Explaining Consumption Excess Sensitivity with Near-Rationality: Evidence from Large Predetermined Payments

Lorenz Kueng*
Northwestern University and NBER†

November, 2015

Abstract

Using new transaction data I show that consumption is excessively sensitive to large, predetermined, regular, and salient payments from the Alaska Permanent Fund, with a large average marginal propensity to consume (MPC) of 30% for nondurables and services and 70% for total expenditures. This deviation from the standard inter-temporal consumption model is concentrated among households for whom the loss from failing to smooth consumption is small in terms of equivalent variation. In particular, the MPC is *increasing* in household income but *decreasing* in the size of the loss. As a result, statistically significant excess sensitivity in response to these large payments is consistent with households following near-rational alternative consumption plans. For macroeconomic policies, such as an economic stimulus program, these near-rational alternatives might be the more relevant behavior than the standard consumption model.

JEL Classification: D12, D91, E21, H31.

Keywords: consumption, excess sensitivity, MPC heterogeneity, near-rationality.

*I thank Brian Melzer, David Card, Jonathan Parker and seminar participants at Northwestern and Wisconsin-Milwaukee for useful comments. I am grateful to Laura Achee at the APFC for providing archival data and for patiently answering many questions about the fund. Ryan Pfirrmann-Powell at the BLS provided expert assistance with the confidential CE data. All remaining errors are mine.

†Department of Finance, Northwestern University's Kellogg School of Management, 2001 Sheridan Road, Evanston, IL 60208, and National Bureau of Economic Research. Phone: 510-860-7544. Fax: 510-255-5013. E-mail: l-kueng@kellogg.northwestern.edu.

Aggregate consumption constitutes more than two-thirds of gross domestic product in almost all developed economies and is thus a major component of the business cycle. Understanding how consumers respond to cash flows is therefore fundamental for designing economic stabilization programs, such as active fiscal policies and automatic stabilizers. The permanent income life-cycle hypothesis—the workhorse model of inter-temporal consumption choices used in one form or another in most of the literature—predicts that in the absence of financial frictions, households will adjust their consumption plans only when they receive new information about their life-time resources. Under this benchmark model, households' optimal consumption plans should smooth out predictable changes in cash flows.

Not surprisingly given the importance of the model, there is an enormous empirical literature that tests this prediction, typically concluding that household consumption systematically reacts to predictable cash flows, exhibiting so-called excess sensitivity. A major limitation of most excess sensitivity tests is that they typically do not point to the reason why consumers do not smooth consumption as much as the model predicts. Moreover, it is often unclear whether the statistical rejection of the theory's sharp prediction is also economically relevant for the households, as the loss from failing to smooth out the cash flow shocks used in these tests is often small in terms of equivalent variation.

However, even if the loss from failing to fully smooth consumption is small for most households, this does not mean that excess sensitivity and similar near-rational alternatives are not relevant for macroeconomics, since those small deviations from the standard model are correlated across households and can therefore add up to large aggregate responses. In fact, many active macroeconomic policies have relatively small *direct* impacts on most households' budgets, such as the widely studied fiscal stimulus payments in 2001 and 2008 for example, but nevertheless provide substantial economic stimulus.

This paper combines new transaction-level financial data from a personal finance website with the repeated quasi-natural experiments provided by the annual payments from the Alaska Permanent Fund, the state's wealth fund, to shed new light on this question. The fund invests the proceeds from the state's oil revenue in broadly diversified financial assets and uses the assets' income stream to pay out the annual Permanent Fund Dividend (PFD).

The experiment has several characteristics that make it a particularly useful research design for testing the basic theory and, importantly, for assessing the economic relevance of potential deviations from the standard model for each household. First, the cash flows are *large* such that deviating from consumption smoothing is potentially costly for a significant fraction of households. Second, the payments are fully *predetermined* and highly predictable. The exact amount of the dividend is announced well before it is paid out in October and is widely discussed in the local media throughout the year. The dividend is therefore very *salient* to consumers. Third, the payments occur *regularly every year*, further limiting surprise as an explanation of excess

sensitivity. Finally, the cash flows are *lump-sum payments* to every person in the household, including children, and are thus relatively more important for lower-income households and for larger families as a fraction of annual income, thereby providing cross-sectional variation in the potential cost of failing to smooth consumption.

Despite the favorable nature of the research design for the standard model, the average household deviates from the model's predictions, exhibiting large and statistically significant marginal propensities to consume (MPC) out of those large, anticipated, regular, salient, and exogenous lump-sum payments.[1] I document the deviations from the standard model using both non-parametric methods and the standard regression framework used in the literature. The average MPC is 30% for nondurables and services after one quarter, and 70% for total expenditures.

Unlike in most previous studies, deviating from the model's predictions by spending the entire income upon arrival instead of smoothing it throughout the year is potentially economically significant in this setting. I quantify the economic power of the quasi-experiment for rejecting the basic permanent-income life-cycle model against near-rational alternatives. The *potential* ex-ante money-metric loss of deviating from the optimal consumption plan (c^*), which fully smoothes consumption during the period, and instead following the near-rational alternative ($\tilde{c}$) that spends the entire dividend upon arrival while also satisfying the inter-temporal budget constraint, is monotonically increasing in the relative size of the payments ($\frac{PFD}{c_T}$) as a fraction of total consumption during the period (c_T),

$$Loss(\tilde{c}, c^*) \approx \frac{\gamma}{2} \cdot (T-1) \cdot \left(\frac{PFD}{c_T}\right)^2,$$

where T is the time distance between announcement or news about the dividend amount and the actual dividend payments, and γ measures the curvature of the time-additive sub-utility function.

Importantly, this potential loss can be calculated *ex-ante* before the arrival of the payment and can therefore be used as a predictor of excess sensitivity, in addition to other predictors such as credit constraints that have previously been used in the literature to explain excess sensitivity. Sorting households by potential loss, I find that this new predictor explains most of the excess sensitivity of higher-income households, consistent with near-rational deviations from the permanent-income life-cycle theory. For instance, moving from the lowest loss quintile to the highest quintile reduces the MPC from 80% to 15%. The predictive power of the potential loss is robust to conditioning on liquid assets, income, education, and age, all of which have previously been proposed to explain excess sensitivity.

[1] More precisely, I estimate the marginal propensity to spend, which might be different from the MPC if goods and services are not fully consumed during the period they are purchased, a point I discuss below.

While it might be obvious that the relative size of the payments should have predictive power in explaining excess sensitivity, near-rationality also has implications that are more surprising. In particular, since the dividend income is smaller share of higher-income households' budget, MPCs that are *increasing* in household income are therefore consistent with near-rationality. I show that the MPCs are indeed *larger* for higher-income households than for lower-income households, and this relationship is monotonic, even after conditioning on liquid wealth. For instance, moving from the lowest to the highest income quintile increases the MPC from 10% to 65%. These differences in MPCs across income quintiles are highly statistically significant thanks to the high quality of the income data, which is based on the households' financial accounts.

The result that the MPC is increasing in income and decreasing in the size of the welfare loss is particularly striking since it provides a new source of heterogeneity in MPCs *in addition* to traditional liquidity constraints. Specifically, I show that lower-income households with very low amounts of liquid wealth also have high MPCs, higher than similar households with sufficient liquid assets. Importantly, this result holds even after conditioning on income or the relative size of the payment. Hence, low utility losses from suboptimal consumption plans predict high MPCs for higher-income households, while low levels of liquid wealth continue to predict high MPCs for lower-income households. MPCs in turn are lowest for lower-income and middle-income households who have sufficient amounts of liquid assets.

The new transaction data have several additional advantages over traditional expenditure surveys that help to address a number of concerns that naturally arise when interpreting the spending responses to the dividend payments. For instance, since the website automatically tracks a household's income and spending after it has linked its credit card and other financial accounts, household expenditures recorded in the transaction data are measured with substantially less error than in surveys, particularly for disaggregated spending categories. This is important since many nondurables and services do have a durable component, especially when looking at frequencies higher than annual, so that spending does not necessarily reflect consumption. One is therefore concerned that households might time the purchase of such goods to the arrival of the dividend cash flows while spreading the consumption of the goods evenly over the year, as predicted by the standard model. Focusing on more disaggregated data, however, shows that spending significantly increases across many categories in response to the dividend, including strictly nondurable categories that have previously been used to address this issue, such as food and dining, and in particular groceries.

An additional strength of the transaction data is the higher frequency compared to survey data, which can be used to better trace out the dynamics of the excess sensitivity and test for anticipation effects. I document that households adjust their spending on nondurables and services within the first three months after receiving the lump-sum payments, and most of the

response occurs in the initial month. The cumulative effect of this additional spending is stable after the first three months and there is no evidence of reversal due to intertemporal consumption shifting. The lack of a decline in spending on nondurables in later periods is additional evidence that potential timing differences between spending and consumption do not explain the large excess sensitivity. Moreover, households do not move their nondurable spending forward to the months prior to the payments, even households with sufficient liquid wealth that do not have to borrow against the future dividend. Hence, even though the payments are preannounced, the spending response resembles that of a typical event study.

In order to provide external validity of the results, I compare the average MPC based on the new transaction data to similar estimates obtained using the Consumer Expenditure Survey (CE), which spans the entire period since the first dividend was paid out in 1982 up to 2013, but covers much fewer Alaskan households per period. The spending response to the dividend payments is similar after accounting for differences in sample composition and the fraction of Alaskans that do not receive the dividend.

Finally, I compare these findings to Hsieh (2003), which was the first study to examine spending responses to the Permanent Fund Dividend. That study, which uses the CE sample from 1980 to 2001 finds a small and insignificant response. The main specification regresses changes in spending on the dividend payments normalized by current family income. Income in the CE survey, however, suffers from substantial measurement error, which attenuates the estimated spending response. In a companion paper (Kueng (2015b)) I replicate the small and insignificant spending responses reported in Hsieh (2003) and show that one can use total expenditures—which are more precisely measured in the CE—to instrument for current income, resulting in a statistically significant spending response that is quantitatively similar to the one reported in this paper.

The paper is organized as follows. Section 1 derives the approximate economic loss of following a near-rational alternative. Section 2 describes the micro data and the Permanent Fund Dividend. Section 3 shows non-parametric and parametric evidence of excess sensitivity. Section 4 uses the economic loss statistic to predict heterogeneity in MPCs across higher-income households. Section 5 shows that credit constraints help predict high MPCs for lower-income households who have low levels of liquid assets. Section 6 performs a thorough robustness analysis of the excess sensitivity results and extends the analysis along several dimensions, including the external validity check using the Consumer Expenditure Survey. Section 7 concludes.

1 Near-Rationality and Excesss Sensitivity

Many studies have used quasi-experiments to document excess sensitivity of household consumption; see e.g., Jappelli and Pistaferri (2010) for a recent survey. The insight that excess sensitivity in consumption could be related to near-rational behavior goes back at least to

Cochrane (1989) who studies early excess sensitivity tests based on aggregate time series, while Browning and Crossley (2001) and Fuchs-Schuendeln and Hassan (2015) focus on more recent tests based on micro data.

The new micro data and the properties of the Alaska Permanent Fund Dividend are ideal for testing the predictions of near-rational behavior for consumer spending out of predetermined cash flows. Since the dividend payments are large, regular and hence salient, many alternative explanations that have previously been proposed to explain excess sensitivity can be ruled out.

For instance, Alaskans are generally not surprised by the cash flows from the dividend, both because they are well communicated in advance and because they occur every year. The high predictably of the dividend well in advance also diminishes motives for additional precautionary saving. The excess sensitivity documented in this paper can also not be the rational response to changes in expected dividends, both because changes in expected dividends are much smaller than the actual cash flows and because forecast errors should on average be zero and hence not systematically lead to a positive spending response. And while I document that lower-income households with few or no liquid assets do respond more than similar households with sufficient liquid assets, the typical household in the sample has high income with lots of liquid wealth and nevertheless strongly responds to the dividend payments. I show that variation in the potential loss from near-rational behavior is highly predictive of variation in MPCs among higher-income households, and hence the degree of excess sensitivity, similar to the predictive power of liquidity constraints for lower-income households.

This paper is therefore one of the first to provide direct evidence at the household level that near-rationality predicts excess sensitivity using a single source of income, i.e., using variation *within* the same research design. Potential wealth-equivalent losses from near-rational behavior in turn are proportional to the relative size of the dividend payments as a fraction of household consumption as shown in the next section and is thus a sufficient statistic that can be used to predict excess sensitivity ex ante. The only other studies I am aware of that analyze whether payment size can explain excess sensitivity in the cross-section using a single income source are Kreinin (1961), Souleles (1999), and Scholnick (2013), all of which use a quadratic function of the payments and find mixed or inconclusive results due to a lack of statistical power.[2] The loss from suboptimal behavior, however, is a monotone function of the *relative* size of the payment scaled by household consumption, as shown above, instead of the payment level. This distinction is not only conceptually important, but it turns out to be also empirically relevant. Using the unscaled squared size of the dividend instead of the relative size of the cash flows, I find that the coefficient on the quadratic term is statistically insignificant and also economically small, while the linear term—the average excess sensitivity—is unaffected by

[2] Parker (1999) also analyzes near-rationality as an explanation for excess sensitivity, but instead focuses on differences in MPCs across different types of goods with different degrees of durability, as they imply different costs from failing to smooth spending.

adding the quadratic term and remains economically and statistically significant.

1.1 Loss from Sub-Optimal Consumption Plans

Denote the optimal consumption plan given wealth W and prices p (i.e., interest rates) by c^*, such that $c^* = \arg\max_c\{U(c) \text{ s.t. } p'c \leq W\}$, with $p'c^* = \sum_t R^{-t}c_t^* = W$ and life-time utility $U(c) = \sum_t \delta^t u(c_t)$. Following Gabaix and Laibson (2002), consider a deviation $\tilde{c}$ from this optimum, which also has to satisfy the inter-temporal budget constraint, $p'\tilde{c} = W$. Using the envelope theorem, i.e., combining the first-order conditions to simplify the second-order approximation of $U(c)$ around c^*, we obtain

$$U(c^*) - U(\tilde{c}) \approx -\frac{1}{2}\sum_t \delta^t \cdot \frac{\partial^2 u(c_t^*)}{\partial c^2} \cdot (c_t^*)^2 \cdot \left(\frac{\tilde{c}_t - c_t^*}{c_t^*}\right)^2. \tag{1}$$

The first-order term $\frac{\partial U(c^*)}{\partial c}'(\tilde{c} - c^*)$ is zero because the first-order conditions imply $\frac{\partial U(c^*)}{\partial c} = \lambda \cdot p$ and $p'(\tilde{c} - c^*) = 0$, since both consumption plans satisfy the intertemporal budget constraint.

To quantify the money value of the loss due to local deviations from the optimal plan, we can calculate the wealth compensation necessary to keep the household at the same utility level under the suboptimal plan $\tilde{c}$ as under c^*. We know that under standard preferences, a proportional change in wealth leads to a proportional change in the optimal consumption profile, i.e., $d\ln(c_s^*) = d\ln(W) \approx \frac{\tilde{W}-W}{W} = \frac{\Delta W}{W}$ $\forall s$. Taking a first order approximation of the value function around the initial optimum, c^*, we obtain

$$U(c^*) - U(c^*(\tilde{W})) \approx -\frac{\Delta W}{W}\sum_t \delta^t \cdot \frac{\partial u(c_t^*)}{\partial c} \cdot c_t^*, \tag{2}$$

where $c^*(\tilde{W})$ is the optimal consumption plan given alternative wealth $\tilde{W}$ in a neighborhood of W. Combining (1) and (2), evaluating $U(c^*(\tilde{W}))$ at $U(\tilde{c})$, and assuming an iso-elastic flow utility $u(x) = x^{1-\gamma}/(1-\gamma)$ yields the money metric used to calculate the relative loss of the sub-optimal plan,

$$Loss(\tilde{c}, c^*) \equiv -\frac{\Delta W}{W} \approx \frac{\gamma}{2}\sum_t \omega_t \cdot \left(\frac{\tilde{c}_t - c_t^*}{c_t^*}\right)^2, \tag{3}$$

with "utility-annuity" weights $\omega_t = \delta^t \frac{\partial u(c_t^*)}{\partial c} c_t^* / \sum_j \delta^j \frac{\partial u(c_j^*)}{\partial c} c_j^* = \frac{\delta^t u(c_t^*)}{U(c^*)}$, because $\frac{\partial u(c)}{\partial c} c = (1-\gamma)u(c)$.

To apply this expression of the loss in practice, and in particular to the setting of the Alaska Permanent Fund Dividend, we need to specify the alternative behavior. Suppose the household deviates from the optimal consumption plan by responding to the Permanent Fund Dividend payments instead of the announcements or news T periods in advance (e.g., four quarters), but otherwise fully optimizes, i.e., follows a near-rational alternative $\tilde{c}$ as defined by Akerlof

and Yellen (1985). Divide the household's finite horizon H in intervals of length T (e.g., four quarters), and to simplify, assume no discounting and zero interest rates ($\delta = R = 1$), such that $c_t^* = c^*$. This alternative consumption plan that does not fully smooth the dividend but instead spends the amount Δ of it in the period in which it is paid out, is defined as

$$\tilde{c}_t = \begin{cases} \tilde{c} & = c^* - \frac{\Delta}{T} & \text{in periods without dividend payments,} \\ \tilde{c} + \Delta & = c^* + (1 - \frac{1}{T})\Delta & \text{in periods with dividend payments.} \end{cases}$$

The alternative plan $\tilde{c}$ is related to the optimal plan c^* by the fact that both have to satisfy the intertemporal budget constraint, hence $\sum_t \tilde{c}_t = \sum_t c_t^* = H \cdot c^*$.

Δ is the degree of excess sensitivity, which corresponds to $\beta \cdot PFD$ in the regression below, i.e., the MPC times the dividend amount. Consumption is higher than the optimum when the dividend is paid out, but in turn has to be lower during the other $T-1$ periods in the interval in order to satisfy the budget constraint. Therefore, relative deviations from the optimal plan are given by

$$\frac{\tilde{c}_t - c_t^*}{c_t^*} = \begin{cases} -\frac{\Delta}{Tc^*} & \text{in periods without dividend payments,} \\ (T-1)\frac{\Delta}{Tc^*} & \text{in periods with dividend payments.} \end{cases}$$

$\frac{\Delta}{Tc^*}$ is the degree of excess sensitivity of consumption plan $\tilde{c}$, expressed as a fraction of total consumption during the interval of length T. Hence, the potential loss from fully spending the dividend payments in the fourth quarter instead of the first quarter, (i.e., having an MPC of 1, or a $\beta = 1$ in the regression below), is

$$Loss(\tilde{c}, c^*) \approx \frac{\gamma}{2} \cdot (T-1) \cdot \left(\frac{PFD}{c_T}\right)^2, \tag{4}$$

where $\frac{PFD}{c_T}$ is the relative size of the dividend as a fraction of total consumption during the interval, $c_T = Tc^*$. This loss can be calculated ex ante as it does not depend on any behavioral response to the dividend. The actual ex-post loss on the other hand depends on the household's degree of excess sensitivity, which can be estimated by regressing changes in spending on changes in dividend income for different subsamples of households with similar potential ex-ante losses.

2 Data and Experiment

The strength of the analysis builds both on the high quality of the new micro data on expenditures and income—in particular for higher-income households—, and the large dividend payments that provide substantial cross-sectional variation in the size of the potential losses from exhibiting excess sensitivity.

2.1 New Transaction Micro Data from Financial Accounts

The main analysis uses new transaction data from accounts at a large personal finance website (PFW) from 2010 to 2014. The micro data is at the user account level, which I will refer to as the household.[3] Households can link their credit card accounts, bank accounts, brokerage accounts and any other major account related to their balance sheet to their online account, giving them a systematic overview of their personal finances. The data and its advantages and shortcomings relative to previous data sources are explained in more detail in Baker (2014). This paper complements the analysis of the data quality in Baker (2014) in an important way by also implementing the same research design in the Consumer Expenditure Survey (discussed below), which is the standard data source used in previous studies. Therefore, the fact that I find similar results using both data sets—after accounting for differences in sample composition and the fraction of Alaskans that do not receive the dividend—provides an external validity check of this new data source.

I restrict the sample of Alaskan households in the PFW data to those that receive dividend payments via direct deposit within the first two days. This covers 97% of all dividend receipts identified based on the transaction label. I further restrict the analysis to households that I observe receiving the dividend since I do not know whether other Alaskan households are not receiving the dividend or whether I do not observe their dividend in the data. The direct deposit recipients make up about 80% of all self-identified Alaskans as inferred from the self-reported ZIP codes. This number is consistent with aggregate take-up statistics based on data from the Alaska Permanent Fund Dividend Division. Since the dividend amount is sensitive to the household size I drop observations with self-reported family sizes above 8 or that receive more than 7 dividend checks, which corresponds to the top 1 percent in both cases. Accounts where the absolute difference between the number of dividends and the self-reported family size is larger than four in any period are also dropped.

I supplement this sample of 1,380 Alaskan households with a random sample of 2,191 households from the State of Washington as a control group. Spending on nondurable goods and service as well as durables are defined to map to the National Income and Product Accounts as closely as possible, and the mapping is provided in Table A.1 of the appendix. To estimate the total impact of the payments on household spending I also include other expenditures (e.g., mortgage and rent, car loan payments, uncategorized expenditures) as well as durable purchases that are paid for with a credit card (e.g., clothes, newspaper and magazines, electronics and

[3] More precisely, the concept of a user account is closer to the concept of a family as defined by the U.S. Census Bureau: "A family consists of two or more people (one of whom is the householder) related by birth, marriage, or adoption residing in the same housing unit. A household consists of all people who occupy a housing unit regardless of relationship. A household may consist of a person living alone or multiple unrelated individuals or families living together." The concept of a family is also closer to the "consumer unit" concept used by the Consumer Expenditure Survey (CE). Nevertheless, following convention in the literature I will refer to both user accounts in the PFW and consumer units in the CE as "households."

software) and hence are typically smaller than say car purchases. These additional categories together with spending on nondurables and services define total expenditures.

Table 1 provides summary statistics of the main variables used in the analysis. All nominal amounts are expressed in dollars of 2014 using the CPI for Alaska and the U.S., respectively. The average size of the dividend per family equals $2,318 and is therefore much larger than the one-time tax rebate of $300 to $600 per households in 2001, which has been studied extensively in the literature. Hence, even in this sample of higher-income households, the annual dividend still represents about 5% of the annual budget, either measured as a fraction of annual income or consumption. Moreover, contrary to the tax rebate, the dividend is paid regularly once every year.

Households in Alaska and Washington are very similar along most other dimensions, including income, demographics, and expenditures. While average household income in the data is high, the median income of $76,219 in the subsample Alaskan households is similar to median household income based on the American Community Survey of 2013 for Alaska ($70,760). The average households in Washington has a higher total balance of savings and checking accounts and taxable brokerage accounts than the average Alaskan household, although the latter still has much more liquid wealth in terms of total bank balances than the average Alaskan household in the Consumer Expenditure Surveys for example ($40,000 versus $13,000).

2.2 The Permanent Fund Dividend

Since 1977, the State of Alaska invests the royalty income it receives from the oil extraction in the state-owned North Slope region in a sovereign wealth fund called the Permanent Fund. This fund, which is managed by the Alaska Permanent Fund Corporation (APFC), has grown considerably over time and had a market value of $52.8 billion as of November 2015; see Goldsmith (2001) for a historical account of the fund. The fund's assets are broadly diversified in domestic and international financial and real assets so that the cash flows generated by the fund are unaffected by local economic conditions. At the end of each fiscal year on June 30, roughly 10% of the fund's generated cash flows over the current and four previous fiscal years is set aside to be paid out by the Permanent Fund Dividend Division based on a public formula set in state law.[4] The funds are paid out to every person who has been a resident of Alaska for the previous year and indicates an intention to remain an Alaskan resident. The rest of the fund's income is typically reinvested in the fund, although the legislature has in principle the authority to use the fund's remaining earnings for any public purpose. Previous attempts by politicians to appropriate more earnings for government funding have resulted in significant

[4] The public formula for the dividend distribution is $\frac{1}{2} \times 21\% \times (\sum_{s=t-4}^{t} SNI_s - \text{Adjustments}_t)$, where SNI is the fund's statutory net income in the current ($s = t$) and previous four ($s = t-4, ..., t-1$) fiscal years. This sum is adjusted for prior year obligations, operating expenses, designated state expenses, and reserves for prior year dividends. The dividend is obtained by dividing the total distribution by the number of eligible applicants.

public backlash so that reinvesting the fund's earnings has become the implicit norm.

Since the dividend is a significant source of income for many Alaskan households, changes in the expected next dividend are frequently discussed in the local media, a fact which I document in Kueng (2015a) using an extensive narrative analysis of all major Alaskan newspapers starting in the early 1980s. The annual dividend amount, which is based on data that is largely known at the end of June of each year, is typically announced by mid-September or earlier, before the first payments are made starting in October. Since the mid-1990s, all information necessary to estimate the dividend is published on the APFC's website. Moreover, since only one fifth of the annual distribution depends on the fund's income in the current fiscal year, the monthly change in the expected dividend is an order of magnitude smaller than the annual changes in the dividend payments, a fact that is confirmed by the narrative analysis as well as new data of the fund's monthly income obtained from the APFC's archive, both documented in Kueng (2015a).[5]

Uncertainty about the next fiscal year's dividend is typically largest in November right after the previous dividend has been distributed because next year's income is unknown. Throughout the fiscal year, this uncertainty gradually declines with each new monthly report of the fund's earnings. The main source of uncertainty about the size of dividend that remains between the end of the fiscal year in June and the dividend payments in October concerns the number of eligible applicants. However, annual changes in the number of eligible applicants are small and can be reasonably well predicted based on state population forecasts. Therefore, additional precautionary savings motives because of uncertainty about the sizes of the dividend cannot account for the substantial average excess sensitivity that I document in the next section, and definitely not for the large MPCs among higher-income households who have sufficient buffers of liquid assets.

3 Average Excess Sensitivity to the Dividend

Under the life-cycle permanent income hypothesis, households should smooth consumption over time, which implies that the level of consumption should be independent of the timing of predictable cash flows. Since spending should not be sensitive to predictable income, finding systematic differences in spending patterns in response to predetermined events such as the Permanent Fund Dividend payments is evidence for so-called excess sensitivity of spending to predictable income changes.

[5] For convenience, Figure A1 in the appendix reproduces the two measures of the expected Permanent Fund Dividend from Kueng (2015a).

3.1 Nonparametric Evidence

Figure 1 shows strong non-parametric evidence of such excess sensitivity by comparing average household spending growth rates of nondurables and services of Alaskans with that of households from Washington, the other nearest state. Figure 1(a) shows averages by calendar month separately for Alaska and Washington, i.e., $E[\ln(c_{i,t}) - \ln(c_{i,t-1})|s]$, $s \in$ {Alaska, Washington}.[6] The average growth rates for the two states are fairly similar except in the month of the dividend payments and the month after, evidence that households in Washington who do not receive the dividend payments are a good control group for seasonal consumption patterns.

Figure 1(b) takes the difference between the two series and shows 95% confidence intervals. Spending on nondurables and services growth is about 10% higher in Alaska during October when dividends are paid out, and this difference is highly statistically significant (t-statistic of 6.2). Consumption growth is 5% lower in the next month (t-statistic of 3.8), suggesting that there is some delayed effect of the dividend.

Figure 1(c) breaks this difference down even more, showing the differential growth rate period-by-period instead of by calendar month. Growth rate differences are highest in months with large dividend amounts per person. For instance, excess growth rates in October of 2010 and 2011 are 13% and 14% (t-statistics of 3.7 and 5.5) when the Permanent Fund paid out nominal dividends of $1,280 and $1,174, respectively, and they are substantially lower in October of 2012 and 2013 (8% and 4% with t-statistics of 3.5 and 1.2), which had smaller dividends ($878 and $900). Excess growth rates are much smaller in all other months and typically not statistically different from zero. The excess growth rate is a function of both the size of the dividend payments and the individual excess sensitivity of each household. However, as shown below, the average response masks substantial heterogeneity in individual responses, and the household-level sensitivity is a function of the size of the dividend, expressed in percent of annual spending.

Figure 1(c) also shows that there are local shocks other than the dividend payments that affect the difference in average consumption growth. For instance, the 3% lower growth in July (t-statistic of 2.6) in Figure 1(b) is mostly driven by a one-time 6% shock in July of 2011. Therefore, in the regressions below I include period fixed effects instead of month effects to take these aggregate shocks into account.

3.2 Parametric Evidence

The analysis above uses differences in the time series of average consumption growth between households in Alaska and Washington to provide non-parametric evidence of excess sensitivity.

[6] I use log changes rather than changes in levels since they are less affected by extreme values and because they are independent of the size of a household. Hence, these non-parametric results are not directly comparable with the parametric MPCs below.

Next, I instead use household-level variation in the dividend payment size while controlling for aggregate period-specific effects. Following the previous literature, the main estimating equation is

$$c_{i,t} - c_{i,t-1} = \sum_s \beta_s \cdot PFD_{i,t-s} + \tau_t + \text{Alaska}_i + \gamma' x_{i,t} + \varepsilon_{i,t} \ , \tag{5}$$

where c_{it} measures expenditures during period t by household i, $PFD_{i,t}$ denotes the dollar amount of the Permanent Fund Dividend received at the beginning of period t, s denotes periods since receiving the dividend (i.e., allowing for leads and lags, such as anticipation effects and delayed responses), τ_t are period fixed effects controlling flexibly for any aggregate effects and seasonality in spending patterns, Alaska_i is an indicator of whether a household is a resident of Alaska, $x_{i,t}$ captures family size and other household characteristics for robustness checks, and $\varepsilon_{i,t}$ are changes in spending not explained by either the dividend or the controls. The coefficient β measures the excess sensitivity of spending to receiving predetermined Permanent Fund Dividend income.

Figure 2 shows the dynamic response of monthly household spending on nondurables and services to the PFD payments by estimating the baseline specification of equation (5), which controls for the main effects of the treatment, i.e., Alaska, family size, and period fixed effects and sets $x_{i,t} = 0$. Figure 2(a) plots the regression coefficients including 6 monthly leads and 8 monthly lags of the dividend payments, $s = -6, -5, \ldots, 8$, or two quarters of leads and three quarters of lags. On average, spending on nondurables and services increases by 11 cents for each dollar of PFD received in October ($s = 0$), and this increase is highly statistically significant (t-statistic of 6.1). Importantly, even though the dividend is predetermined at least by September, and there is substantial speculation in the media throughout the year about the likely size of the next dividend, Figure 2(a) shows no evidence of any anticipation effects. The point estimates of all leads are close to zero and reasonably precisely estimated for the month prior to the dividend, for example ruling out any announcement effect larger than 2 cents at a 95% confidence level. When calculating the dynamics of the marginal effect of the dividend on nondurable and service consumption—the marginal propensity to consume (MPC)—, I therefore do not report the leads although all regressions at monthly frequency do control for those effects.

Consumption in November ($s = 1$) is only 7 cents lower than in the previous month, hence the dividend has a delayed effect on consumption by another 5 cents ($= 0.114 + (-0.067)$) relative to September, the month before the dividend payments. This is the net or marginal effect of the dividend in later months. Figure 2(b) shows the evolution of those net effects together with two standard error bands (dashed lines). The net effect is largest and most precisely estimated in the month of the dividend payment. The dividend additionally increases consumption significantly in November and December by 5 and 8 cents, respectively. The point estimates of all subsequent net effects are small and not statistically significant.

Figure 2(c) cumulates the net effects to provide the dynamic MPC together with two standard error bands. It highlights that the MPC stabilizes within one quarter of the dividend receipt, and the effect remains statistically significant up to a lag of six months. The total steady-state MPC of nondurable and service spending is about 25 cents. Half of it occurs immediately on impact, the other half occurs with a two month lag.

Overall, Figure 2 documents an economically large and statistically significant response of household spending to the Permanent Fund Dividend payments. The response is concentrated in the month in which dividend payments are transferred via direct deposits, and most of the effect occurs within the first quarter after the arrival of the cash flows. In the following, I therefore restrict the analysis to the spending response within the first quarter of receiving the dividend. Column 1 in Table 2 shows that these results are robust to controlling for liquid assets (bank balances), income, relative size of the cash flows (discussed below), and fixed effects for age, years of schooling, residential ZIP code, homeownership status, marital status, and occupation. Alaskan households spend on average about 30% of the dividend on nondurables during the quarter in which they receive the payments relative to the previous quarter.

4 Near-Rationality and Higher-Income Household MPCs

The high average response documented in the previous section is striking for several reasons. First, the nature of the dividend payments should in principle by favorable to the standard model. After all, those cash flows are highly predictable, occur regularly every year, are salient to households living in Alaska, and are fairly large. Second, the typical household in the sample has a substantial amount of liquid assets and relatively high income as shown in Table 1. Since the standard explanations of excess sensitivity, such as liquidity constraints or precautionary savings, do not seem to apply to these households, I use the potential wealth-equivalent loss of deviating from perfect consumption smoothing to test whether near-rationally can explain the high MPCs among those higher-income households, and whether the MPCs are falling in the relative size of the cash flows.

4.1 Relative Size of the Cash Flow and MPC Heterogeneity

Equation (4) shows that the potential loss is a function of three factors, the relative size of the dividend as a fraction of household consumption, the time distance between news about future dividends and the actual dividend payments, and the curvature of the sub-utility function. Since the local media report frequently on the expected dividend amount the throughout the year, differences in expected dividends across households are likely only minor. I will also abstract from heterogeneity in risk- or time-preferences, although they might be related to binding liquidity constraints for a small subset of households, which I study in the next section.

The relative size of the dividend payment is therefore a sufficient statistic of the costs of near-rational behavior. To test the predictive power of near-rationality, I modify the baseline specification of equation (5) and interact the dividend payments with the quintiles $I(q)$ of the relative dividend size,

$$c_{i,t} - c_{i,t-1} = \sum_q \beta_q \cdot I(q)_{i,t} \times PFD_{i,t-s} + \sum_q \gamma_q I(q)_{i,t} + \tau_t + \text{Alaska}_i + \gamma' x_{i,t} + \varepsilon_{i,t} \ , \qquad (6)$$

or with the relative dividend size itself, which is calculated as the amount of the dividend received divided by the household's annualized expenditures. The first specifications using quintiles is more robust to extreme values. On average, the relative dividend size for Alaskan households in each quintile is 1.6%, 2.7%, 3.7%, 5.4%, and 10.3%, respectively.

Figure 4 shows that the MPC indeed falls significantly with the increase in the predicted loss, consistent with near-rational behavior. Households in the lowest quintile for whom the dividend is a small fraction of annual spending and hence a less important source of (permanent) income have an estimated MPC of 81%, substantially higher than the MPC of 15% for households in the highest relative dividend size quintile, for whom the dividend is a much more substantial income source. Moreover, the MPC declines steadily as we move from lower to higher quintiles. The precision of the estimated excess sensitivity is increasing in the relative size of the dividend as it explains a larger fraction of spending for households in higher quintiles.

Table 3 documents that these results are robust to controlling for other potential mechanisms. In particular, the sharp decrease in MPCs as a function of the relative size of the dividend is robust to controlling for income, age, and the amount of liquid assets, measured as the total balance of savings and checking accounts. In addition to Alaska and period fixed effects, I also include fixed effects for education, residential ZIP code, homeownership status, marital status, and occupation. Colum 1 provides the baseline average MPC across all households. The estimate is similar to the ones shown in Figure 2, and the slightly larger magnitude is a result of the additional controls.

Column 2 and 3 show the effect of a change in the size of the relative dividend on the MPC. A one standard deviation increase in the relative size decreases the MPC by 12 percentage points on average. Alternatively, moving up one relative size quintile reduces the MPC by 15 percentage points. Hence, on average, moving from the lowest to the highest quintiles reduces the MPC from about 75% to 15%, which is very similar to the unconditional estimates shown in Figure 4.

Column 4 shows that measuring the potential costs using the relative size of the dividend instead of the squared dividend in levels, as previously done in the literature (e.g., Souleles (1999) and Scholnick (2013)), is important. The coefficient for the squared dividend payment is both economically and statistically insignificant, and including the squared term does not

affect the coefficient on the relative size of the dividend, i.e., the linear term.

Therefore, Figure 4 and Table 3 provide strong evidence that the relative size of the cash flow is an important factor for explaining excess sensitivity, and in particular for explaining the dispersion of excess sensitivity in the cross-section of higher-income households. However, note that the cost from failing to smooth consumption alone cannot fully explain excess sensitivity since the response is statistically significant in each relative-size quintile in Figure 4.

4.2 Household Income and MPC Heterogeneity

Figure 3(b) makes the same point using a different source of variation. If near-rationality explains excess sensitivity, then we would expect the MPC to be *larger* for high-income households for whom the dividend is only a small fraction of annual income. Indeed, panel 3(b) shows that the MPC is *increasing* in household after-tax income per household equivalent.[7] Households in the top income quintile have an average MPC of 64% compared with an MPC of only 8% for households in the lowest quintile, for whom the dividend is a substantial source of total family income. The average Alaskan income per household equivalent in each quintile is \$16,000, \$30,000, \$42,000, \$58,000, and \$104,000, respectively. Note that the slope of the MPC is steeper in the relative dividend than in the household income. This is presumably due to the fact that average annual spending is a better measure of permanent income and hence of the intertemporal budget constraint than annual income.

While this result might appear obvious after realizing that income is inversely related to the costs from failing to smooth the dividend, it is noteworthy that it stands in sharp contrast to results in other studies who use income as an alternative measure of liquidity constraints, e.g. Johnson, Parker and Souleles (2006). This is typically necessary since liquid assets tend to be poorly measured in expenditure surveys. The next section shows that low levels of liquid wealth indeed predict a high MPC for households in the bottom two income quintiles who have no or few liquid assets relative to households with similar income but sufficient levels of liquidity, consistent with the importance of borrowing constraints. However, the results in Figure 3(b) show that while income might be a reasonable predictor of (the lack of) liquid assets for lower-income households, the relationship with excess sensitivity breaks down for higher-income households. At higher levels, income has the opposite implication for excess sensitivity since the welfare-cost effect dominates the liquidity effect.

Columns 5 and 6 of Table 3 show that these results are robust to including various controls. Importantly, the steep increase of the MPCs as a function of income is preserved when controlling for the relative size of the dividend. Hence, the two source of variation—household income and the relative size of the cash flow—provide independent evidence for the same mech-

[7] I use the OECD household equivalence scale to adjust for differences in family size, which assigns a value of 1 to the first household member, and adds 0.7 for each additional adults and 0.5 for each child below the age of 16.

anism, the predictive power of near-rationality. For instance, a one standard deviation increase in income raises the MPC by 17 percentage points on average. Alternatively, moving up one income quintile increases the MPC by 14 percentage points. Hence, on average, moving from the lowest to the highest quintiles increases the MPC from about 3% to 60%, which is again very similar to the unconditional estimates shown in Figure 3(b).

5 Credit Constraints and Lower-Income Household MPCs

Credit constraints are the main explanation proposed in the literature for excess sensitivity. Households might want to borrow against future income, but in the case of the Permanent Fund Dividend, a law implemented in 1989 prevents individuals from assigning the dividend to any third party other than the government. The dividend can therefore not be used as legal collateral for any debt contract.

Borrowing constraints are of course only binding for households with low levels of liquid assets. On average, the total balance of savings and checking accounts combined is $40,000 for Alaskan households in the sample, with another $67,000 invested in liquid, taxable brokerage accounts. Hence, liquidity constraints cannot explain the large MPCs, especially for households in the top two income quintiles in Figure 3(b), who have on average $55,000 respectively $98,000 in their bank accounts, with another $84,000 respectively $236,000 in their taxable brokerage accounts.

I therefore focus on the left part of the distribution in Figure 3(b), i.e., on households with an annual income below the sample median of $75,000 (or $42,000 per household equivalent). The average bank account balance is $17,000 for Alaskan households in this subsample, with a sizeable standard deviation of $7,000, which implies that there is sufficient variation in liquid assets to test whether liquidity constraints predict higher MPCs for lower-income households.

I divide the sample into three bins. The first bind contains households with no or almost no liquid assets, defined as having less than $100 in total bank account balances. Households in the second bin have balances between one and three times the amount of their dividend payments. Although such households are appear unconstrained, precautionary savings motives might cause them to respond excessively to the dividend payments. The third bin contains households with bank balances of more than three times their annual dividend payments.

Figure 4 shows that households with low or no liquid assets indeed have a very high MPC, with a statistically significant point estimate of 88%. The 95-percent confidence interval is wide because there are few households with such low levels of liquid assets in the sample, reflecting the fact that such households are underrepresented relative to the population of Alaskan households. Households with bank balances of more than three times their dividend payment (but with below-median income) have an economically and statistically insignificant MPC of just 8%. Such households are unconstrained and welfare loss calculations therefore

dominate their behavior, leading them to spend less of the dividend on impact since the dividend is a substantial fraction of their annual income. The MPC of households with bank balances between one and three times their dividend payment lies in between, suggesting that they trade off precautionary savings motives with welfare loss calculations.

6 Extensions and Robustness

In this section I address a number of concerns that typically arise in the literature, and I discuss a couple of extensions of the main analysis, in particular an external validity check using the Consumer Expenditure Survey.

6.1 Consumption vs. Spending

Many nondurables and services do have a durable component, especially when looking at frequencies higher than annual, so that spending does not necessarily reflect consumption. One is therefore concerned that households might time the purchase of such goods to the arrival of the dividend cash flows while spreading the consumption of the goods evenly over the year, as predicted by the standard model.

To address this concern I take advantage of the high quality of the expenditure transaction data, especially for disaggregated categories, compared to traditional expenditure surveys. In panel A of Table 3, I follow Lusardi (1996) and show that strictly nondurable goods, i.e., goods that arguably have an low durable component, also strongly respond to the dividend payments. Column 1 documents that spending on food and dining substantially increases in the first quarter after the dividend is paid out. The magnitude is in line with previous research, such as Broda and Parker (2014) for example, who estimate the spending responses to the smaller economic stimulus payments in 2008 using the Nielsen Consumer Panel. Column 2 shows that a large fraction of this amount is spent on groceries. Column 3 uses expenditures on a service item, personal care (spa, massages, and hair salons), to make the same point. Similarly, spending on children's activities and toys (column 4) and on gasoline (column 5) also increases significantly.

Therefore, focusing on more disaggregated data shows that spending significantly increases across many categories in response to the dividend, including strictly nondurables. Moreover, the absence of a reversal of the response of nondurables shown in Figure 2 strongly suggest that the excess sensitivity of spending on nondurables cannot be explained by intertemporal substitution of expenditures while smoothing the consumption of the service flows.

6.2 Robustness of Excess Sensitivity Results

Colums 1 to 4 in panel B of Table 3 show that the estimated response of nondurables and services is robust. The quantile regression in column 1 shows that the median response is very

similar to the average response in column 1 of Table 2. Columns 2 and 3 show that controlling for family size and other household characteristics, which include income, liquid assets, and fixed effects for 5-year age bins, residential ZIP code, homeownership status, marital status, and occupation, does not affect the estimated consumption response. The point estimates of about 28 cents is similar to the cumulative effect of 24 cents after the first three month shown in Figure 2.[8] Column 4 limits the sample to Alaskans and hence estimates the effect based on variation in the size of the dividend and relative to quarters without dividend payments. The point estimate is almost identical, although less precisely estimated due to the smaller sample size.

6.3 Response of Durables and Total Expenditures

It is worth noting that changes in durables spending do not necessarily reflect (large) changes in the consumption of the service flow from those durables. Therefore, finding a positive effect of the dividend on durable spending is not necessarily evidence against the standard model. However, one might still be interested in the total effect of the dividend payments on spending.

Figures 5(a) and 5(b) show the dynamic response of household spending on durables that are purchased with a credit card and hence can be classified accordingly. While the overall pattern is similar to that of nondurables and service consumption in Figure 2, there are some notable differences. First, the effect is slightly smaller both on impact (9 cents) and after one quarter (16 cents), which is due to the fact that those transactions occur less frequent and only capture smaller durables purchases. Second, there is some evidence of intertemporal substitution. Purchases of smaller durables falls slightly by 3 cents the month before the dividend payments and this drop is marginally significantly with a t-statistic of 2.2. Both results suggest that Alaskan households might time the purchase of durables to the predictable arrival of the dividend cash flows. The MPC of smaller durables is therefore slightly hump-sharped. Finally, the effect of durables is measured less precisely, especially the lagged response. The cumulative effect becomes insignificant already at the 5th monthly lag, although the point estimates remain fairly stable at around 10 cents.

Column 6 of Table 3, panel A, estimates the response of such smaller durables using changes between the fourth and third quarter, resulting in an MPC of about 12%. Some durables purchases are not financed by a credit card but with cash instead. Column 7 therefore adds cash withdrawals from ATMs to the measure of durable spending, presumably with the intent to spend in, although this can only be conjectured. The MPC within the first quarter of this more comprehensive measure of durables is 19%.

Column 8 uses all forms of spending, including other spending such as gifts and charity, stu-

[8]Differencing at quarterly frequency drops more observations than with monthly spending, which explains the small differences in the two point estimates.

dent and car loans, mortgage and rent payments, etc.[9] The average MPC of total expenditures is 71%, which is large.

6.4 External Validity using the Consumer Expenditure Survey

To provide external validity of the excess sensitivity results presented above, I use the Consumer Expenditure Survey (CE), which is the standard data set used in previous research. The CE sample spans the entire period since the first dividend was paid out in 1982 up to 2013, but covers fewer Alaskan households per period and follows them only for at most four quarters.

Since the CE does not ask Alaskan households directly whether they received the Permanent Fund Dividend and how large the payment was, the payment has to be imputed based on family size, state of residence, calendar year, and the fraction of households in any year that do not receive the dividend at all or in full either as a check or direct deposit. This fraction can be calculated for each year based on aggregate statistics provided by the Alaska Permanent Fund Dividend Division. Because the state identifier for Alaska is suppressed in the public-use CE sample before 1996, I use the confidential data at the BLS. As is standard in the literature, I aggregate the expenditures for each household-interview to "three-monthly" aggregates. Spending on nondurables and services is defined to be comparable to the concept used for the PFW sample, which in turn approximates the NIPA definition (see Table A.1).

Column 5 of Table 3, panel B, shows a statistically significant MPC of 8% using the CE sample. In order to compare this estimate to the average MPC based on the PFW sample, I make two adjustments. First, I apply the same imputation procedures that are necessary for the CE sample also to the PFW sample. Specifically, I impute the dividend payments based on family size, state of residence, and calendar year, thereby ignoring the information about the exact size of the payments based on the direct deposits.[10] Column 6 shows that the added measurement error reduces the MPC from 29% to 19%.

Second, I take into account the difference in sample compositions. While the CE is designed to be representative for the entire U.S., it is not representative for subpopulations such as single states. Hence, neither the CE sample nor the PFW sample is a representative sample of the population of Alaskan households. In particular, Alaskan households in the CE have a lower average family income of $63,000 in local dollars of 2014, compared to a mean income of $94,000 in the PFW sample. Since the MPC is increasing in income as shown above, these differences in sample composition matter. Column 7 therefore adds an interaction term of the dividend with

[9] See Table A.1 for a description of the different spending categories.

[10] This approach follows the idea used in a series of papers by Romer (1986b,a, 1991) who compares pre- and post-WWII macroeconomic time series by making the cleaner post-war data as noisy as the pre-war data. Here, I make the cleaner dividend income measure in the PFW sample as noisy as the imputed income in the CE.

after-tax family income to the specification used for column 6. For each $100,000 of income, the MPC increases by about 19 percentage points, which implies an average MPC of 10% using average Alaskan family income in the CE. This number is very close to the point estimate of 8% in column 5.

Therefore, measurement error introduced by the necessity to impute dividend payments in the CE can explain about half of the difference between the response estimated in the PFW and CE samples, and the other half is explained by composition effects, the fact that both samples represent different segments of the population.

The CE sample also allows me to reconcile these new results with the estimates provided by Hsieh (2003), who was the first to use this quasi-natural experiment to test the standard consumption model. His analysis found no response of spending to the dividend payments using the CE. One difference is that throughout this paper I estimate the effect of the PFD on changes in spending (i.e., MPCs), while he estimates the effect on log-changes in spending (i.e., an elasticity). In order to estimate an elasticity, he divides the PFD payments by self-reported family income. Unfortunately, family income in the CE suffers from substantial measurement error and under-reporting. In Kueng (2015b) I replicate his estimate of the spending elasticity to the PFD payments using his shorter CE sample from 1980 to 2001. The measurement error in family income explains 70% of the difference between his estimated elasticity of 0 and the elasticity of 14% estimated using total household expenditures to normalize PFD payments, an alternative measure of (permanent) income that is more precisely quantified in the CE. The other 30% of the difference between is due to the fact that the latter is based on a longer sample from 1980 to 2013, which contains the much larger variation in annual dividend payments during the 2000s, while the dividend grew almost linearly in his sample period from 1983 to 2001.

6.5 Anticipation Effects

In order to obtain a complete picture of the impact of the Permanent Fund Dividend on household behavior, we need to assess what households do with the part of the dividend that they do not spend. One possibility is that they respond to news about future dividends before the dividends are paid out. Note that although there is no systematic response to leads of the dividend payments at any horizon, this does not fully rule out forward-looking behavior. Instead, under the rational-expectation version of the standard model, households would rationally only respond to new information about the dividend, and such news shocks should in fact not be systematically related to any predictable variable, in particular not to the calendar month or any lead of the actual dividend. Moreover, rational forecast errors have an average of zero and are as likely positive as negative, contrary to the actual dividend payment and any of its leads or lags.

One way to assess how much households could on average respond in advance to the pay-

ments is to estimate how much of the dividend is left after taking into account both the amount spent when receiving the dividend (i.e., the excess sensitivity of total expenditures) and the additional amount of federal income taxes due in the next year.

Although the PFD is a program run by the state of Alaska, the dividends for adults are fully taxable for federal income tax purposes, and depending on the amount of the dividend, children's dividends may be taxable too. I estimate the average marginal tax rate (AMTR) paid by Alaskan households on the marginal dollar of dividends received by regressing tax expenses in the current year on the previous year's dividend, PFD_i^{lag}, fully interacted with calendar month fixed effects,

$$T_{it} = \sum_{q=1}^{12} \tau_q \cdot \left(m_q \times PFD_i^{lag} \right) + \sum_{q=1}^{12} m_q + u_{it}, \tag{7}$$

where T are the household's tax payments (or refunds) in period t, and m_q are 12 monthly fixed effects. τ_q is the average marginal tax rate paid in month q. Adding up all 12 tax rates over the year yields the AMTR paid on one dollar of additional Permanent Fund Dividend income.

Figure 5(c) shows the monthly marginal tax rates τ_m from estimating (7). I restrict the sample of Alaskans to the 86% households who receive the entire dividend amount in form of a direct deposit. This makes sure that they did not elect to have federal taxes withheld from their dividend checks, which would induce a downward bias. Since Alaska does not have a state or local income tax, the federal AMTR is all we need. Although the individual coefficients are not very precisely estimated, we see that households pay between 8, 6, and 5 cents of additional taxes in February, March, and April for each additional dollar of PFD received. The response in all other months is small and statistically insignificant, consistent with the pattern of federal income tax revenues from aggregate statistics.

Adding up the coefficients of the federal marginal taxes across all calendar months yields a point estimate for the federal AMTR of 22.5%, with a standard error of 10.3%. This estimate is similar to independent estimates provided by Mertens (2013), who extends the estimates in Barro and Redlick (2011), finding an AMTR across all tax units in the U.S. between 22.1% and 23.5% in 2010.

Combining the estimated AMTR with the estimated MPC of total expenditures of 71% (Table 3, column 8) leaves only 6 cents unaccounted for per dollar of dividend received. Since the budget constraint has to be satisfied on average, and since households presumably save some amount of this residual income, there is little room for large anticipation effects, consistent with the direct evidence provided in Figures 2(a) and 5(a).

6.6 Consumption Commitments

Models with consumption commitments offer an additional interpretation of excess sensitivity. In these models, in order to consume the service flow from durable goods such as housing,

households have to commit ex-ante to purchasing the stock (e.g., Chetty and Szeidl (2007)), which is illiquid and involves a transaction cost to turn it into a liquid asset (e.g., Kaplan and Violante (2014)). Therefore, they optimally tolerate larger deviations in more flexible spending, such as consumption of nondurables.

Consistent with these predictions, I find that homeowners have on average a 25 percent point larger MPC than non-homeowners, with a standard error of 7%. More importantly, consumption commitments interact with the expected size of the income shock such that homeowners' tolerate predictable changes in flexible consumption, i.e., nondurables and services, but respond less to larger shocks. Therefore, a central prediction of these models is that the relative difference in excess sensitivity between homeowners and non-homeowners should decrease as the size of the shock increases.[11] To test this prediction, I interact homeownership status with quartiles of the relative size of the dividend in order to analyze how the differential responses change as the relative size of the dividend payments increases. Figure 5(d) shows that the estimated difference in the response is indeed declining in the relative size of the dividend, consistent with the theory, although the power is low.

While these results are consistent with models of consumption commitments and wealthy hand-to-mouth consumers, there are two important caveats. First, homeowners in this sample have sufficient money in their bank accounts, on average $65,000, so that they do not appear to be credit constrained. This is also the case for each of the relative dividend size quartiles, in which average bank balances range from $81,000 in the lowest quartile to $35,000 in the highest quartile. Consequently, conditioning on liquid assets and income does not substantially affect the results. Therefore, other frictions that differentially affect homeowners, either economic frictions or behavioral biases, probably also play an important role.

Second, an important limitation of optimization-based models of excess sensitivity—which in addition to consumption commitment models also include models of rational inattention (e.g., Reis (2006) and Luo (2008)) and optimization frictions (e.g., Chetty (2012))—is that they cannot explain why household spending responds to fully predetermined cash flows rather than to new information about those future cash flows. In those models, households' optimal responses typically smooth consumption between dates at which they rebalance their portfolios (consumption commitments), update their information sets (rational inattention), or pay their fixed cost of re-optimization (optimization frictions), which is inconsistent with the large MPCs estimated in this paper. Moreover, the size of new information about future Permanent Fund Dividends is much smaller than the dividend itself, because the size of the dividend is highly predictable and hence rational forecast errors are small. Therefore, even if households did rebalance or update only infrequently and only at dates at which dividends are paid out,

[11] This is a distinct prediction for excess sensitivity relative to models with habit formation; see Chetty and Szeidl (2014).

they should still only react to the new information they received since the last time they changed their consumption plans. However, the observed MPCs are an order of magnitude larger than what one would expect if households only responded to new information about the dividend. Moreover, consumption should as likely decrease as increase on those dates and need not necessarily be related to the nominal amount of the dividend, since forecast errors due to new information should on average be zero. Instead, households respond to the entire dividend amount, not just the news component.

7 Conclusion

This paper finds significant evidence of consumption excess sensitivity in response to nominally large, predetermined, salient cash flows. The *potential* loss from failing to smooth consumption and instead consuming the entire predictable cash inflow varies systematically across households. This potential loss is shown to be monotonically increasing in the relative size of the cash flow as a fraction of permanent income, measured as annual household consumption. The *realized* loss on the other hand is endogenous and depends on the response of each household to the cash flow. Sorting households according to their potential loss shows that households for which the loss would be the largest violate standard the permanent income life-cycle hypothesis the least. Households for which the loss is trivial violate the basic model's prediction the most, consistent with them following near-rational alternatives.

This paper is among the first to use a single income source to show that the relative size of the payment and hence the potential loss from not smoothing out this cash flow is a crucial factor in explaining excess sensitivity in household consumption. The relative size of the cash flows together with liquidity constraints, which help to explain excess sensitivity among lower-income households with few or no liquid assets, could therefore help to reconcile the large literature that tests for excessively sensitive in household consumption.

The highly statistically significant deviation of household consumption from the theory's main prediction shown in this paper does not imply a significant deviation in terms of equivalent variation for most households. The potential loss of deviating from the nominal model one likes to test, which in the case of the permanent income life-cycle model is captured by the relative size of the predetermined payments, can therefore be used as a measure of economic power of a research design in testing the underlying theory.

Finally, while the failure of the standard theory documented in this paper is not economically significant for most households, it has important implications for policy. Many policies have a large predictable component, such as an economic stimulus program, automatic stabilizers, and many more. According to the standard model without credit constraints, one would expect households to adjust to the news about such policies, and only do so if it affects the households' permanent income or life-time budget constraints. However, not smoothing out

those cash flows has trivial costs for most households and for most policies. Therefore, gaining a better understanding of the near-rational alternatives that households follow in response to such relatively small payments is important in order to design effective and robust policies. At the same time, the fact that the deviations from the standard model documented in this paper are consistent with households following near-rational alternatives implies that optimization-based extensions of the standard model might have limited economic power and thus might not be very robust. Modelling near-rational behavior in a parsimonious and robust way thus remains an important challenge for future research.

References

Akerlof, George A. and Janet L. Yellen, "A Near-Rational Model of the Business Cycle, with Wage and Price Inertia," *Quarterly Journal of Economics*, 1985, *100* (Supplement), 823–838.

Baker, Scott R., "Debt and the Consumption Response to Household Income Shocks," *working paper*, 2014.

Barro, Robert J. and Charles J. Redlick, "Macroeconomic Effects from Government Purchases and Taxes," *The Quarterly Journal of Economics*, 2011, *126*, 51–102.

Broda, Christian and Jonathan A Parker, "The economic stimulus payments of 2008 and the aggregate demand for consumption," *Journal of Monetary Economics*, 2014, *68*, S20–S36.

Browning, Martin and Thomas F. Crossley, "The Life-Cycle Model of Consumption and Saving," *Journal of Economic Perspectives*, 2001, *15* (3), pp. 3–22.

Chetty, Raj, "Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply," *Econometrica*, 2012, *80* (3), 969–1018.

_ and Adam Szeidl, "Consumption Commitments and Risk Preferences," *The Quarterly Journal of Economics*, 2007, *122* (2), 831–877.

_ and _ , "Consumption Commitments and Habit Formation," *working paper*, 2014.

Cochrane, John H., "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *American Economic Review*, 1989, *79* (3), 319–337.

Fuchs-Schuendeln, Nicola and Tarek A. Hassan, "Natural Experiments in Macroeconomics," *working paper*, 2015.

Gabaix, Xavier and David Laibson, "The 6D bias and the equity-premium puzzle," in "NBER Macroeconomics Annual 2001, Volume 16," MIT Press, 2002, pp. 257–330.

Goldsmith, Scott, "The Alaska Permanent Fund Dividend Program," *working paper*, 2001.

Hsieh, Chang-Tai, "Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund," *American Economic Review*, March 2003, *93* (1), 397–405.

Jappelli, Tullio and Luigi Pistaferri, "The Consumption Response to Income Changes," *Annual Review of Economics*, 2010, *2*, 479–506.

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

Kaplan, Greg and Giovanni L Violante, "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 2014, *82* (4), 1199–1239.

Kreinin, Mordechai E, "Windfall income and consumption: Additional evidence," *The American Economic Review*, 1961, pp. 388–390.

Kueng, Lorenz, "Expected Permanent Fund Dividends: New Measures Based on a Narrative Analysis and Archival Data," *working paper*, 2015.

_ , "Revisiting the Consumption Response to the Alaska Permanent Fund Dividend using CE Data," *working paper*, 2015.

Luo, Yulei, "Consumption dynamics under information processing constraints," *Review of Economic Dynamics*, April 2008, *11* (2), 366–385.

Lusardi, Annamaria, "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets," *Journal of Business & Economic Statistics*, 1996, *14* (1), 81–90.

Mertens, Karel, "Marginal Tax Rates and Income: New Time Series Evidence," *National Bureau of Economic Research*, 2013.

Parker, Jonathan A., "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, 1999, *89* (4).

Reis, Ricardo, "Inattentive Producers," *Review of Economic Studies*, July 2006, *73* (3), 793–821.

Romer, Christina D., "Is the Stabilization of the Postwar Economy a Figment of the Data?," *The American Economic Review*, 1986, pp. 314–334.

_ , "Spurious volatility in historical unemployment data," *The Journal of Political Economy*, 1986, pp. 1–37.

_ , "The cyclical behavior of individual production series, 1889-1984," *The Quarterly Journal of Economics*, 1991, pp. 1–31.

Scholnick, Barry, "Consumption Smoothing after the Final Mortgage Payment: Testing the Magnitude Hypothesis," *Review of Economics and Statistics*, 2013, *95* (4), 1444–1449.

Souleles, Nicholas S., "The Response of Household Consumption to Income Tax Refunds," *American Economic Review*, 1999, *89* (4), 947–958.

Table 1: Summary statistics

	State of Alaska		*State of Washington*	
	Mean	St.Dev.	Mean	St.Dev.
Permanent Fund Dividend:				
- annual payments (real, 2014)	2,318	1,732	--	--
- per annual household income (in %)	4.20	8.22	--	--
- per annual household consumption (in %)	4.86	4.00	--	--
Quarterly expenditures (real, 2014):				
- nondurables and services	8,264	5,833	8,042	6,085
- "small" durables	3,086	3,046	2,971	3,015
- other items in total expenditures	12,737	15,412	12,815	16,009
Demographics:				
- Family size	2.70	1.42	2.71	1.30
- Age	32.20	10.71	30.94	10.32
- Education (years of schooling)	15.34	2.22	16.03	2.12
- Annual after-tax household income (real, 2014)	93,801	72,811	94,829	75,825
- Savings and checking account (real, 2014)	39,959	81,388	61,130	104,846
- Taxable brokerage accounts (real, 2014)	67,447	235,977	109,014	308,570
Number of households	1,380		2,191	

Notes: Nominal variables are in local dollars of 2014 and, expect for annual dividend payments, are winsorized at 1%.

Table 2: Heterogeneity of MPCs

Dep. var.: Δc_{it}, quarterly nondurables and services	average MPC	by shock size			by income	
		linear	quintile	squared PFD	linear	quintile
	(1)	(2)	(3)	(4)	(5)	(6)
PFD payments	0.297***	0.490***	0.744***	0.288***	0.067	0.032
	(0.044)	(0.078)	(0.113)	(0.095)	(0.069)	(0.052)
PFD x shock size		-2.875***				
		(0.775)				
PFD x shock size quintile			-0.152***			
			(0.032)			
squared PFD/100				-0.014		
				(0.196)		
PFD x income / $100,000					0.485***	
					(0.144)	
PFD x income quintile						0.143***
						(0.027)
Observations	46,807	46,807	46,807	46,807	46,807	46,807
R-squared	0.108	0.109	0.110	0.109	0.109	0.109
- Alaska FE	YES	YES	YES	YES	YES	YES
- Period FEs	YES	YES	YES	YES	YES	YES
- Shock size	YES	YES	YES	--	YES	YES
- Income	YES	YES	YES	YES	YES	YES
- Liquid assets	YES	YES	YES	YES	YES	YES
- Household characteristics	YES	YES	YES	YES	YES	YES

Notes: To simplify interpretation, all quintiles have values from 0 to 4. For robustness, the linear interactions as well as the dependent variable are winsorized at the 1% level. Household characteristics include fixed effects for age, education, residential ZIP code, homeownership status, marital status, and occupation. Robust standard errors in parentheses, clustered at the household level, are adjusted for arbitrary within-household correlations and heteroskedasticity.

Table 3: Spending Across Goods, Robustness, and External Validity

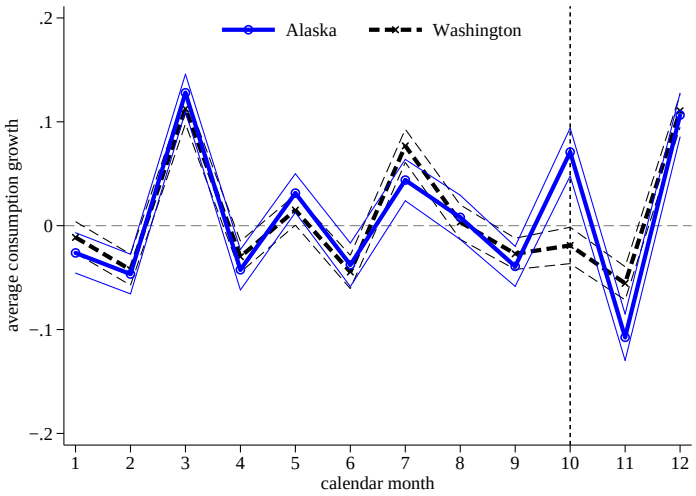
	food and dining					smaller durables		
	all	groceries	personal care	kids activities	gasoline	cc txns	incl. withdrawals	total exp
Panel A : Spending across goods	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PFD payments	0.075***	0.058***	0.007***	0.005***	0.020***	0.123***	0.185***	0.714***
	(0.014)	(0.011)	(0.002)	(0.001)	(0.005)	(0.028)	(0.040)	(0.151)
- Alaska FE	YES	YES	YES	YES	YES	YES	YES	YES
- Period FEs	YES	YES	YES	YES	YES	YES	YES	YES
Observations	46,807	46,807	46,807	46,807	46,807	46,807	46,807	46,807
R-squared	0.140	0.109	0.013	0.011	0.060	0.060	0.042	0.062

	Robustness				External validity			
	median	family size	hh charact.	Alaskans only	CE	PFD imputation	sample composition	IV
Panel B : Robustness and CE	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PFD payments	0.265***	0.282***	0.286***	0.284***	0.079**			
	(0.032)	(0.043)	(0.044)	(0.051)	(0.036)			
PFD x family size						0.190***	-0.021	0.264***
						(0.030)	(0.048)	(0.040)
PFD x family size x income/$100,000							0.187***	
							(0.044)	
predicted MPC using average CE income							*0.097*	
- Alaska FE	YES	YES	YES	--	YES	YES	YES	YES
- Period FEs	YES	YES	YES	YES	YES	YES	YES	YES
- Family size	--	YES	YES	--	--	--	--	--
- Other household characteristics	--	--	YES	--	--	--	--	--
Observations	46,807	46,807	46,807	17,899	385,800	46,807	46,807	46,807
R-squared	0.068	0.107	0.109	0.117	0.006	0.107	0.108	0.106

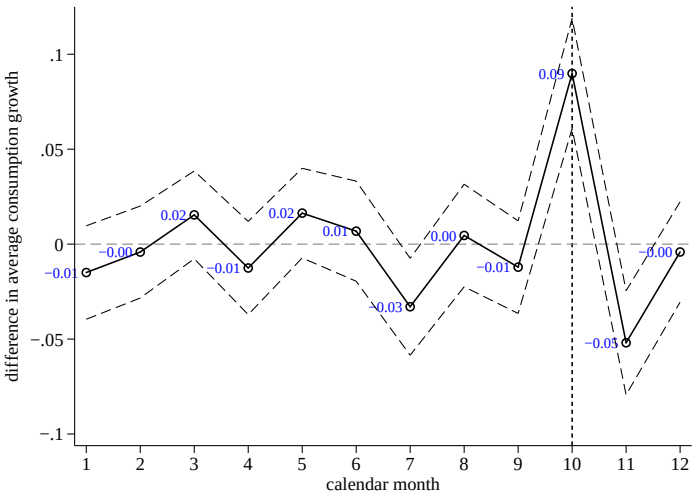
Notes: Other household characteristics include income, bank account balances, and fixed effects for age, education, residential ZIP code, homeownership status, marital status, and occupation. Robust standard errors in parentheses, clustered at the household level, are adjusted for arbitrary within-household correlations and heteroskedasticity. Durable expenditures exclude check transactions, which cannot clearly be assigned to a spending category. Expenses totaling the amount of the annual dividend are excluded in order to avoid any mechanical effects due to misclassifications. For robustness, all dependent variables are winsorized at the 1% level, except for (1) in panel B.

Figure 1 – Non-parametric evidence: Average household nondurables spending growth

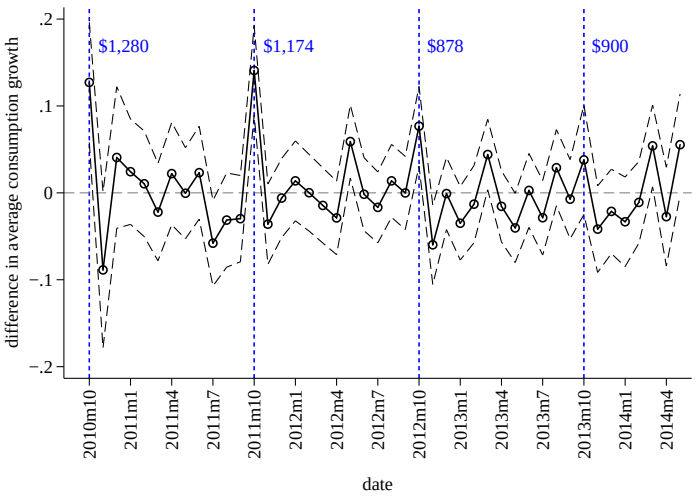
(a) average consumption growth by month



(b) difference in average consumption growth by month



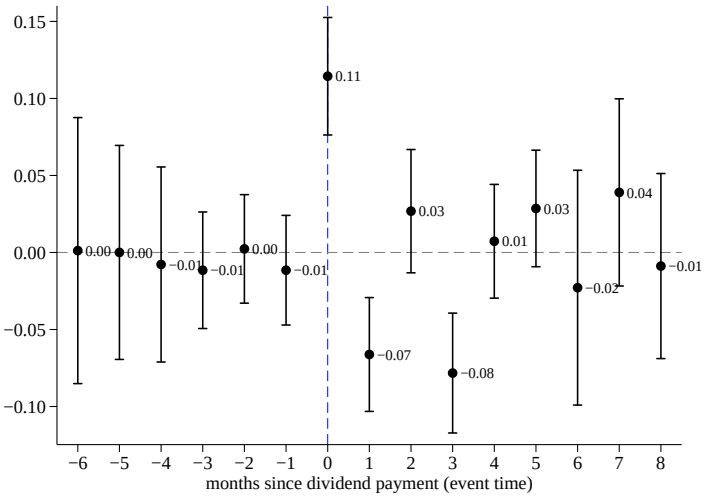
(c) difference in average consumption growth by period



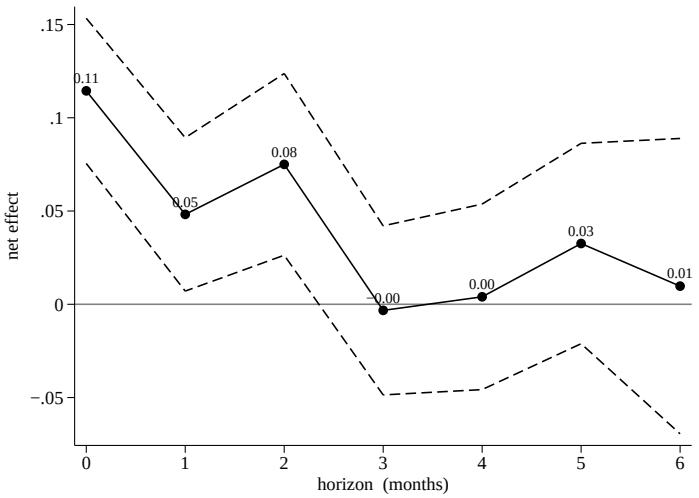
Notes: These figures show the average growth rate (i.e., log-changes) of nondurable and service consumption for households in Alaska vs. Washington: $E[\Delta ln(c_{it})|s]$, $s \in$ {Alaska, Washington}. Vertical dashed blue lines show the arrival of the Permanent Fund Dividend payments, with the numbers in (c) referring to the dividend amount per person. Black and red dashed lines are 95% confidence intervals.

Figure 2 – Parametric evidence: Average nondurables MPC

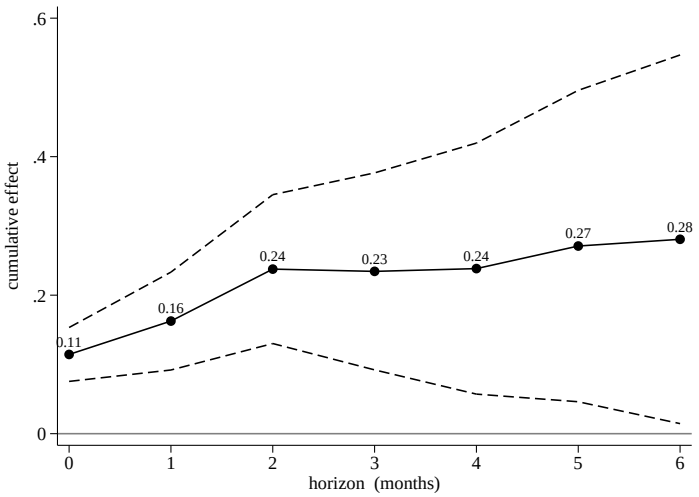
(a) βs from $\Delta c_{it} = \sum_s \beta_s PFD_{i,t-s} + \tau_t + \text{Alaska}_i + \varepsilon_{it}$



(b) impulse response function: $MPC(s) \equiv \frac{\partial E[c_{i,s} - c_{i,-1} | \tau_t, \text{Alaska}_i]}{\partial PFD_{i,0}}$

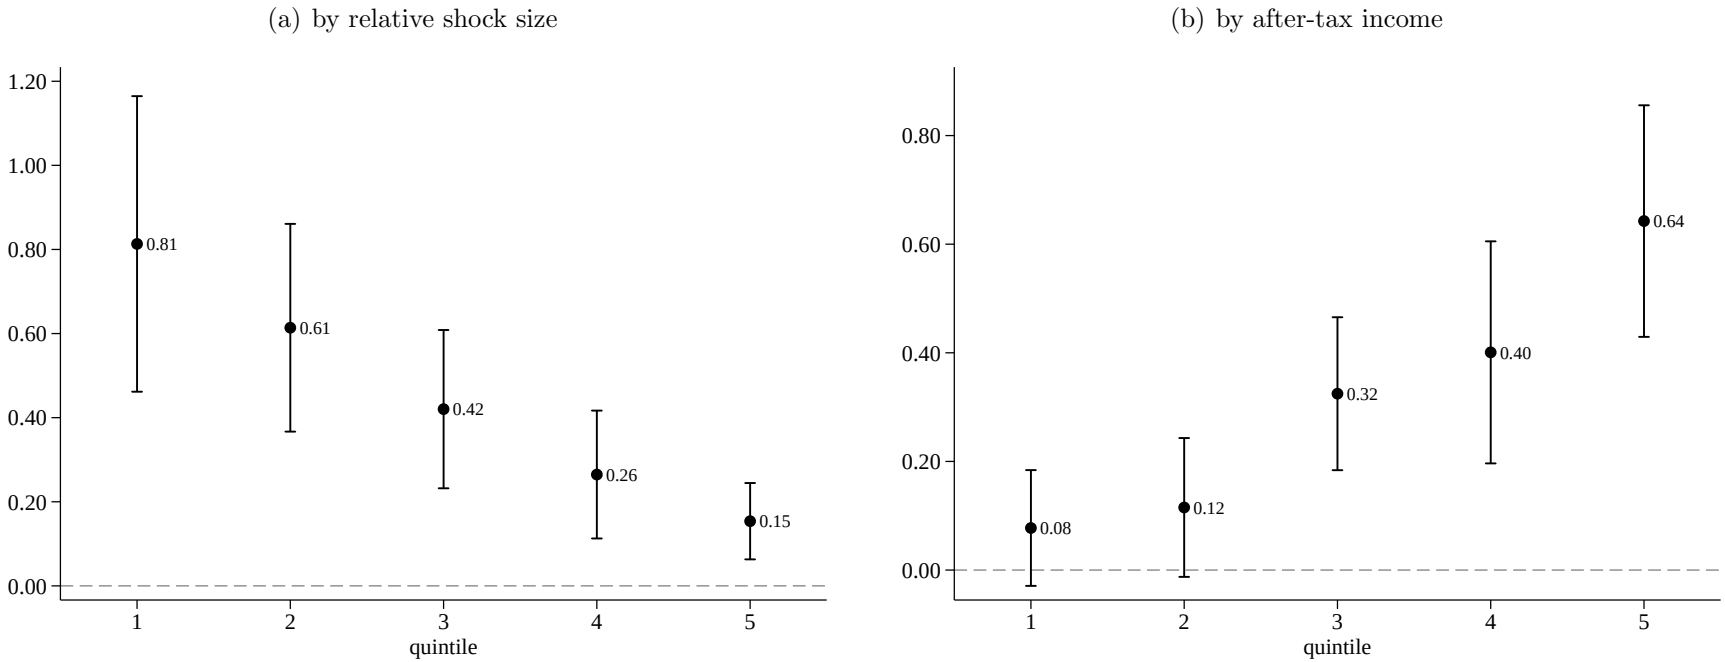


(c) cumulative MPC $= \sum_s MPC(s)$



Notes: These figures show the the response of nondurable consumption and services to the receipt of the Alaska Permanent Fund Dividend (PFD). All specifications use changes in levels as the dependent variable. Panel (a) shows all the regression coefficients on the dividend, including all leads. Panel (b) shows the additional delayed contributions to the spending response over the first seven months, while panel (c) cumulates those responses. Bars and dashed lines show two standard error bands, using robust standard errors clustered at the household level, which adjust for arbitrary within-household correlations and heteroskedasticity.

Figure 3 – Near-Rationality and Heterogeneity in MPCs



Notes: This figure decomposes the average quarterly MPC in nondurables and services across the two predicted dimensions of near-rational losses: (a) uses the relative shock size, which is the PFD divided by annualized household total expenditures; and (b) uses after-tax income dividend by the households OECD equivalence scale (see text for details). Bars show two standard error bands, using robust standard errors clustered at the household level, which adjust for arbitrary within-household correlations and heteroskedasticity.

Figure 4 – Liquidity Constraints and Heterogeneity in MPCs among Lower-Income Households

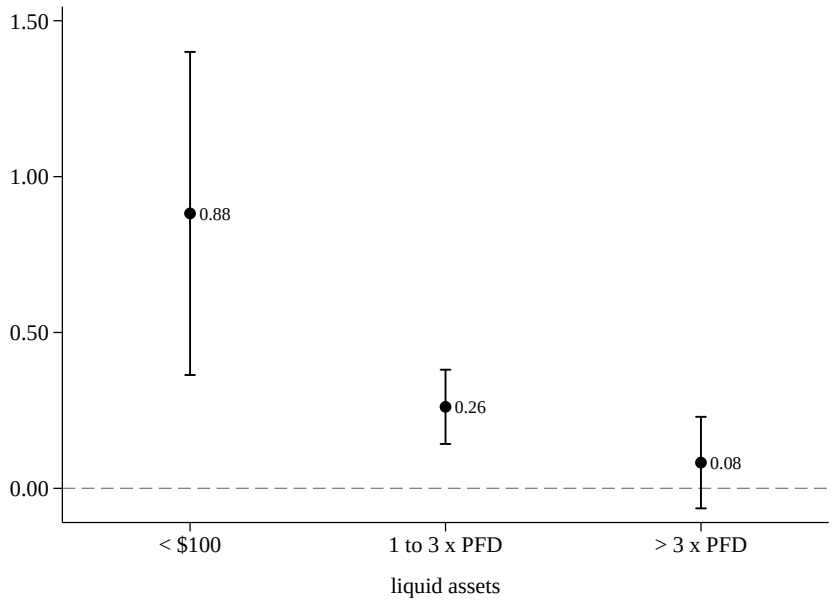
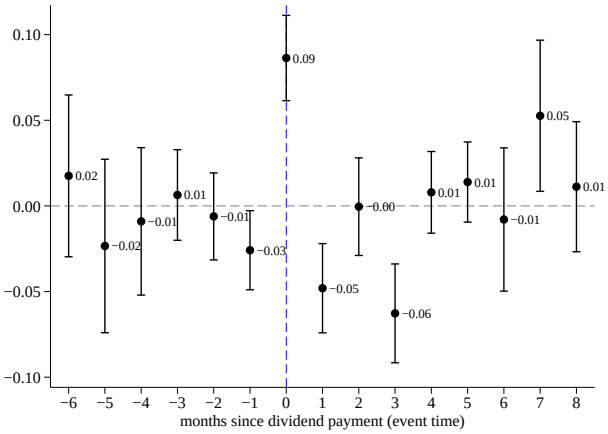
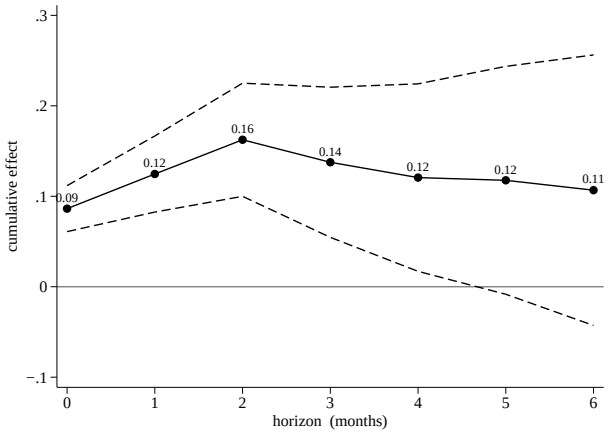


Figure 5 – Extensions

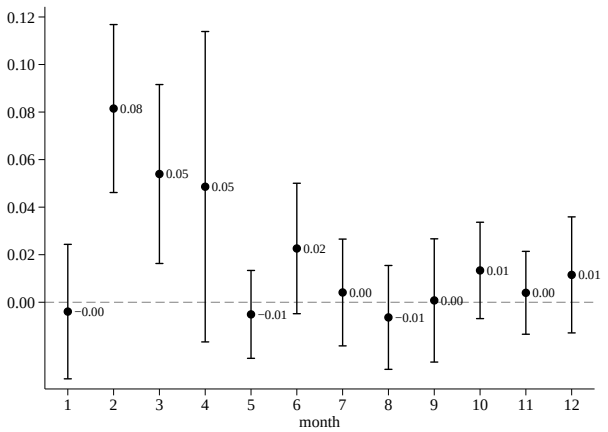
(a) small durables: regression βs



(b) small durables: cumulative MPC



(c) additional income taxes caused by dividend



(d) cons. commitments: homeowners vs renters

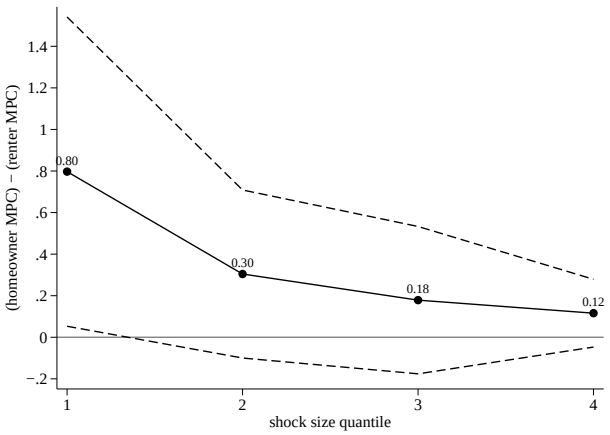
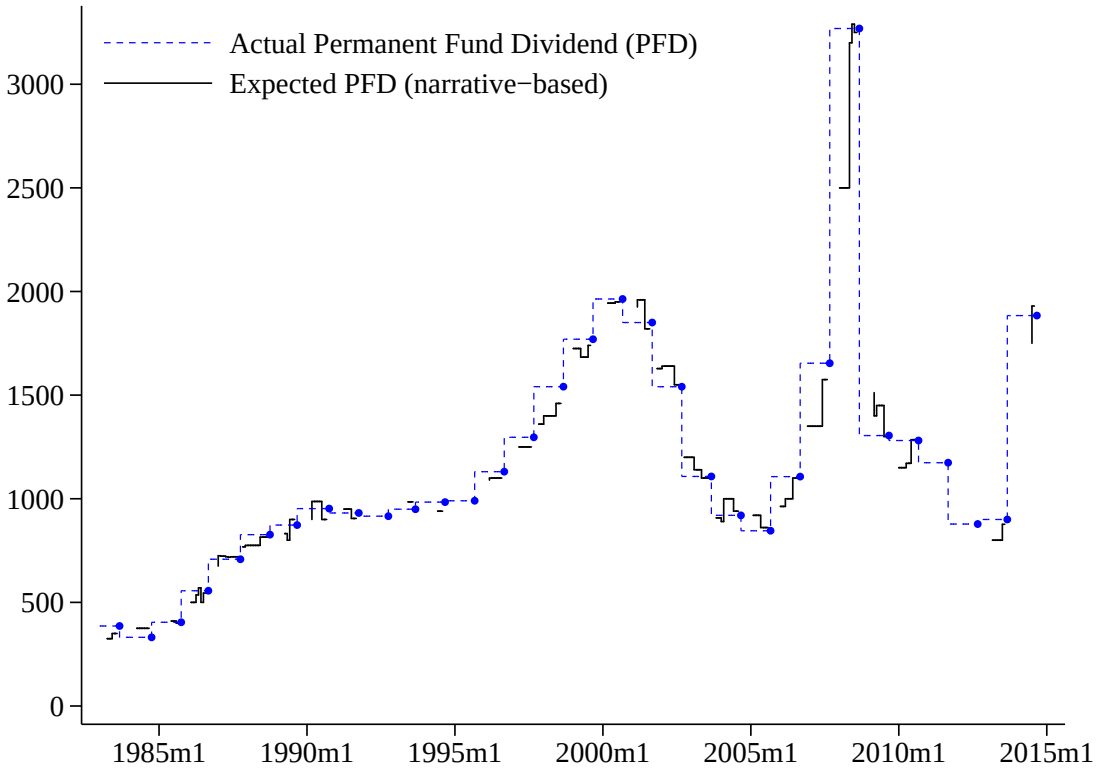


Table A1. Expenditure aggregation

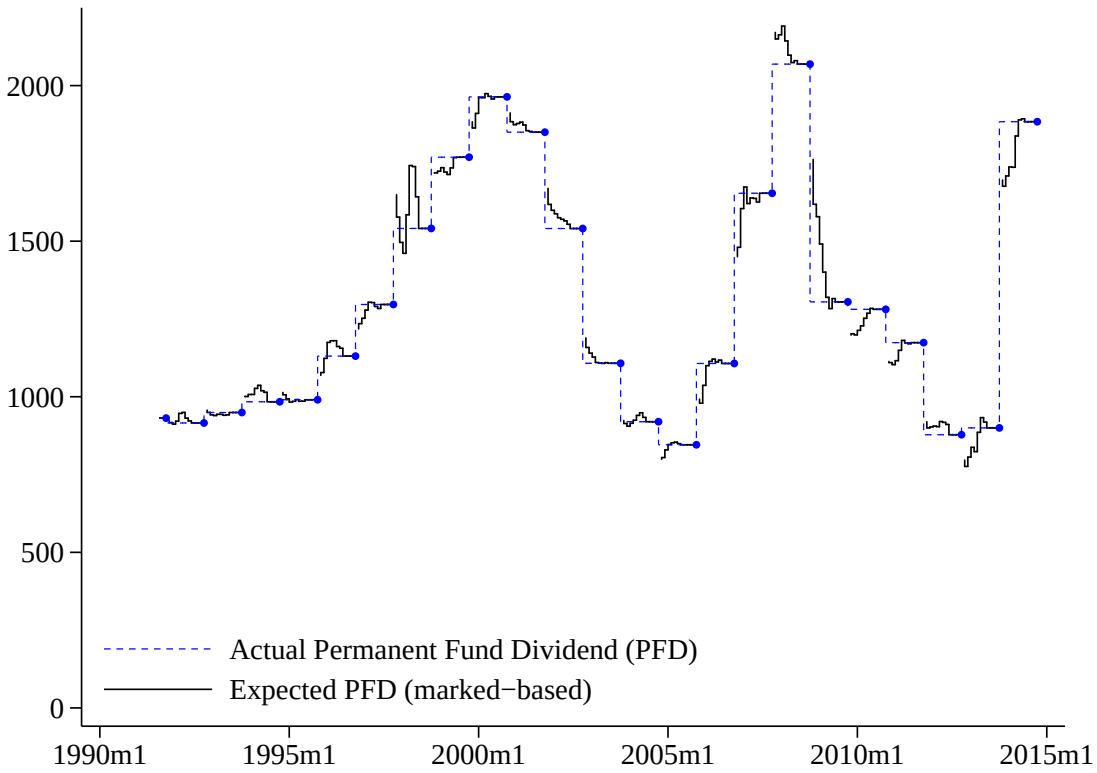
Personal Finance Website Data (PFW)				Consumer Expenditure Survey (CE)
PFW ID	Expenditure	PFW ID	Expenditure	
Non-durables:		*Services:*		*Non-durables and services:*
611	Baby Supplies	1	Entertainment	food at home
7	Food & Dining	101	Arts	food away
701	Groceries	102	Amusement	alcohol at home
704	Coffee Shops	4	Personal Care	alcohol away
706	Fast Food	403	Hair	tobacco
707	Restaurants	404	Spa & Massage	personal care services
708	Alcohol & Bars	406	Laundry	child care
1401	Gas & Fuel	5	Health & Fitness	adult care
1702	Office Supplies	501	Dentist	domestic services
		502	Doctor	gas
Durables:		503	Eyecare	electricity
103	Music	505	Pharmacy	fuel
104	Movies & DVDs	506	Health Insurance	phone
105	Newspapers & Magazines	507	Gym	water
2	Shopping	508	Sports	public transport
201	Clothing	6	Kids	vehicle services
202	Books	602	Babysitter & Daycare	vehicle insurance
204	Electronics & Software	609	Kids Activities	gasoline
206	Hobbies	9	Pets	rental cars
207	Sporting Goods	901	Pet Food & Supplies	rental furniture
606	Toys	902	Pet Grooming	clothes
1003	Books & Supplies	903	Veterinary	tailors
12	Home	10	Education	textiles
1201	Furnishings	1001	Tuition	fees and charges
1203	Home Improvement	11	Financial	occupation expenses
1208	Home Supplies	1102	Life Insurance	entertainment services
1403	Service & Parts	1105	Financial Advisor	reading material
		1202	Lawn & Garden	educational services
Other expenditures:		1204	Home Services	health insurance
603	Child Support	1206	Home Insurance	health care services
610	Allowance	13	Bills & Utilities	life insurance
8	Gifts & Donations	1301	Television	home maintenance
801	Gift	1302	Home Phone	home repairs
802	Charity	1303	Internet	home management
1002	Student Loan	1304	Mobile Phone	home security
1207	Mortgage & Rent	1306	Utilities	home insurance
1404	Auto Payment	14	Auto Services & Transport	parking
16	Fees & Charges	1402	Parking	
1601	Service Fee	1405	Auto Insurance	
1602	Late Fee	1406	Public Transportation	
1604	Finance Charge	15	Travel	
1605	ATM Fee	1501	Air Travel	
1606	Bank Fee	1502	Hotel	
1607	Trade Commissions	1503	Rental Car & Taxi	
20	Uncategorized	1504	Vacation	
2001	Cash & ATM	17	Business Services	
		1701	Advertising	
		1703	Printing	
		1704	Shipping	
		1705	Legal	

Figure A1 – Expected vs. actual Permanent Fund Dividend

(a) narrative-based expected dividend



(b) market-based expected dividend



Notes: This figure reproduces the expected dividends from Kueng (2015a), which are (i) based on a narrative analysis of all major Alaskan newspapers and (ii) based on monthly income from the fund's asset obtained from the archives and the website of the Alaska Permanent Fund Corporation (APFC). See Kueng (2015a) for more detail about the two series of expected dividends. The figure shows that changes in expected dividends using either measure are an order of magnitude smaller than the dividend itself.