mean value over our sample period. Hence, $\widetilde{P}_{s,t}$ captures in percentage terms how much store-level prices at date t deviate from their average geometric price index level. As in the expenditure 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 6 plots the coefficients on the payday dummies γ_j for a two-week interval around pension arrival dates. On 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. Panel B shows estimates of equation (5.2) where we replace prices on the left-hand side with the number of goods sold $\widetilde{n}_{s,t}$ normalized by its mean value over our sample period. On payday, the number of goods purchased increases by 6.2%. The results are impervious to the inclusion of store chain by month-year and Census region by month-year fixed effects, indicating that these patterns are not driven by retailers' seasonal

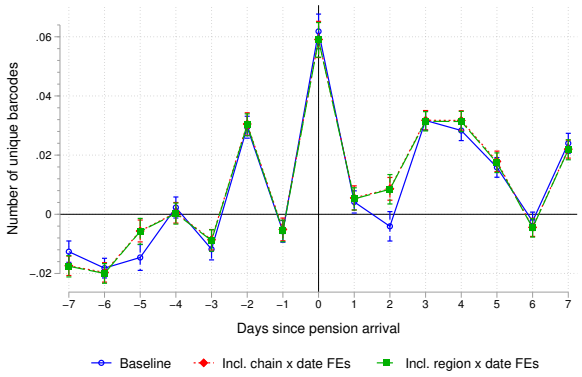
[28] The prices $p_{k,s,t}$ included in our sample are unit prices in the sense that the same product sold in different quantity will be attached to a distinct barcode k.

FIGURE 6. Store-level Price Index and Variety Responses to Payday

A. Average Price Index



B. Number of Unique Goods Purchased



Notes: Panel A of the figure plots the event study coefficients $\widehat{\gamma}_j$ obtained from estimating regression (5.2). The y-axis records the percent increase in store-level prices around payday, using the average price index defined in (5.1). Panel B plots the estimates from (5.2) where we use as the dependent variable the log of the daily number of goods sold (identified by unique barcodes) in a store relative to the average number sold over the sample time period. Point estimates in red obtained from augmenting (5.2) to include store chain $\times$ month-year fixed effects, while estimates in green result from including Census region $\times$ month-year fixed effects. Bars indicate 99% confidence intervals, with standard errors clustered at the store level.

sales promotions or regional shocks such as fluctuations in supply chain resilience.[29]

The positive response of store-level prices to payday could be due to stores 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, average prices paid by consumers could increase without any change in prices of individual goods set by stores to capture the incidence of 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

[29]Stores in our sample vary in the fraction of regular pensioner shoppers, and so one might expect stores clustered in areas with more pension recipients would be more likely to price discriminate. The fact that our estimates of store-level pricing and sales patterns are insensitive to chain and region-specific trends echoes the findings in DellaVigna & Gentzkow (2019) that chains do not price to market due to managerial inertia and concerns about alienating consumers.

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. 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.

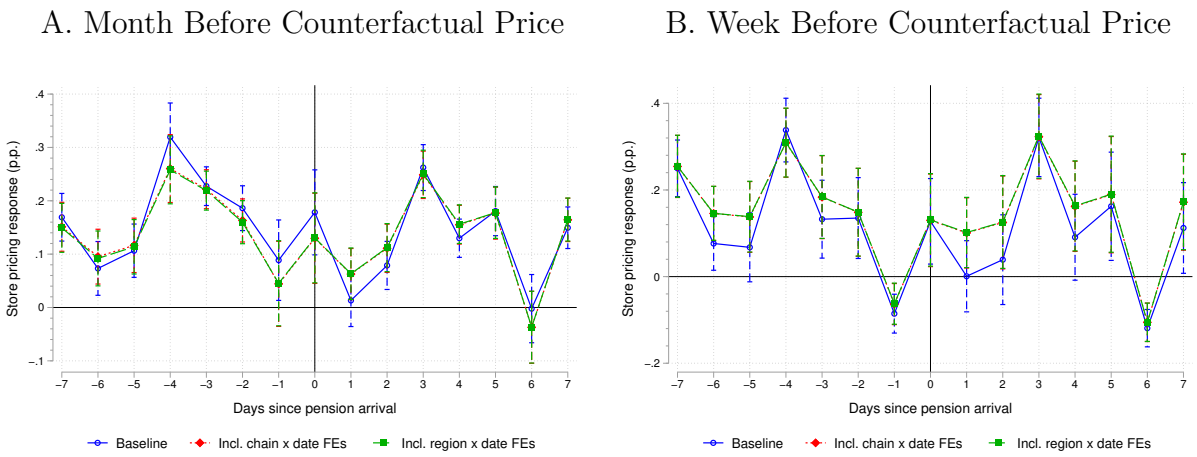
Figure 7 plots the estimates of the coefficients on *Payday* obtained from equation (5.2) using observed prices minus the same coefficients from using the store-level counterfactual price indices as the dependent variable. The differences are positive for each day within a two-week window around paydays, implying 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 our baseline counterfactual prices are 1.40% higher. The 0.18 p.p. gap between the two estimates captures the retailer's pricing response, which accounts for only 11.39% of the overall spike in payday prices (Panel A).

Similarly, for an alternative counterfactual where we use last observed prices exactly one week before payday (Panel B), we compute only a 0.13 p.p. gap, or just 11.40% of the payday spike in the average price index.[30] We provide a more complete analysis of store pricing responses in Appendix F.3 and conclude that for most goods subcategories there is no statistical difference between the average price index and the counterfactual on payday. Thus, 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.

In Appendix D, we explore whether retailers engaging in more aggressive sales promotions on paydays explains these substitution and variety effects. Consistent with an established IO literature on the seasonality of markdowns (e.g. MacDonald 2000; Chevalier, Kashyap, & Rossi 2003), we show that stores are slightly more likely to offer temporary sales on paydays, driven by higher discounts offered on *below-median* price barcodes within a four-digit goods category. Together with the results on consumer substitution towards splurge goods on payday, our evidence supports the explanation in Chevalier & Kashyap (2019) that retailers can attract more price-sensitive "bargain hunters" by offering discounts during periods of

[30]The week before counterfactual has the advantage of holding fixed the day of the week and the month-year in comparing the payday price index to the last price. We use the preceding odd month as our baseline counterfactual because the week before one might reflect forward-looking behavior by retailers trying to capture the incidence of payments, which downward biases our estimates of the pricing response when we do not include chain fixed effects.

FIGURE 7. Store Pricing Responses around Payday



Notes: We construct store-level price indices as in equation (5.1) and then estimate regression (5.2) to obtain estimates $\widehat{\gamma}_{1,j}$. We then construct a counterfactual last price index, using either the last observed price for each barcode in the odd month preceding the payday (Panel A) or the last observed price in the week prior to payday (Payday B), and rerun (5.2) on the exact same sample using this price index as the dependent variable to obtain $\widehat{\gamma}_{2,j}$. The figure plots the point estimates and 99% confidence interval bars for the differences $\widehat{\gamma}_{1,j} - \widehat{\gamma}_{2,j}$, with standard errors in each regression clustered at the store level. Point estimates in red obtained from augmenting each version of (5.2) to include store chain × month-year fixed effects, while estimates in green result from including Census region × month-year fixed effects. The differences capture the p.p. increase in prices on payday that cannot be explained by changes in the composition of goods purchased.

peak demand. Overall, we uncover limited evidence of price discrimination in response to customers receiving payments, which is consistent with more recent evidence from Goldin, Homonoff, & Meckel (2022) on retailer responses to SNAP in the U.S.[31] To the extent some price discrimination does occur, in Appendix A.4, we incorporate temporary sales into our optimal frequency framework and discuss how pricing responses to payday alter the government's choice of optimal pay cycle length.

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.

[31] Our failure to uncover substantive price discrimination is unlikely to be due to imperfect targeting of goods on store's side – that is, raising prices on goods with large expenditure shares in pensioners' baskets. Stores appear in our dataset if they contract with a marketing firm which offers them the necessary information to engage in this type of behavior.

6.1 Estimating Administrative Costs

Our identification of the slope of the administrative cost function in our model $\mu'(T)$ relies on a 1988 reform to the pension system which reduced the length of the pay cycle from three months to two months.[32] In February 1988, poor elderly recipients of the Old Age Welfare Pension benefit and existing JPS claimants who only qualified for the national pension but not the employee-based pension benefit 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 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.

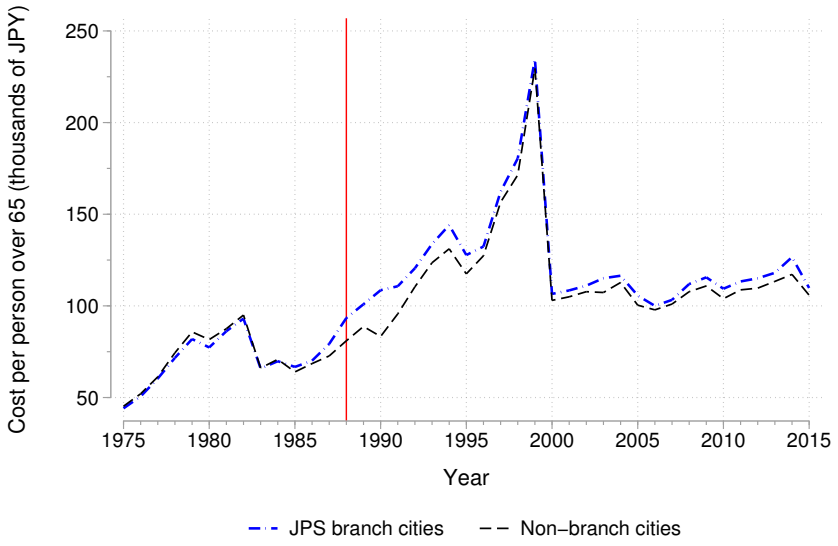
Given this natural experiment, we run a standard difference-in-differences regression:

$$\log \mu_{j,t} = \beta \cdot Branch_j \times Post_t + \gamma_j + \delta_t + \epsilon_{j,t} \tag{6.1}$$

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 includes costs in support of the pension system. We consider this measure to be good a proxy for administrative costs per claimant, since over 90% of individuals over age 65 receive benefits from JPS but only a small share of this population receive other forms of welfare benefits. Under the assumption of parallel trends in our cost per claimant measure, β captures the mid-run increase in costs due to increasing the frequency of benefits from

[32] In Appendix E, we discuss the details related to the data used in this section, as well as our procedures for identifying pension branch office locations.

FIGURE 8. Municipal Government Expenditures on the Elderly



Notes: The figure plots time series of municipal spending (in thousands of real 2012 JPY) on administering the pension system and elderly welfare benefits divided by the number of persons over age 65 residing in the municipality. The blue dashed line refers to the average of this ratio over all municipalities with a JPS branch office (N = 239), while the black dashed line averages over municipalities without a JPS branch office (N = 424). The vertical red line indicates the year (1988) when JPS switched from a quarterly to bimonthly payment schedule. We use branch office locations as of 1980 and exclude from the sample municipalities with gaps in the cost data. See Appendix E for details.

every three months to every two months, holding fixed the formula for benefit amounts.

Figure 8 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.[33] 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 (passed in 2000) which aimed to reduce costs and aggregate spending and benefits.[34] 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.

[33]Since 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.

[34]Major 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).

TABLE 3. Effect of Pension Frequency Reform on Municipal Admin Costs

	(1)	(2)	(3)	(4)	(5)	(6)
$Branch \times Post$	0.043**	−0.012	0.034	−0.003	0.073***	0.021
	(0.017)	(0.014)	(0.023)	(0.016)	(0.026)	(0.017)
City & year FEs	✓	✓	✓	✓	✓	✓
Incl. Tokyo	✓		✓		✓	
Incl. major cities	✓		✓		✓	
1985 population bin × year FEs			✓	✓	✓	✓
1985 per capita income bin × 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 in parentheses. 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. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$

Table 3 presents results from estimating equation (6.1). Administrative costs increase by approximately 4.3% in branch office cities in the post-reform years, or 7.3% conditional on non-parametric trends by municipal population and income bins. 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 do not uncover evidence that the switch from a quarterly to bimonthly payment schedule had a large effect on government costs of administering the public pension system.[35]

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%

[35]We 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.

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.[36]

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 (2.6) and characterized by the first-order condition in (2.7). 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 (2.16). This assumption is consistent with our empirical results in Section 4.3, where we showed that the spike in expenditures on payday increases by 0.13 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, $\overline{B}$. 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.[37]

[36] A 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.

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

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.[38] In FY 2011 there were 67.37 million contributors to the system and 38.67 million 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 $\overline{B} = 3{,}462$ 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.[39] 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.

Lastly, we choose preference parameters that match the results from estimating expenditure regressions of the form in equation (4.1). 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 $x'(T) = 0.0013$, reflecting that the payday spike in expenditures tends to be 0.13 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

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

[39] More 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 $\ell^* = \frac{\log(1-\widehat{x}\cdot n/c)}{\log(2/3)}$, where $\widehat{x}$ 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 $\ell = 0.28$.

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 9 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.

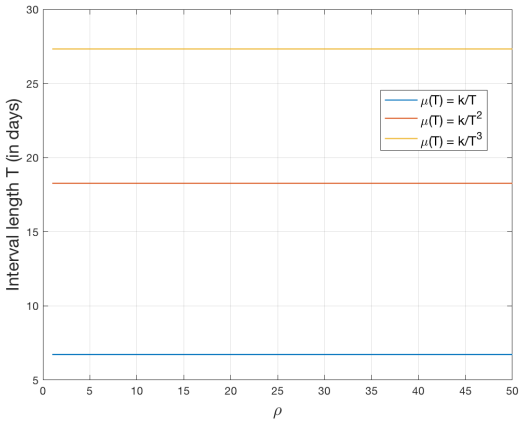
Second, for each of the administrative cost functions across both model versions, the implied optimal frequency is less than one month. For $\ell = 1$, the optimal frequency is $T^* = 6.72$ days, whereas for $\ell = 2$ we obtain $T^* = 18.27$ days, and for $\ell = 3$ we obtain cycles of length $T^* = 27.33$ days. In Panel B, we show how T^* varies continuously as a function of the convexity of administrative costs for a fixed value of the inverse IES ($\rho = 10$).[40] 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.

[40] Again, 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.

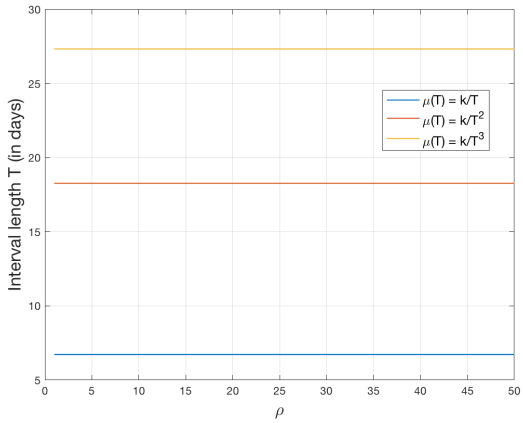
FIGURE 9. Results from Calibrated Optimal Frequency Model

A. Optimal frequency by Inverse IES

Quasi-hyperbolic Discounting

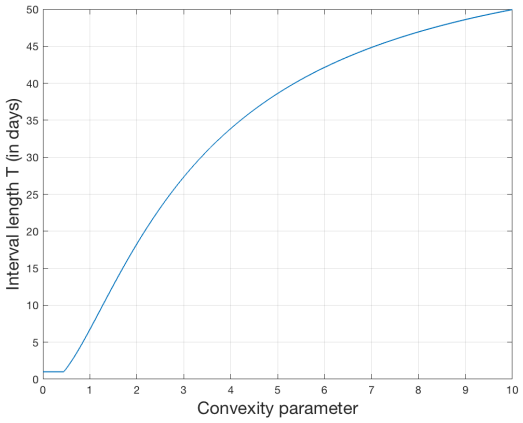


Payday Liquidity

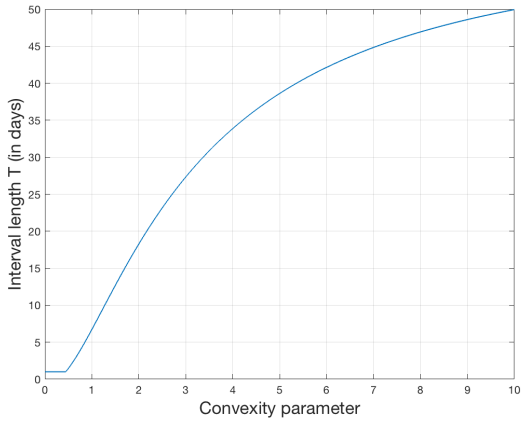


B. Optimal frequency by Convexity of Administrative Costs

Quasi-hyperbolic Discounting



Payday Liquidity



Notes: The figure plots the optimal frequency implied by quasi-hyperbolic discounters and by payday liquid consumers as a function of either the consumers' inverse intertemporal elasticity of substitution (Panel A) or the convexity of administrative costs (Panel B). We solve for T in the government's FOC in equation (2.7) for each version of the model using standard non-linear optimization packages. See main text for details on calibration and Appendix A for derivations.

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 subject to a tradeoff between minimizing the welfare loss from consumption non-smoothing and incurring increased administrative costs. We apply our model to the bimonthly Japanese pension system and find, under a battery of alternative assumptions about consumer preferences and administrative costs, that the government could improve welfare by moving to a monthly disbursement schedule.

Our framework applies more generally to other public benefits programs. We demonstrate under several empirically plausible models of individual consumption paths that welfare losses from non-smoothing are likely to be small. However, the main lesson here is that for public benefits programs where administrative costs vary minimally with respect to the frequency of disbursement, the government can achieve a welfare improvement by splitting up entitlements into smaller, more frequent payments. Even under extreme assumptions about consumer preferences or administrative costs, the optimal interval length rarely exceeds one month, which supports the worldwide prevalence of monthly payment schedules.

Further, our application to high-frequency retail data sheds light on consumer and retailer responses to benefit payment dates. Shoppers sharply increase expenditures within two days of scheduled payment dates, but most of this response is concentrated towards purchases 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. While within-store average prices spike on paydays, this is almost entirely due to consumers substituting towards a basket of more numerous, higher quality goods, rather than retailers raising prices or suspending temporary sales on specific goods. Our work therefore points to documenting consumers' substitution and variety responses to income payments as a policy-relevant path for future research.

References

Aguiar, M. & E. Hurst (2005): "Consumption versus Expenditure," *Journal of Political Economy*, 113(5): 919-948.

Anderson, E., B. A. Malin, E. Nakamura, D. Simester, & J. Steinsson (2017): "Informational Rigidities and the Stickiness of Temporary Sales," *Journal of Monetary Economics*, 90: 64-83.

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.

Baker, S., B. Baugh, & L. Kueng (2021): "Income Fluctuations and Firm Choice," *Journal of Financial and Quantitative Analysis*, 56(6): 2208-2236.

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., I. Ben-David, H. Park, & J.A. Parker (2021): "Asymmetric Consumption Smoothing," *American Economic Review*, 111(1): 192-230.

Baugh, B. & F. Correia (2022): "Does Paycheck Frequency Matter? Evidence from Micro Data," *Journal of Financial Economics*, 143(3): 1026-1042.

Baugh, B. & J. Wang (2021): "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*, 91(4): 832-857.

Best, M., J. Cloyne, E. Ilzetzki, & H. Kleven (2020): "Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches," *Review of Economic Studies*, 87(2): 656-690.

Bils, M. & P.J. Klenow (2001): "Quantifying Quality Growth," *American Economic Review*, 91(4): 1006-1030.

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. & M.D. Collado (2001): "The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates," *American Economic Review*, 91(3): 681-692.

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.

Coibion, O., Y. Gorodnichenko, & D. Koustas (2021): "Consumption Inequality and the Frequency of Purchases," *American Economic Journal: Macroeconomics*, 13(4): 449-482.

DellaVigna, S. & M. Gentzkow (2019): "Uniform Pricing in U.S. Retail Chains," *Quarterly Journal of Economics*, 134(4): 2011-2084.

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.

Evans, W.N. & T.J. Moore (2012): "Liquidity, Economic Activity, and Mortality," *Review of Economics and Statistics*, 94(2): 400-418.

Farhi, E. & X. Gabaix (2020): "Optimal Taxation with Behavioral Agents," *American Economic Review*, 110(1): 298-336.

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 (2022): "Issuance and Incidence: SNAP Benefit Cycles and Grocery Prices," *American Economic Journal: Economic Policy*, 14(1): 152-178.

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.

Huffman, D. & M. Barenstein (2005): "A Monthly Struggle for Self-Control? Hyperbolic Discounting, Mental Accounting, and the Fall in Consumption between Paydays," IZA Discussion Paper.

Kueng, L. (2015): "Revisiting the Response of Household Spending to the Alaska Permanent Fund Dividend using CE Data," *mimeo*, Swiss Finance Institute.

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.

Nakamura, E. & J. Steinsson (2008): "Five Facts About Prices: A Reevaluation of Menu Cost Models: Supplementary Material," *Quarterly Journal of Economics*, 123(4): 1415-1464.

Nevo, A. & K. Hatzitaskos (2006): "Why Does the Average Price Paid Fall during High Demand Periods," *mimeo*, Wharton.

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(2): 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.

Shoji, T. (2020): "Liquidity Constraints and a Surge in Storable Goods Demand: Evidence from Japan," *mimeo*, Seikei University.

Stephens Jr., M. (2003): "'3rd of tha 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.

Online Appendix to

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

by Cameron LaPoint (Yale SOM) & Shogo Sakabe (Columbia)

Contents

A Model Derivations **46**

A.1 General Expression for the Welfare Loss 46

A.2 Details for Quasi-hyperbolic Discounting Case 47

A.3 Details for Payday Liquidity Case 48

A.4 Menu Costs and Retailer Responses 48

B Further Details on Expenditure Data **52**

B.1 Constructing Spending Subcategories 52

B.2 Sample Size and Statistics by Trip Frequency 55

B.3 Comparison to JSTAR Survey Data 56

C Using Google Data to Test for Inattention **56**

C.1 Constructing Daily Google Search Data 56

C.2 Estimation Procedure 58

D Temporary Sales around Payday **59**

E Japanese Pension System Branch Offices **66**

F Additional Results on Payday Responses **68**

F.1 Expenditure Reponses 68

F.2 Spending Responses around Retirement Age 68

F.3 Pricing Responses 68

A Model Derivations

A.1 General Expression for the Welfare Loss

We can rearrange the compensating variation equation in (2.5) to obtain an intuitive expression for the welfare loss:

$$1 - \lambda(T) = 1 - \frac{T \cdot u^{-1}\left\{ \frac{1}{T}\int_0^T u\Big(c_0 \cdot \exp\Big(-f(t)\Big)\Big)dt \right\}}{\int_0^T c_0 \cdot \exp\Big(-f(t)\Big)dt} \tag{A.1}$$

where $c_0 \equiv \exp(\theta)$ is the value of consumption on payday, as defined in (2.4). 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:

$$1 - \lambda(T) = 1 - \frac{T \times \overline{C}}{C^{tot}} = \frac{C^{tot} - T \times \overline{C}}{C^{tot}} \tag{A.2}$$

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

$$\lambda'(T) = \frac{(T \times \overline{C})' \cdot C^{tot} - (T \times \overline{C}) \cdot (C^{tot})'}{(C^{tot})^2} \tag{A.3}$$

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

$$\begin{aligned}(\overline{C})' &= \frac{\partial}{\partial T} u^{-1}\Bigg\{ \underbrace{\frac{1}{T}\int_0^T u\Big(c_0 \cdot \exp(-f(t))\Big)dt}_{\equiv \overline{U}(T)} \Bigg\} \\ &= \frac{\partial}{\partial T} u^{-1}\Big(\overline{U}(T)\Big) \cdot \frac{\partial \overline{U}(T)}{\partial T} = u^{-1}\Big(\overline{U}(T)\Big)(u^{-1})'\Big(u(\overline{U}(T))\Big) \cdot \frac{\partial \overline{U}(T)}{\partial T} \\ &= u^{-1}\Big(\overline{U}(T)\Big)(u^{-1})'\Big(u(\overline{U}(T))\Big) \cdot \frac{1}{T}\left[u\Big(c_0 \cdot \exp(-f(T))\Big) - \frac{1}{T}\int_0^T u\Big(c_0 \cdot \exp(-f(t))\Big)dt \right]\end{aligned} \tag{A.4}$$

$$(C^{tot})' = \frac{\partial}{\partial T}\int_0^T c_0 \cdot \exp(-f(t)) = c_0 \cdot \exp(-f(T)) = \overline{B} \tag{A.5}$$

Note that if the Envelope theorem held here, the $\overline{C}'$ term would be zero, since the agent would be maintain the same daily average level of utility regardless of the interval length.

Thus, combining (A.3)-(A.5) we can write the FOC to the government's problem in (2.7) as

$$\frac{p(\overline{C} \times C^{tot}) + p(T \times (\overline{C})' \cdot C^{tot}) - p(T \times \overline{C}) \cdot c_0 \cdot \exp(-f(T))}{(C^{tot})^2} = \mu'(T) + p \cdot \overline{B} \tag{A.6}$$

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 (A.1) 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:

$$1 - \lambda(T) = \frac{c_0}{\overline{B}} \cdot \exp\Big(- \nu \cdot T/2\Big) \tag{A.7}$$

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

$$\overline{C} = u^{-1}\left\{ \left(\frac{c_0^{1-\rho}}{\nu \cdot T(1-\rho)^2} \right) \Big[1 - \exp\Big((\rho - 1)\nu \cdot T \Big) \Big] \right\} \tag{A.8}$$

$$C^{tot} = \frac{c_0}{\nu} \cdot \Big(1 - \exp(-\nu \cdot T) \Big) = \overline{B} \cdot T \tag{A.9}$$

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

$$1 - \lambda(T) = \begin{cases} 1 - \frac{1}{\overline{B}} \cdot \exp\Big(c_0 - \nu T/2 \Big) & \text{if } \rho = 1 \\ 1 - \frac{c_0}{\overline{B}} \cdot \left[\frac{1 - \exp\Big((\rho-1)\nu T \Big)}{\nu T(1-\rho)} \right]^{\frac{1}{1-\rho}} & \text{if } \rho \neq 1 \end{cases} \tag{A.10}$$

The marginal welfare loss is therefore:

$$-\lambda'(T) = \begin{cases} \frac{c_0 \cdot \nu}{2\overline{B}} \cdot \exp(-\nu T/2) & \text{if } \rho = 1 \\ \frac{c_0}{\overline{B}\nu(1-\rho)^2 T^2} \left[\frac{\exp(\nu(\rho-1)T) - 1}{\nu(\rho-1)T} \right]^{\frac{\rho}{1-\rho}} \frac{\Big(\exp(\nu(\rho-1)T)(\nu(\rho-1)T - 1) + 1 \Big)}{1 - \frac{c_0}{\overline{B}} \exp(-\nu T/2)} & \text{if } \rho \neq 1 \end{cases} \tag{A.11}$$

A.3 Details for Payday Liquidity Case

Expression (2.15) 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:

$$-\lambda'(T) = \frac{(\bar{c}/\overline{B})(1+x(T))^{1/T-1}}{T^2}\Big[\Big(1+x(T)\Big)\cdot\log\Big(1+x(T)\Big) - Tx'(T)\Big] \tag{A.12}$$

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 (A.1) to compute the marginal welfare loss. In this context $\overline{C}$ and C^{tot} are given by:

$$\overline{C} = \frac{\bar{c}}{T^{1/(1-\rho)}}\Big[\Big(1+x(T)\Big)^{1-\rho} + (T-1)\Big]^{\frac{1}{1-\rho}} \tag{A.13}$$

$$C^{tot} = \overline{B} + \frac{\overline{B}\cdot T}{T+x(T)} = \overline{B}\cdot T \tag{A.14}$$

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

$$-\lambda'(T) = \frac{\bar{c}}{\overline{B}}\cdot\frac{\Big((1+x)^{\rho}-(1+x)\Big)T^{\frac{2-\rho}{\rho-1}}\Big((1+x)^{1-\rho}+x-1\Big)^{\frac{1}{1-\rho}}}{(\rho-1)\Big(T(1+x)^{\rho}-(1+x)^{\rho}+(1+x)\Big)} \tag{A.15}$$

A.4 Menu Costs and Retailer Responses

We now incorporate monopolistic retailers into the optimal frequency framework. Retailers set prices to maximize real profits, subject to menu costs that must be paid for increasing prices on payday. This puts further downward pressure on the optimal payment frequency, as the government can render the gains from price discrimination small relative to the menu cost by shortening the pay cycle length. We consider two variants of this model. One in which there is a single monopolistic retailer who sets the price for the entire consumption basket, and another version in which retailers specialize in a particular variety. The intuition is largely the same across the two versions.[1]

[1]Note in focusing on payday liquidity we adhere to the discrete time representation of the model used in Section 2.3 in the main text.

A.4.1 A Single Monopolistic Firm

There is a single monopolistic retailer who sets prices for a common basket of goods purchased by all consumers. The retailer maximizes pay cycle profits according to the problem:

$$\max_{\{P_t\}} \left\{ \sum_{t=0}^{T-1} P_t \cdot Y_t - W_t \cdot L_t - \kappa \cdot W_t \times \mathbb{1}_t - P_t \cdot \Gamma \right\} \tag{A.16}$$

where $\kappa \cdot W_t$ represents the real menu cost denoted in κ units of labor L_t, and $\mathbb{1}_t$ is an indicator for whether the retailer changes the price P_t at time t. Firms use labor L_t provided by the working population, which in this setting is the non-pensioners or non-recipients (indexed hereafter by "NR" of benefits). Labor is used to satisfy non-recipient demand $F(L_t) = C_t^{NR}$, and the retailer pays a fixed real inventory cost Γ each period to external suppliers who stock their shelves and inelastically satisfy the residual demand of the recipients (C_t^R).[2]. The pay cycle runs from payday at $t = 0$ until $T - 1$, where T is the number of days in between payments. We set the discount rate on future profits to unity for ease of exposition.

We assume claims to retailer profits are held by external (e.g. foreign) shareholders. If we assume all profits were rebated to the working population, then the problem would collapse to the original one outlined in Section 2, since any price discrimination would ultimately be to the benefit of the working population, and then subsequently taxed away by the government to finance the benefits system. Alternatively, if profits are proportionally split between recipients and non-recipients, then whether there is an additional welfare loss depends on how much of the surplus profits accrued to the non-recipients gets eroded by real wage deflation from payday price hikes. While interesting to consider, profit redistribution significantly complicates the model without changing the basic tradeoffs faced by the government.

The fraction p of the population who are non-recipients choose consumption and labor over the pay cycle to solve the following problem:

$$\max_{\{C_t, L_t\}} \left\{ \sum_{t=0}^{T-1} u(C_t) - \nu(L_t) \right\} \quad \text{s.t.} \quad P_t \cdot C_t = S_t + W_t \cdot L_t - \frac{\tau(b)}{T} \tag{A.17}$$

Suppose the fraction $(1 - p)$ of the population who are pensioners – or more generally, those consumers receiving regular benefits – have payday liquidity. That is, as described in equations (2.12) and (2.13) in Section 2.3, they smooth real consumption throughout the pay cycle at a level $\bar{c}$, but splurge by consuming $\Big(1 + x(T)\Big) \cdot \bar{c}$ on payday $(t = 0)$.[3]

[2]This business structure is akin to a retailer who has their own private label brand of products sold alongside products received from a third-party wholesaler which they price at a markup. If instead we simply assumed $F(L_t) = C_t$, the retailer would always find it profitable to raise prices on payday, meaning the government could influence price changes along the intensive margin instead of the extensive margin.

[3]The qualitative results of this model obtain for quasi-hyperbolic discounters as well, since the welfare loss from non-smoothing consumption will balloon with respect to the interval length T. Focusing on the payday liquidity case helps illustrate how retailer price discrimination affects the optimal frequency with closed-form expressions for prices and consumption.

The government runs a balanced budget as in (2.1) and finances benefits through lump-sum taxation of the non-recipients. Because the government's choice of interval length influences the retailer's pricing strategy, relative to equation (2.6) there is an additional term in the objective function which accounts for the fact that if retailers choose to price discriminate, benefit non-recipients experience a reduction in their maximized level of utility $U^*(\cdot)$:

$$\min_T \Big\{ -p \cdot \lambda(T) + \gamma \cdot \Big(p \cdot b(T) + \mu(T)\Big) + \Big(U^*(C^{1,NR}(T)) - U^*(C^{0,NR}(T))\Big)\Big\} \tag{A.18}$$

As in the model without menu costs, the first term represents the purely paternalistic motivation of the government to set a lower T to mitigate the non-smoothing behavior of the payment recipients, and the second term captures the costs of imposing taxes on the working population, with marginal cost of funds $\gamma = 1$ in the case of lump-sum taxes. The final term is the gap between the maximized level of utility when the non-recipient's consumption is either altered by payday pricing $\mathbb{1}_0 = 1$ or smooth for $\mathbb{1}_0 = 1$.

If the government sets T such that the retailers do not engage in a temporary price hike, then $C^{0,NR}(T) = C^{1,NR}(T)$ and the non-recipients experience no welfare loss relative to the original problem without menu costs and monopolistic pricing. On the other hand, if administrative costs are sufficiently convex, it may make sense for the government to allow price discrimination by setting T such that excess payday demand attracts payday price hikes and generates a loss for the non-recipients.

In any case, this problem depends on a number of additional parameters not discussed in the main text, such as the disutility of labor $\nu(L)$, firm production $F(L)$, and the size of menu costs κ. To illustrate the solution to this problem via closed-form expressions, we now consider a very simple parameterization where retailers have linear production $F(L) = A \cdot L_t$, with a constant productivity A, and workers have log utility from consumption with constant marginal disutility of labor: $u(C_t) - \nu(L_t) = \log(C_t) - \omega L_t$.

First, under this parameterization, we can combine the non-recipient's FOC for consumption and labor supply to show that the real wage is proportional to consumption:

$$\frac{1}{C_t^{NR}} = P_t \cdot \lambda_t \tag{A.19}$$

$$W_t \cdot \lambda_t = \omega \tag{A.20}$$

$$\implies \frac{W_t}{P_t} = \omega \cdot C_t^{NR} \tag{A.21}$$

Using this fact and the linear production function, we can then rewrite the retailer's real profits over the pay cycle as

$$\sum_{t=0}^{T-1} C_t - \omega C_t^{NR} \cdot \frac{C_t^{NR}}{A} - \kappa \cdot \omega C_t^{NR} \times \mathbb{1}_t - \Gamma \tag{A.22}$$

where C_t is aggregate demand for the retailer's products, which is a weighted average of the

demands of the non-recipients (NR) and recipients (R):

$$C_t = p \cdot C_t^R + (1-p) \cdot C_t^{NR} \tag{A.23}$$

Since the payday liquid consumers are price-insensitive, in that they pick consumption based on a heuristic rule of thumb, the firm's problem reduces to picking a point on the non-recipient's demand curve to maximize real profits. The firm's FOC with respect to C_t^{NR} yields equilibrium real expenditures of benefit non-recipients:

$$C_t^{NR} = \begin{cases} \frac{A}{2\omega} & \text{if } \mathbb{1}_t = 0 \\ \frac{A\Big((1-p)-\kappa\cdot\omega\Big)}{2\omega\cdot(1-p)} & \text{if } \mathbb{1}_t = 1 \end{cases} \tag{A.24}$$

Hence, non-recipients smooth consumption unless there is a temporary price hike on payday ($t = 0$), and they consume more in real terms in the absence of a price hike ($\mathbb{1}_t = 0$).

Combining non-recipient expenditures with the condition in equation (A.21) yields the pricing strategy as a function of the nominal wage:

$$P_t = \begin{cases} \frac{2W_t}{A} \text{ if } \mathbb{1}_t = 0 \\ \frac{2W_t\cdot(1-p)}{A\cdot(1-p-\kappa\omega)} \text{ if } \mathbb{1}_t = 1 \end{cases} \tag{A.25}$$

where for a given nominal wage W_t, prices are always higher when $\mathbb{1}_t = 0$, meaning retailers raise prices rather than engage in temporary sales on payday.

Because recipients maintain a smooth consumption profile beyond payday, the retailer will only find it profitable to pay the menu cost to change prices at $t = 0$. This means retail price discrimination, if it does occur, will be concentrated on payday, and the relevant comparison from the retailer's perspective is profits at time zero when $\mathbb{1}_0 = 0$ vs. profits when $\mathbb{1}_0 = 1$. Denote real profits on payday with the price hike as $\Pi_0^{R,1}$, and real profits on payday without the price hike as $\Pi_0^{R,0}$. Then the retailer will find it worthwhile to pay the menu cost to capture some of the excess demand from the recipients on payday whenever the following condition is met:[4]

$$x(T) > \frac{1-p}{p\bar{c}\cdot\Big((2-A)(1-p)-A\cdot\kappa\omega)\Big)}\cdot\left\{A\kappa + \frac{A(1-p-\kappa\omega)}{(1-p)}\cdot\left(p\bar{c} + \frac{A(1-p)-\kappa\omega}{2\omega}\right)\right\} \tag{A.26}$$

This condition says that the retailers will find it profitable to price discriminate whenever excess demand from the benefit recipients, captured by $x(T)$, is sufficiently high that the

[4]The common component of profits across the two pricing strategies over $1 \leq t \leq T-1$ is:

$$\frac{T-1}{2}\cdot\Big[p\cdot\bar{c} + (1-p)\cdot A/2\omega\Big]$$

increase in profits offsets reduced demand from the non-recipients plus the menu cost.

Finally, the difference in the non-recipients' optimized level of utility when there is price discrimination is given by:

$$\log(C_0^{NR,1}) - \log(C_0^{NR,0}) + \omega \cdot \left(L_0^1 - L_0^0\right) = 0 \tag{A.27}$$

The difference in the optimized level of utility is ultimately zero in this special version of the model where all price discrimination occurs on the extensive margin. That is because the loss in utility over consumption associated with the price hike is completely offset by the reduction in disutility from labor supply. Therefore there exist parameter sets (p, A, κ, ω) such that even if the government optimally chooses the pay cycle length to promote consumption smoothing among the benefit recipients, retailers may respond by raising prices on payday.

B Further Details on Expenditure Data

B.1 Constructing Spending Subcategories

Our data contain 968 major goods categories provided by Magee Co., with some categories containing thousands of barcodes. 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 B.1. Table B.2 through Table B.4 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	Confectionery	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, liqueurs, 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

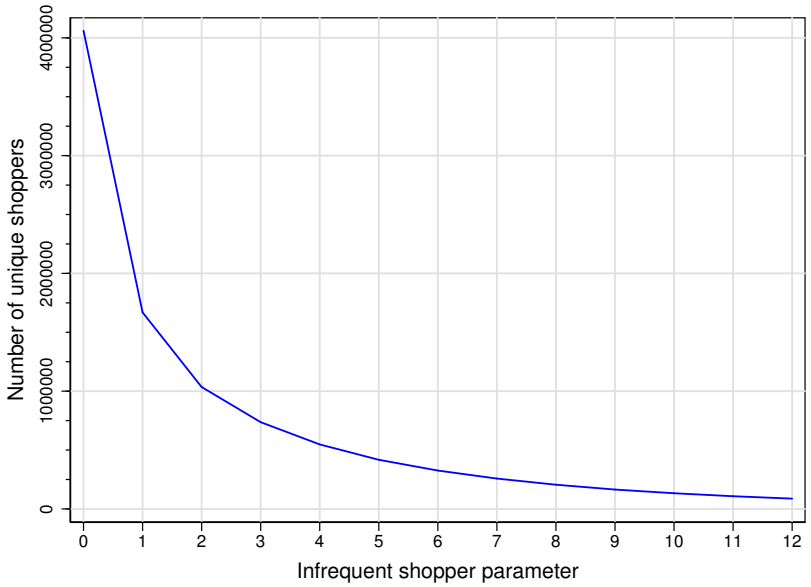
	Fruits/Vegetables	Fish	Grains	Other Processed
Avg. monthly expenditures	981	2,203	2,734	5,661
% Normal retirement age	25.0%	36.0%	26.8%	28.8%
Avg. number of monthly trips	4.2	5.5	6.8	8.0
Avg. periodicity	4.7	4.3	3.5	3.0
# Stores	503	505	509	510
# Shoppers	30,952	127,599	417,280	782,680

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

B.2 Sample Size and Statistics by Trip Frequency

FIGURE B.1. Shoppers who Visit Store k Times per Month



Notes: The figure shows the number of unique shoppers, identified by loyalty point card number, who make at least k visits (x-axis) to a store in our sample during each month of our sample period.

Table B.5. Retail Expenditures Summary Statistics

	All Goods	Raw Foods
Avg. monthly expenditures	23,102	10,127
	(15,752)	(7,260)
Avg. expenditures per trip	2,320	1,086
	(1,180)	(579)
Avg. number of monthly trips	9.1	8.5
	(6.2)	(5.7)
Avg. periodicity	2.7	2.9
	(2.5)	(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 summary statistics for purchases on all goods and purchases on raw foods, which include all perishable food items. Monthly and per-trip expenditures are reported in nominal Japanese yen. 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 attained the normal retirement age (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.

B.3 Comparison to JSTAR Survey Data

C Using Google Data to Test for Inattention

This appendix outlines how we construct 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. Baugh & Wang (2021) argue that consumers may be unaware of or ignore rules governing regular Social Security payment schedules, even when such rules are public information. Contrary to this notion, we find a clear uptick in searches on the eve of a scheduled payment date.

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. Google Trends data are available from January 1, 2004.[5] 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.[6]

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:

$$\text{pension search index}_{r,t} = 100 \times \frac{\text{SP}_{r,t}}{\max_r \text{SP}_{r,t}} \tag{C.1}$$

where the search propensity (SP) is computed as:

$$\text{SP}_{r,t} = \frac{\Big[\text{Google searches containing "pension"}\Big]_{r,t}}{\Big[\text{total Google searches}\Big]_{r,t}} \tag{C.2}$$

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 using total page views

[5]Note 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

[6]Our results are qualitatively similar for non-interacted keyword searches of "public pension payments."

from the 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 the index level is relative to the week with the greatest search volume within that subperiod, 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.

Therefore 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:

$$Z^*_{i,j} \equiv \underbrace{\frac{\text{SVI}_{i,\tau}}{\text{SVI}_{j',\tau_k}}}_{\text{correction factor}} \times \underbrace{\text{SVI}_{i,j,\tau_k}}_{\text{downloaded data}} \tag{C.3}$$

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

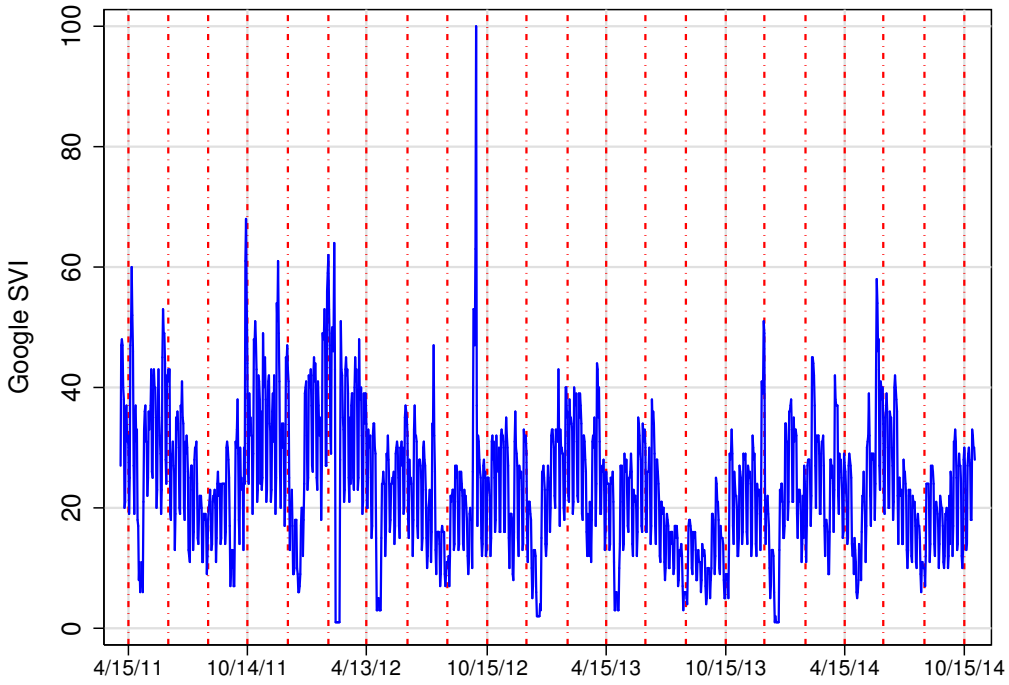
6. Create the new series $SVI^*_{i,j,\tau}$ from the data frame of scale factors Z^*:

$$SVI^*_{i,j,\tau} \equiv \text{round}\left(\frac{Z^*_{i,j}}{\max_{(i,j)}\{Z^*\}} \times 100\right) \tag{C.4}$$

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

FIGURE C.1. Raw Daily Google Searches for "Public Pension Payments"



Notes: The figure displays the daily time series of the Japanese Google SVI for "public pension payments." We construct the index according to the steps summarized in equations (C.3) and (C.4). Dashed red lines indicate scheduled pension payment dates during our sample period for the scanner data.

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.

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:

$$\widetilde{SVI}_t = \sum_{j=-7}^{+7} \beta_j \cdot Payday_{t+j} + \gamma \cdot t + \delta_{dow} + \phi_{wom} + \psi_{my} + \xi_h + \alpha_p + \epsilon_t \qquad \text{(C.5)}$$

where we control for the full set of time fixed effects capturing cyclicality 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 affected some 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 occur close to scheduled delivery dates.[7]

Figure C.2 plots the estimated coefficients $\widehat{\beta}_j$ 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 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.

D Temporary Sales around Payday

In this appendix, we investigate whether retailers respond to increased demand on pension paydays by engaging in price discrimination through reducing the frequency of temporary sales. To the contrary, we find that, if anything, stores are more likely to offer more generous sales on paydays. Moreover, this conclusion holds for several methods we use to identify temporary sales which are common in the literature on sticky prices in high-frequency retail data. These results complement our analysis in Section 5.2, where we use our counterfactual "last price" indices to estimate that approximately 90% of the observed increase in average store-level prices on payday is due to consumers substituting towards an expanded basket of higher quality goods.

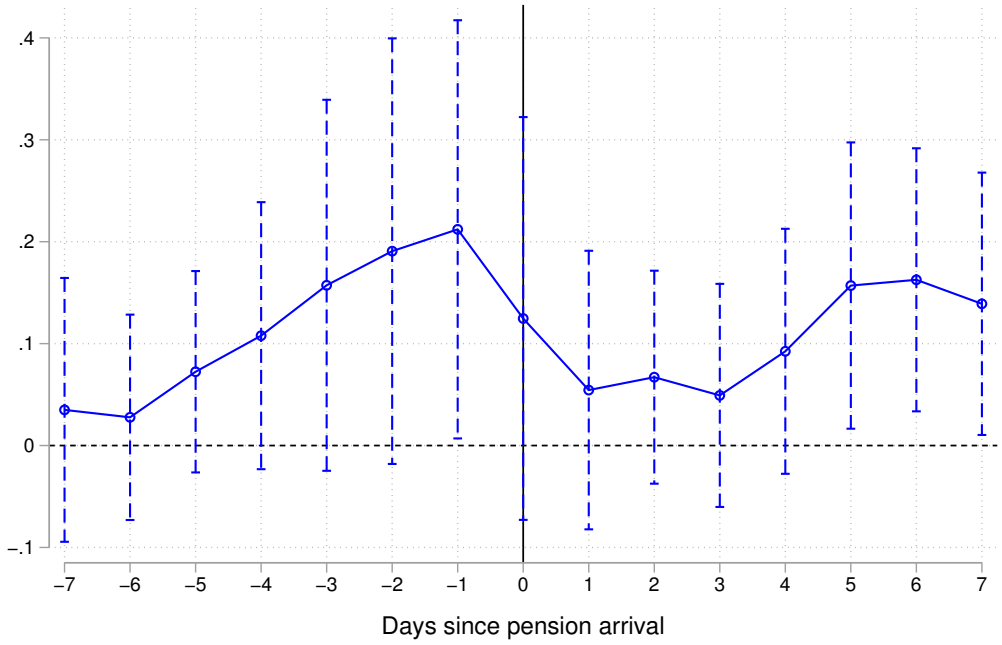
In particular, we employ two sets of approaches to identify temporary sales at the store-good level:

(i) The first approach follows Eichenbaum, Jaimovich, & Rebelo (2011) and Kehoe & Midrigan (2008, 2015), who define the most common price during a certain time interval as the regular price. For each product and date, we define the regular price as the three-month centered rolling mode (42 days on either side of the date, or 85 days in total) of the original store-good price series. In what follows we refer to this method as the "rolling mode" method. As noted in Abe & Tonogi 2010, who also use daily scanner data, using a weekly modal price coincides with the definition of bargain sales which are excluded by Statistics Japan when constructing the official Japanese CPI. We obtain qualitatively similar results when we instead use a 7-day rolling mode to characterize temporary sales.

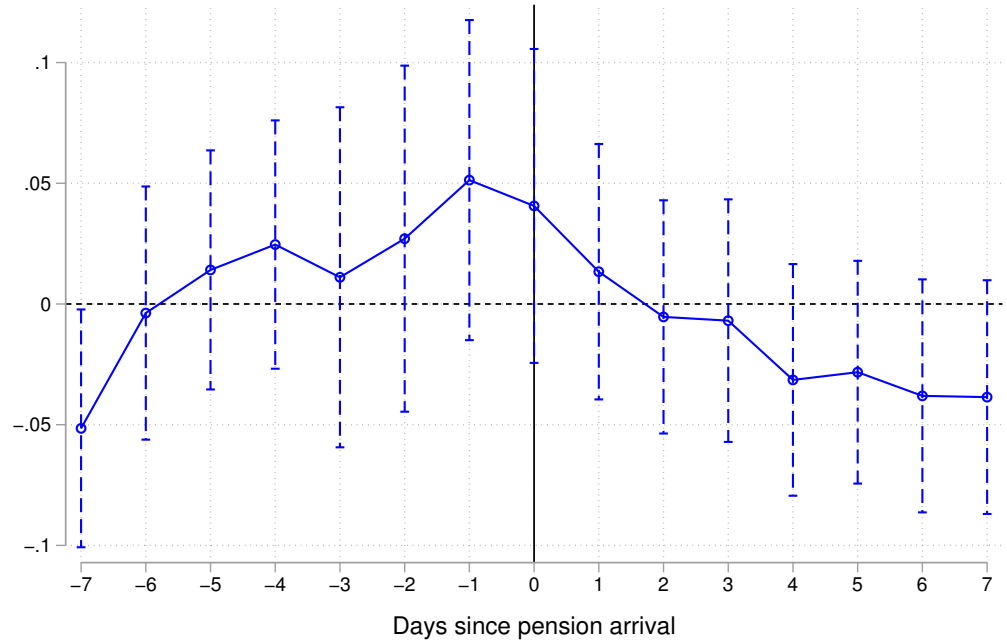
[7] We obtained a list of pension program announcements from the official JPS website: http://www.nenkin.go.jp/oshirase/taisetu/index.html. From this list of 58 announcements during our sample time period, we also tried excluding announcements that were purely clerical in nature (e.g. new application 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.2. Google Searches for "Pension Payment" around Payday

A. Sample Time Period: April 2011 – October 2014



B. Full Time Period: February 2004 – December 2018



Notes: The figure plots in blue the estimated event study coefficients $\widehat{\beta}_j$ obtained from estimating regression equation (C.5). The y-axis records the log points increase in the Google SVI for "public pension payments" around a scheduled payday. 95% confidence intervals plotted in red.

(ii) Our second approach is the sales filter proposed by Nakamura and Steinsson (2008) that classifies V-shaped patterns (i.e. symmetric dips and rebounds) in the price data as temporary sales. This V-filter approach has been shown to pick up actual temporary sales confirmed by field economists, such as in the case of the sales flag in the BLS microdata underlying the CPI (Bils & Klenow 2004). Due to the daily frequency of our data, our parameter choices deviate from the literature which has predominantly used either monthly data (Nakamura & Steinsson 2008) or weekly data like the Kilts-Nielsen retail panel (Chahrour 2011). We set the tuning parameters of the algorithm to be $L = K = J = 42$, or equivalently, we search for V-shaped pricing patterns over a 42-day (1.5 months) window as in Sudo et al. (2018). In what follows we refer to this procedure as the "V-shaped filter."

For each product within a store, we compute the frequency of temporary sales and the average rate of discount. We define the discount rate for a good at date t as $d_t = 1 - p_t/r_t$, where p_t and r_t denote, respectively, the actual price of the good and the regular price computed over the interval containing date t. To be concrete, we compute a store-level average temporary sales frequency $\overline{f}_s$ and discount rate $\overline{d}_s$ via the following expressions:

$$\overline{f}_s = \frac{1}{|T|} \sum_{t \in T} \left(\frac{1}{|K_s|} \sum_{k \in K_s} \mathbb{1}_t \{ p_{s,t,k} < r_{s,t,k} \} \right) \tag{D.1}$$

$$\overline{d}_s = \frac{1}{|T|} \sum_{t \in T} \left(\frac{1}{|K_s|} \sum_{k \in K_s} \left(1 - p_{s,t,k}/r_{s,t,k} \right) \right) \tag{D.2}$$

where $|T|$ is the number of dates contained in the set T, which may correspond to a set of either paydays or non-paydays in our sample. K_s consists of goods k which are sold in store s during our sample time period on both paydays and non-paydays.

In Figure D.1, we plot the kernel density of the store-level average frequency of sales and the average rate of discount on paydays vs. non-paydays. We restrict the sample to goods that are sold both in paydays and non-paydays (about 92% of the original sample). Overall, across all goods we find that stores are equally likely to offer sales on payday vs. non-paydays, but conditional on offering a sale, the discount is likely to be larger on paydays. Under the rolling mode filter, the average payday discount is 0.65 p.p. higher than the average non-payday discount (p-value = 0.000).

Figure D.2 displays the densities of store-level average sales frequency and discount rates for four-digit goods whose average prices are above the median price within a four-digit goods code, while Figure D.3 does the same for below-median price goods.[8] Dividing goods into above and below-median prices within a four-digit category is analogous to comparing, for example, the propensity of stores to offer discounts on a blended whiskey (below-median price) versus a single-malt whiskey (above-median price) produced by the same brand. As shown in Bils & Klenow (2001), prices within a goods category are increasing in household

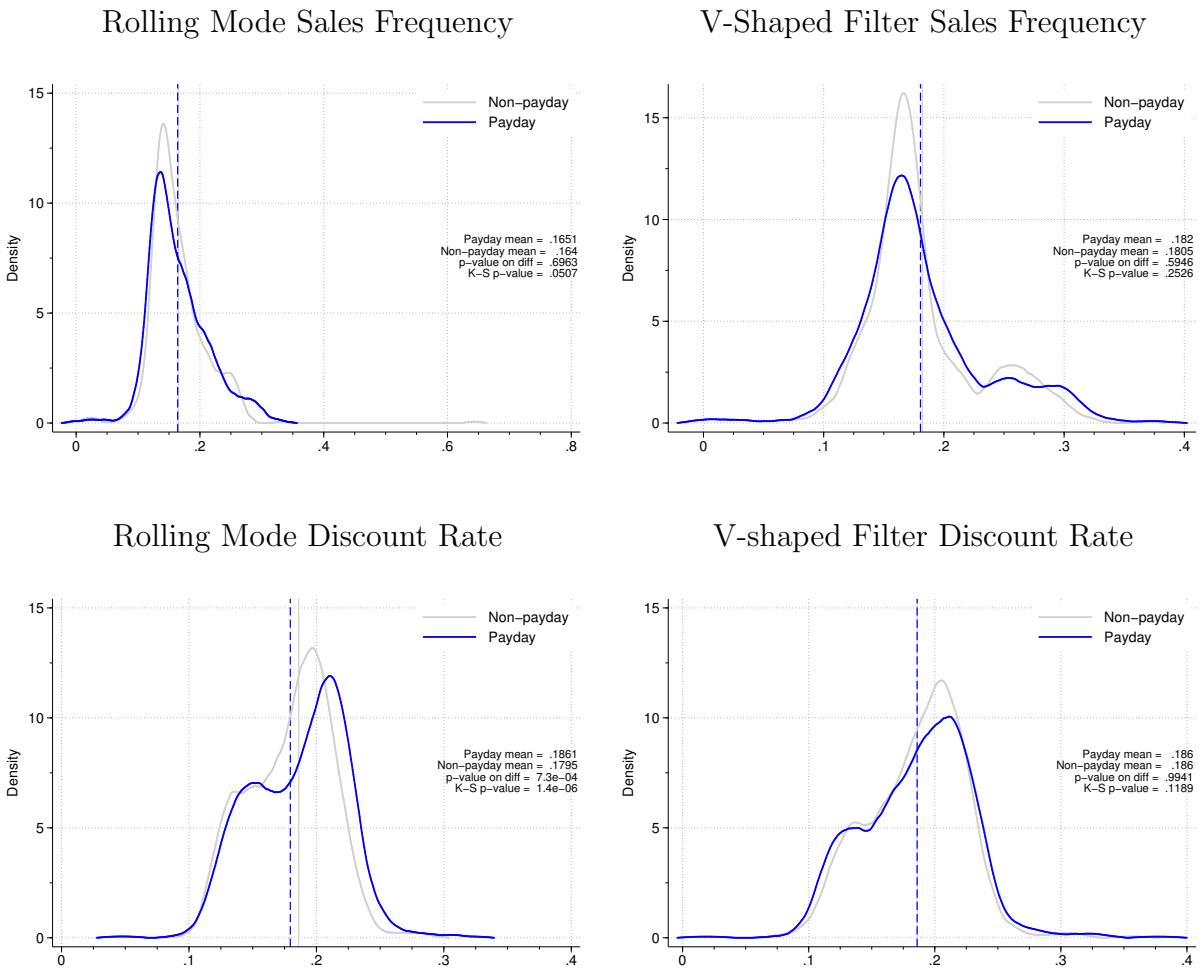
[8]See Appendix B for details on how goods categories are defined in the scanner data. For example, four-digit categories within the two-digit category of "raw meat and poultry" include "wagyu" and "eggs."

expenditures – that is, quality Engel curves slope upward. Splitting fine categories of goods into price quantiles is therefore analogous to sorting goods according to quality, or the income profile of their consumers.

We uncover stark differences in retailer pricing strategies by product quality that get partially washed out when we look at temporary sales pooled across all goods. For above-median price goods and across both sales filters, average temporary sales frequency is about 1 p.p. lower on paydays compared to non-paydays, while discounts are between 0.5 p.p. (rolling mode) to 1 p.p. (V-shape) higher on non-paydays. In other words, Figure D.2 shows that for high-quality goods, stores offer fewer sales promotions on paydays, and the ones they do offer feature less generous discounts. This pattern is completely reversed in Figure D.3 where we examine sales discounts for below-median price goods. Average temporary sales frequency is about 1.5 p.p. higher on paydays, while discounts are between 0.8 p.p. (V-shape) and 1.5 p.p. (rolling mode) higher on paydays.

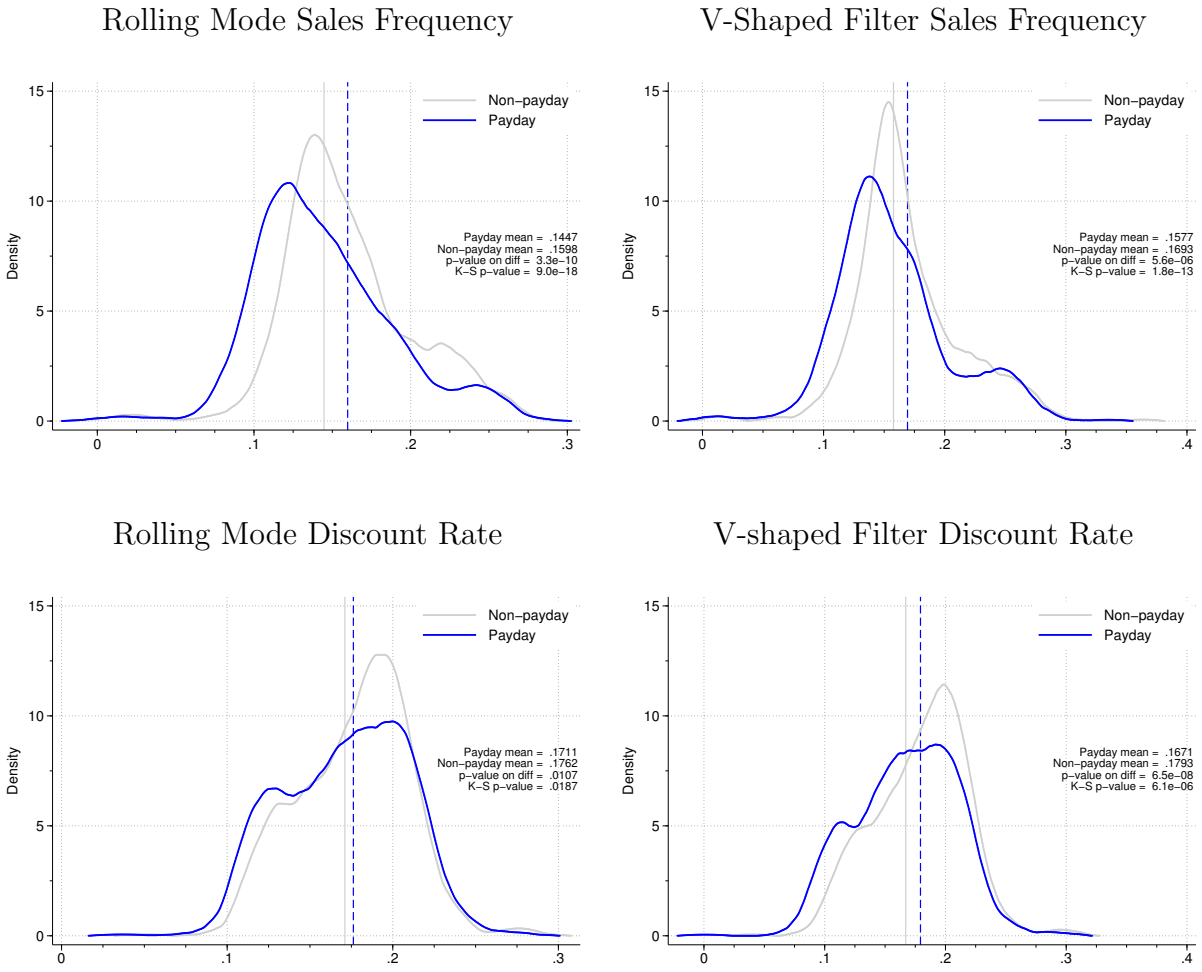
That stores offer more aggressive sales promotions for lower-quality products on paydays is consistent with the idea that price discriminating retailers can generate profit by lowering prices during periods of peak demand to attract more price-sensitive "bargain hunter" types without affecting their profits from "brand loyalists" who always buy the higher quality version of the same good (Chevalier & Kashyap 2019). In our main analysis, we show that over 90% of the observed increase in store-level average prices is due to consumer substitution towards discretionary purchases of higher-end goods, or "splurge" goods. Given that retailers keep prices on splurge goods high during paydays by forgoing temporary sales, it is unlikely that the government's choice of payment frequency would have an additional effect on intra-cycle consumption patterns through a supply-side channel. However, in Appendix A we consider some extensions of our model with monopolistic firms offering a price menu which depends on the government's payment interval choice.

FIGURE D.1. Store-level Temporary Sales Frequency and Discounts



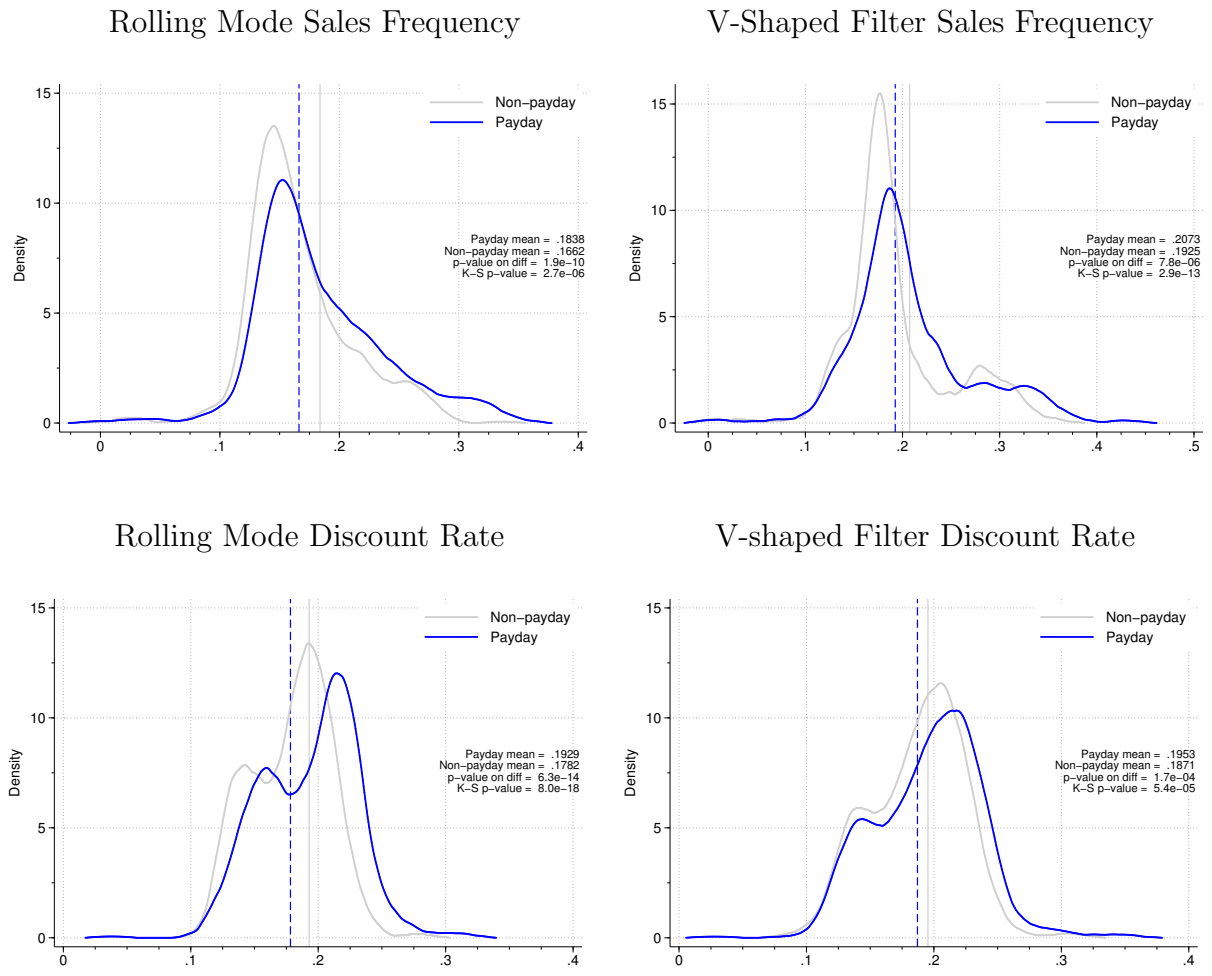
Notes: The top two panels display the kernel densities for average store-level temporary sales frequencies pooling all products sold at a store across both paydays and non-paydays. A frequency of 0.2, for example, implies the average product at a store goes on sale once every $1/0.2 = 5$ days. The bottom two panels perform the same exercise for the average store-level discount rate, where the discount rate is 1 minus the ratio of the current price to the regular price. The left-hand side panels show the frequencies and discount rates under the rolling mode filter, while the right-hand side panels show the distributions when we use the V-shaped filter to identify sales. In both algorithms we search for temporary sales over a 42-day window on either side of a calendar date. Solid grey vertical lines indicate the mean daily frequency or discount rate across stores on non-paydays, while blue dashed lines show the mean across stores on paydays. The K-S p-value shows the two-sided exact p-value from a Kolmogorov-Smirnov test of equality for the payday vs. non-payday distributions.

FIGURE D.2. Temporary Sales Frequency and Discounts on Above-Median Price Goods



Notes: The top two panels display the kernel densities for average store-level temporary sales frequencies, including only products which have an above-median average price within their four-digit goods category. A frequency of 0.2, for example, implies the average product at a store goes on sale once every $1/0.2 = 5$ days. The bottom two panels perform the same exercise for the average store-level discount rate, where the discount rate is 1 minus the ratio of the current price to the regular price. The left-hand side panels show the frequencies and discount rates under the rolling mode filter, while the right-hand side panels show the distributions when we use the V-shaped filter to identify sales. In both algorithms we search for temporary sales over a 42-day window on either side of a calendar date. Solid grey vertical lines indicate the mean daily frequency or discount rate across stores on non-paydays, while blue dashed lines show the mean across stores on paydays. The K-S p-value shows the two-sided exact p-value from a Kolmogorov-Smirnov test of equality for the payday vs. non-payday distributions.

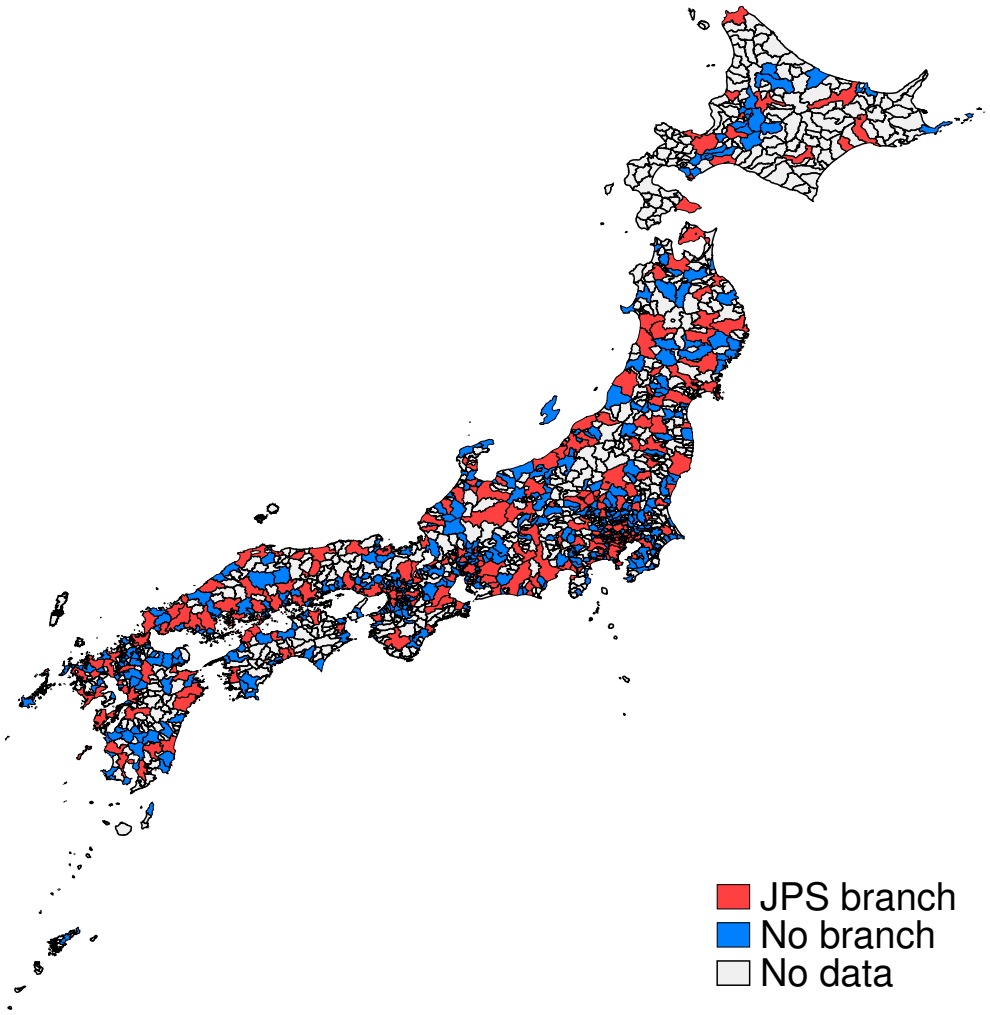
FIGURE D.3. Temporary Sales Frequency and Discounts on Below-Median Price Goods



Notes: The top two panels display the kernel densities for average store-level temporary sales frequencies, including only products which have a below-median average price within their four-digit goods category. A frequency of 0.2, for example, implies the average product at a store goes on sale once every $1/0.2 = 5$ days. The bottom two panels perform the same exercise for the average store-level discount rate, where the discount rate is 1 minus the ratio of the current price to the regular price. The left-hand side panels show the frequencies and discount rates under the rolling mode filter, while the right-hand side panels show the distributions when we use the V-shaped filter to identify sales. In both algorithms we search for temporary sales over a 42-day window on either side of a calendar date. Solid grey vertical lines indicate the mean daily frequency or discount rate across stores on non-paydays, while blue dashed lines show the mean across stores on paydays. The K-S p-value shows the two-sided exact p-value from a Kolmogorov-Smirnov test of equality for the payday vs. non-payday distributions.

E Japanese Pension System Branch Offices

FIGURE E.1. Map of Japanese Pension System Branch Office Locations



Notes: The map displays municipalities with (red) and without (blue) Japanese Pension System (JPS) offices as of 1980, when our local government balance sheet data are first widely available. We exclude from our city-level panel any local jurisdictions with gaps in public expenditure data around the 1988 reform which reduced the payment interval length from three months to two months; such cities are indicated on the map in gray ("no data"). In our analysis, we impose modern municipal boundaries defined according to the city code crosswalk available through RIETI (Kondo 2019). Our estimates of the slope of the administrative cost function for the JPS are nearly unchanged when we instead impose historical 1980 municipal boundaries to assign treatment status.

Table E.1. Summary Statistics for JPS Branch vs. Non-Branch Office Cities

	Branch (N = 239)	Non-branch (N = 424)	Difference	p-value
Log Census population	12.16	11.07	1.09	0.00
	(0.06)	(0.03)	(0.06)	
CBD population density (1000s/km^2)	7.66	5.11	2.55	0.00
	(0.46)	(0.15)	(0.38)	
Fraction population > 65 y.o. (%)	10.71	11.17	−0.46	0.09
	(0.17)	(0.18)	(0.26)	
Fraction population > 75 y.o. (%)	4.05	4.26	−0.21	0.07
	(0.07)	(0.08)	(0.12)	
$\%\Delta^{75-85}$ population > 65 y.o.	43.00	46.73	−3.73	0.02
	(0.94)	(1.12)	(1.65)	
Fraction female residents (%)	51.14	51.09	0.05	0.63
	(0.00)	(0.00)	(0.00)	
Fertility rate	2.33	2.28	0.05	0.04
	(0.02)	(0.02)	(0.03)	
Log per capita income	7.82	7.79	0.03	0.01
	(0.01)	(0.01)	(0.01)	
Labor force participation rate (%)	50.24	49.71	0.53	0.06
	(0.21)	(0.16)	(0.28)	
Unemployment rate (%)	3.48	3.08	0.40	0.00
	(0.09)	(0.06)	(0.10)	
Ratio of govt. expenditures to revenues	0.97	0.97	0.00	0.89
	(0.00)	(0.00)	(0.00)	
Log welfare spending per person > 65 y.o.	4.10	4.05	0.05	0.18
	(0.03)	(0.02)	(0.04)	
Log welfare spending per person > 75 y.o.	5.08	5.03	0.05	0.19
	(0.03)	(0.02)	(0.04)	

Notes: The table provides the mean and standard deviation, the unconditional difference in means and standard error, and p-value on the two-sided t-test of the difference between municipalities with a branch office of the Japanese Pension System (JPS) and those without a branch office. All non-growth rate variables are measured as of the pre-reform period in 1985. Population variables come from the quinquennial Census. To obtain per capita income (in 1,000s of JPY), we use the Cabinet Office local statistics for taxable income and divide by total 1985 Census population. Government expenditure ratios and unemployment rates also come from the Cabinet Office local statistics. To compute these statistics, we impose modern municipal boundaries using the historical city code crosswalk available through RIETI (Kondo 2019).

F Additional Results on Payday Responses

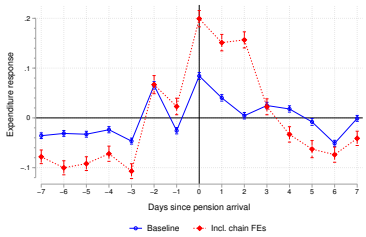
F.1 Expenditure Reponses

F.2 Spending Responses around Retirement Age

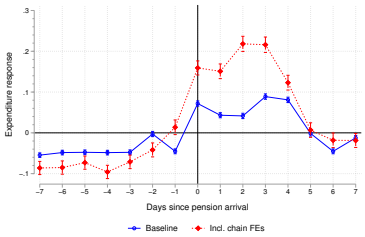
F.3 Pricing Responses

FIGURE F.1. Response of Major Subcategory Expenditures to Payday

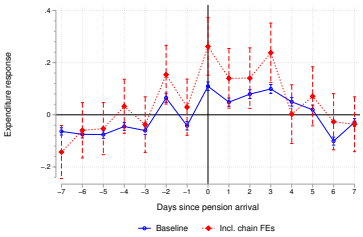
(a) Prepared Foods



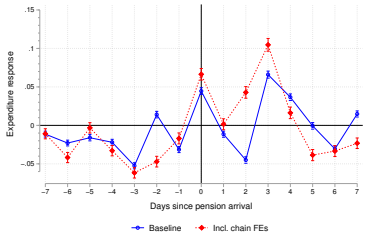
(b) Sweets/Desserts



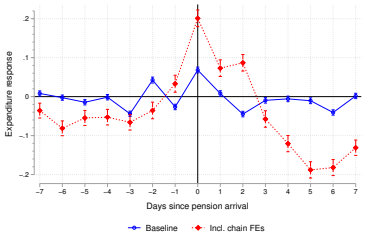
(c) Alcohol



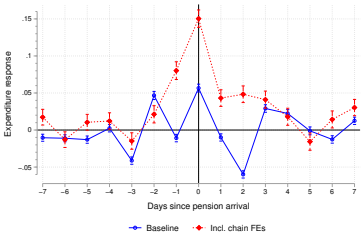
(c) Fresh Produce



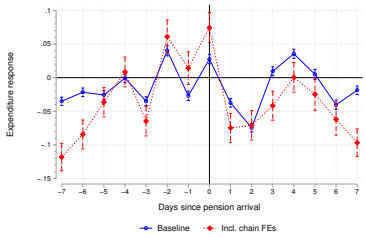
(d) Fresh Fish



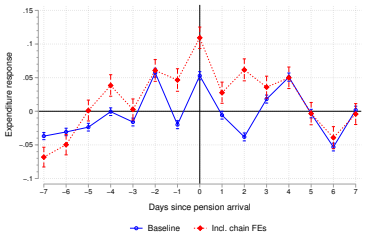
(e) Meat & Poultry



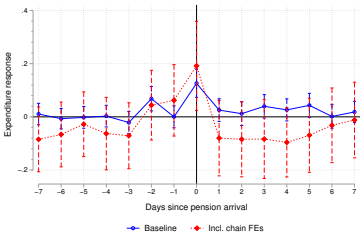
(f) Grains



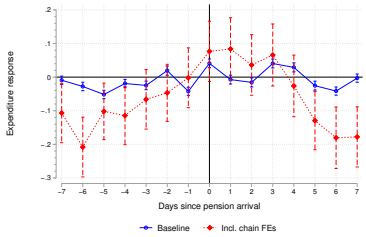
(g) Non-alcoholic beverages



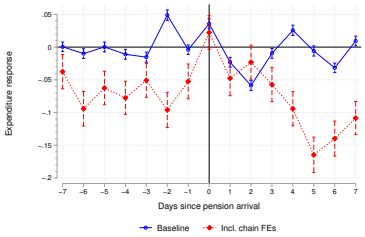
(h) Tobacco



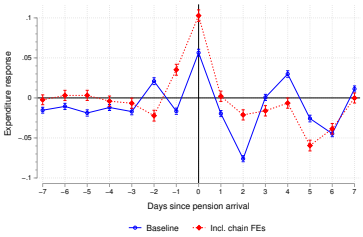
(i) Processed Fruits/Vegetables



(j) Preserved Fish



(k) Other Processed Foods



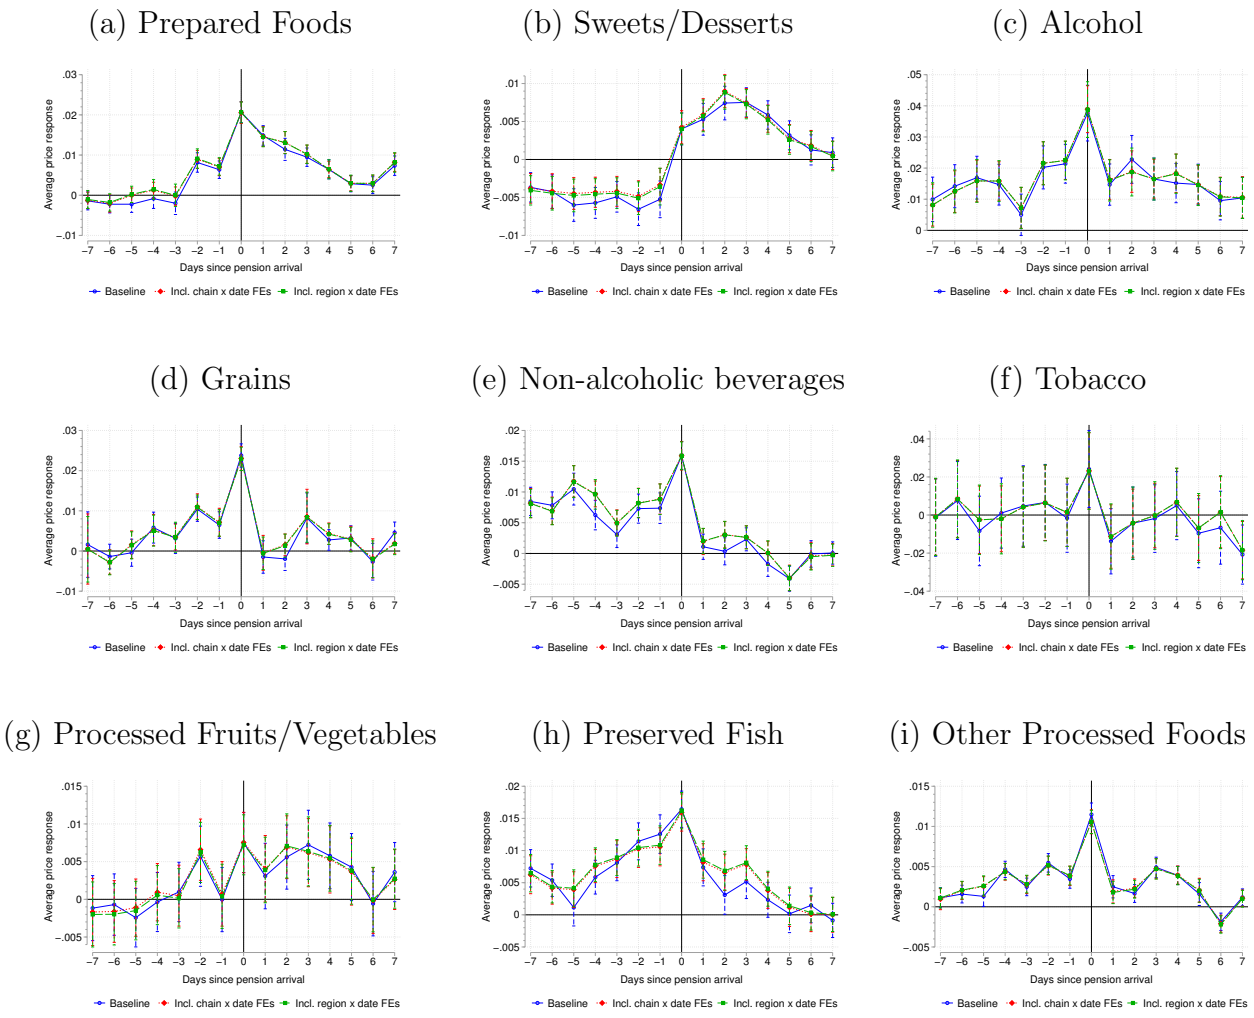
Notes: Each panel plots in blue the estimated event study coefficients $\widehat{\beta}_j$ obtained from estimating regression equation (4.1) for a major goods expenditure category over a panel of weekly shoppers which remains fixed across categories. Point estimates in red obtained from a version of equation (4.1) which includes store chain fixed effects to account for variation in spending patterns due to retailer-specific pricing decisions. The y-axis records the percent increase in expenditures relative to average daily expenditures within that category over the panel. We winsorize the top 1% of total daily expenditures within each subcategory. Bars indicate 99% confidence intervals, with standard errors obtained from clustering by shopper ID. See Appendix B for details on how we sort goods into each expenditure category.

Table F.1. Payday Expenditure Responses by Permanent Income Decile

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

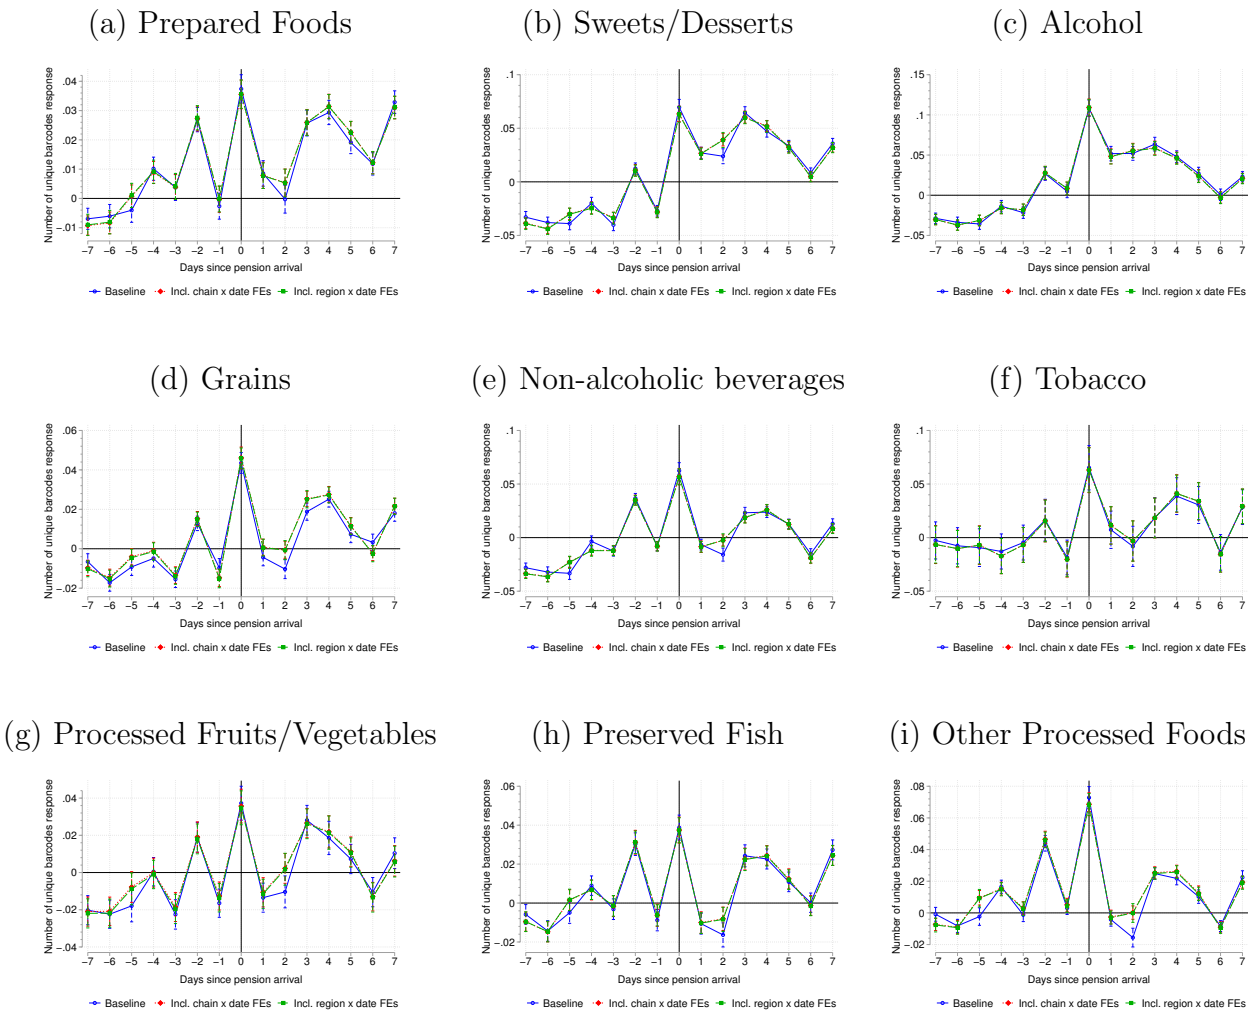
Notes: The table reports expenditure responses of regular shoppers of age 65 or older to payday, sorted by decile of average total pay cycle expenditures. Each column number refers to results for a decile of average total expenditures over bimonthly pay cycles. Each row refers to a version of our baseline specification in equation (4.1) 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. We winsorize the top 1% of daily expenditures. Robust standard errors clustered by shopper ID in parentheses. $^{***}p < 0.001, ^{**}p < 0.01, ^{*}p < 0.05, +p < 0.1$

FIGURE F.2. Response of Store-level Major Subcategory Average Prices to Payday



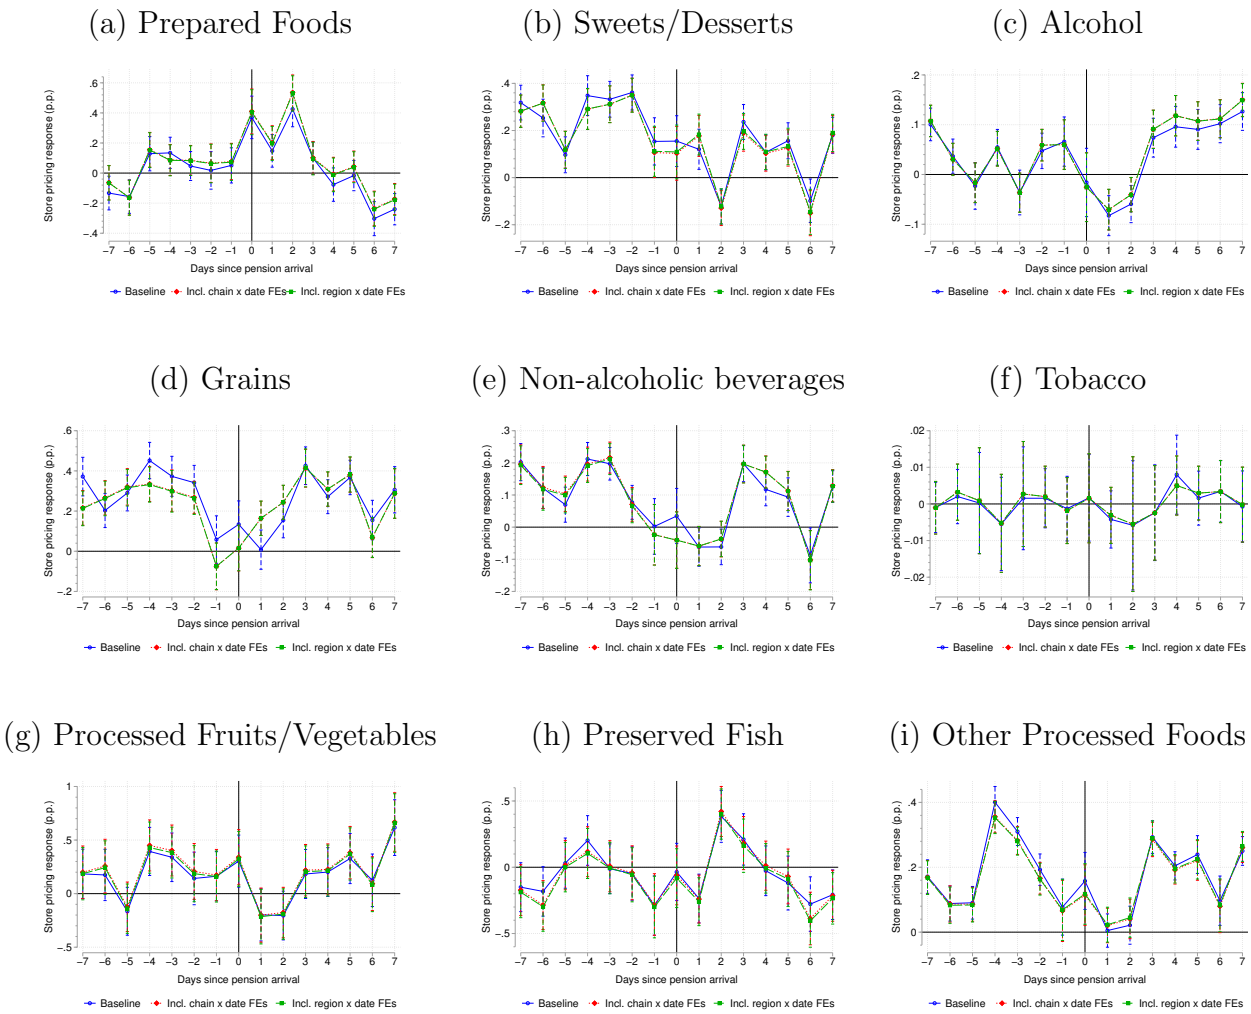
Notes: Each panel plots the event study coefficients $\widehat{\gamma}_j$ obtained from estimating regression (5.2) for spending within one of the one-digit goods categories defined in Appendix B. The y-axis records the percent increase in store-level prices around payday, using the average price index defined in (5.1). Point estimates in red obtained from augmenting (5.2) to include store chain $\times$ month-year fixed effects, while estimates in green result from including Census region $\times$ month-year fixed effects. Bars indicate 99% confidence intervals, with standard errors clustered at the store level.

FIGURE F.3. Store-level Variety Responses within Each Major Goods Subcategory



Notes: Each panel plots the event study coefficients $\widehat{\gamma}_j$ obtained from estimating regression (5.2) where we use as the dependent variable the log of the daily number of goods sold (identified by unique barcodes), within one of the one-digit goods categories defined in Appendix B, and in a store relative to the average number sold over the sample time period. Point estimates in red obtained from augmenting (5.2) to include store chain × month-year fixed effects, while estimates in green result from including Census region × month-year fixed effects. Bars indicate 99% confidence intervals, with standard errors clustered at the store level.

FIGURE F.4. Store Pricing Responses aournd Payday by Major Goods Subcategory



Notes: We construct store-level price indices as in equation (5.1) and then estimate regression (5.2) to obtain estimates $\widehat{\gamma}_{1,j}$. We then construct a counterfactual last price index, using the last observed price for each barcode in the odd month preceding the payday event, and rerun (5.2) on the exact same sample using this price index as the dependent variable to obtain $\widehat{\gamma}_{2,j}$. Each panel in the figure repeats this procedure restricting to spending within one of the one-digit goods categories defined in Appendix B. The figure plots the point estimates and 99% confidence interval bars for the differences $\widehat{\gamma}_{1,j} - \widehat{\gamma}_{2,j}$, with standard errors in each regression clustered at the store level. Point estimates in red obtained from augmenting each version of (5.2) to include store chain $\times$ month-year fixed effects, while estimates in green result from including Census region $\times$ month-year fixed effects. The differences capture the p.p. increase in prices on payday that cannot be explained by changes in the composition of goods purchased.

Appendix References

Abe, N. & A. Tonogi (2010): "Micro and Macro Price Dynamics in Daily Data," *Journal of Monetary Economics*, 57(6): 716-728.

Aguiar, M. & E. Hurst (2007): "Life-Cycle Prices and Production," *American Economic Review*, 97(5): 1533-1559.

Baugh, B. & J. Wang (2018): "When Is it Hard to Make Ends Meet?, *mimeo*, University of Illinois at Urbana-Champaign.

Bils, M. & P.J. Klenow (2001): "Quantifying Quality Growth," *American Economic Review*, 91(4): 1006-1030.

Bils, M. & P.J. Klenow (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112(5): 947-985.

Chahrour, R.A. (2011): "Sales and Price Spikes in Retail Scanner Data," *Economics Letters*, 110(2): 143-146.

Eichenbaum, M., N. Jaimovich, S. Rebelo (2011): "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review*, 101(1): 234-262.

Kehoe, P.J. & V. Midrigan (2008): "Temporary Price Changes and the Real Effects of Monetary Policy," NBER Working Paper, No. 14392.

Kehoe, P.J. & V. Midrigan (2015): "Prices Are Sticky After All," *Journal of Monetary Economics*, 75: 35-53.

Kondo, K. (2019): "Municipality-level Panel Data and Municipal Mergers in Japan," RIETI Data Management Project, https://www.rieti.go.jp/en/publications/summary/19030013.html.

Nakamura, E. & J. Steinsson (2008): "Five Facts About Prices: A Reevaluation of Menu Cost Models: Supplementary Material," *Quarterly Journal of Economics*, 123(4): 1415-1464.

Sudo, N., K. Ueda, K. Watanabe, & T. Watanabe (2018): "Working Less and Bargain Hunting More: Macroimplications of Sales during Japan's Lost Decades," *Journal of Money, Credit and Banking*, 50(2-3): 449-478.