2014

Essays in Consumption, Behavioral and Applied Microeconomics

Christina Yiwei Zhang
University of Pennsylvania, christina.y.zhang@gmail.com

Essays in Consumption, Behavioral and Applied Microeconomics

Abstract

This dissertation consists of three self-contained chapters which broadly address issues related to individual and household decision making. The first essay investigates household budgeting and cash flow management decisions. Specifically, I focus on household consumption decisions when there is a mismatch between the frequency at which their income (pay) arrives and the frequency for which they make their consumption decisions. Exploiting monthly variation in income arising from bi-weekly pay schedules, I find that household expenditures respond significantly to such variation with the response coming almost entirely from durables. These responses cannot be explained by the presence of binding liquidity constraints. I consider several behavioral explanations including time inconsistency with sophistication, mental accounting, and a new model of budgeting heuristics in which individuals naively extrapolate their current income into the future.

The second essay, joint with Judd Kessler and Katherine Milkman, explores public recognition in a charitable giving context and provides new empirical evidence on how public recognition can increase pro-social behavior. Using observational data on alumni giving to a large university, we find that the enactment of recognition programs for consecutive giving increased both the probability of giving and the dollar amount donated directly to the University. Furthermore, such recognition crowded in donations to other University priorities. Exploiting differences in the timing of various recognition programs, we find that while individual value public recognition generally, they value recognition more so when such recognition can be used to convey information regarding personal traits of the donor.

Finally, the third essay, joint with Jeremy Tobacman, examines the determinants of intrinsic motivation. While intrinsic motivation has important applications in many areas, it is often difficult to separate intrinsic motivation from extrinsic motivation due to potential correlation between pecuniary and non-pecuniary incentives. Using an online word-hunting game, we isolate the role of performance as a source of intrinsic motivation and study how past performance affects an individual's persistence of play. Using individual variation in the ability to exploit the presence of prefixes and suffixes, we estimate the causal effect of past performance on the length of a game spell (continued play) and the probability of ending a spell. We find that higher past performance significantly increases spell length and significantly decreases the probability of a spell ending.

Degree Type

Dissertation

Degree Name

Doctor of Philosophy (PhD)

Graduate Group

Applied Economics

First Advisor

Jeremy B. Tobacman

Keywords

Behavioral Economics, Charitable Giving, Household Finance, Intrinsic Motivation

ESSAYS IN CONSUMPTION, BEHAVIORAL AND APPLIED MICROECONOMICS

Christina Yiwei Zhang

A DISSERTATION

in

Applied Economics

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2014

Supervisor of Dissertation

Jeremy Tobacman, Assistant Professor of Business Economics and Public Policy

Graduate Group Chairperson

Eric Bradlow, Professor of Marketing, Statistics and Education

Dissertation Committee

Jeremy Tobacman, Assistant Professor of Business Economics and Public Policy

Todd Sinai, Associate Professor of Real Estate and Business Economics and Public Policy

Judd Kessler, Assistant Professor of Business Economics and Public Policy

Daniel Gottlieb, Assistant Professor of Business Economics and Public Policy

Nikolai Roussanov, Assistant Professor of Finance

ESSAYS IN CONSUMPTION, BEHAVIORAL AND APPLIED MICROECONOMICS

To Mom and Dad,

whose support, encouragement, and love never failed

and to whom I owe everything.

ACKNOWLEDGEMENT

This dissertation would not have been possible without the effort and support of many people. I am incredibly grateful to Jeremy Tobacman who has been my main advisor and, more importantly, my mentor during my graduate studies. I will always remember the time I brought preliminary results for what would eventually be my job market paper to his office, and he ended the meeting with "Congratulations! You have a job market paper!" The amount of time and effort Jeremy has dedicated to meeting with me, reading over drafts, and being a general sounding board while I tried to work things out over the past few years is simply unbelievable. I can only hope that other students will be so lucky as to have a mentor like I had with Jeremy. I am also forever grateful to Judd Kessler, Daniel Gottlieb, Nick Roussanov, Todd Sinai, and Katy Milkman who supported me over the past few years and without whom, I would never have survived graduate school.

I am indebted to the many students in my program who celebrated (and sometimes commiserated) with me over the past five years. Graduate school would never have been the same without Hessam Bavafa, William Mann, Anita Mukherjee, Mike Punzalan, and Dan Sacks. I will never forget cabin trips, late nights at the office, and a weekly happy hour rotation for our "Dissertator's Anonymous" group which was mostly an excuse to just hang out. I of course cannot imagine my graduate school experience without Anthony DeFusco whose patience, love, and willingness to go over entire presentations and proofread papers in the early hours of the morning when I'm scrambling to get things done can never be repaid. To him I owe a special thanks for being the best thing other than my PhD that I got out of my graduate school experience and for always being the other half of my team.

Finally, I owe everything to my family, without whom I would never have reached where I am today. I am grateful for the never-ending love and support from my parents and for both their complete faith in me as well as their reassurance that if all things fell apart, I could always come home. They have always been my inspiration. If I can ever be even half

as amazing as they are, I will know that I have done well. And last but certainly not least, I am eternally grateful to be so lucky as to have as amazing a brother as Hang-Hang whose constant teasing and love always provides me with perspective on what truly matters in life.

ABSTRACT

ESSAYS IN CONSUMPTION, BEHAVIORAL AND APPLIED MICROECONOMICS

Christina Yiwei Zhang

Jeremy Tobacman

This dissertation consists of three self-contained chapters which broadly address issues related to individual and household decision making. The first essay investigates household budgeting and cash flow management decisions. Specifically, I focus on household consumption decisions when there is a mismatch between the frequency at which their income (pay) arrives and the frequency for which they make their consumption decisions. Exploiting monthly variation in income arising from bi-weekly pay schedules, I find that household expenditures respond significantly to such variation with the response coming almost entirely from durables. These responses cannot be explained by the presence of binding liquidity constraints. I consider several behavioral explanations including time inconsistency with sophistication, mental accounting, and a new model of budgeting heuristics in which individuals naively extrapolate their current income into the future.

The second essay, joint with Judd Kessler and Katherine Milkman, explores public recognition in a charitable giving context and provides new empirical evidence on how public recognition can increase pro-social behavior. Using observational data on alumni giving to a large university, we find that the enactment of recognition programs for consecutive giving increased both the probability of giving and the dollar amount donated directly to the University. Furthermore, such recognition crowded in donations to other University priorities. Exploiting differences in the timing of various recognition programs, we find that while individual value public recognition generally, they value recognition more so when such recognition can be used to convey information regarding personal traits of the donor.

Finally, the third essay, joint with Jeremy Tobacman, examines the determinants of intrinsic motivation. While intrinsic motivation has important applications in many areas, it is often difficult to separate intrinsic motivation from extrinsic motivation due to potential correlation between pecuniary and non-pecuniary incentives. Using an online word-hunting game, we isolate the role of performance as a source of intrinsic motivation and study how past performance affects an individual's persistence of play. Using individual variation in the ability to exploit the presence of prefixes and suffixes, we estimate the causal effect of past performance on the length of a game spell (continued play) and the probability of ending a spell. We find that higher past performance significantly increases spell length and significantly decreases the probability of a spell ending.

TABLE OF CONTENTS

ACKNOWLEDGEMENT iv

ABSTRACT vi

LIST OF TABLES xii

LIST OF ILLUSTRATIONS xiii

CHAPTER 1 : Consumption Responses to Pay Frequency: Evidence from 'Extra' Paychecks 1

1.1 Introduction 1

1.2 Misalignment in a Consumption-Savings Framework 6

1.3 Data 8

1.4 Empirical Methodology 11

1.5 Results 13

1.6 Potential Explanations 19

1.7 Conclusion 30

CHAPTER 2 : Social Recognition in Charitable Giving: In Pursuit of Perfection . 50

2.1 Introduction 50

2.2 Data Description 54

2.3 Empirical Setup 58

2.4 Effect of Recognition for Consecutive Giving on Donation Behavior 59

2.5 Do donors exhibit strategic behavior in response to recognition programs? . 62

2.6 Recognition versus Information 64

2.7 Discussion 66

CHAPTER 3 : Play for Performance 90

3.1 Introduction 90
3.2 Game Background and Data 91
3.3 Relationship between Performance and Game Board Features 93
3.4 Estimation Framework and Results 94
3.5 Conclusion 99

APPENDIX 117

BIBLIOGRAPHY 138

LIST OF TABLES

TABLE 1 : Summary Statistics 35

TABLE 2 : Response to Extra Paychecks by Consumption Categories 36

TABLE 3 : Response to Extra Paychecks by Durable Consumption Categories 37

TABLE 4 : Response to Extra Paychecks by Vehicle Consumption Categories . 38

TABLE 5 : Summary Statistics for Car Purchases 39

TABLE 6 : Response to Extra Paychecks by Vehicle Financing Categories . . . 40

TABLE 7 : Response to Extra Paychecks by Vehicle Financing Categories Conditional on Ever Purchasing a Car 41

TABLE 8 : Response to Extra Paychecks by Pay Frequency across Consumption Categories 42

TABLE 9 : Response to Extra Paychecks by Consumption Categories with Time Adjustment 43

TABLE 10 : Response to Extra Paychecks by Consumption Categories by Liquid Asset Holdings 44

TABLE 11 : Response to Extra Paychecks by Consumption Categories by Total Before-Tax Income 45

TABLE 12 : Response to Extra Paychecks by Consumption Categories by Age . 46

TABLE 13 : Response of Aggregate Consumption Measures by Committed Consumption as Fraction of Monthly Wages 47

TABLE 14 : Response to Extra Paychecks by Consumption Categories by Consumption Volatility 48

TABLE 15 : Comparison of Vehicle Loan Characteristics by Timing of Purchase 49

TABLE 16 : Summary Statistics 73

TABLE 17 : Categories of Support ($) by Years Since Graduation 74

TABLE 18 : Donations at Lower Bound (Bunching) 75

TABLE 19 : Summary Statistics by Ivy Stone Society (ISS) Eligibility 76
TABLE 20 : Summary Statistics by Perfect Giving (PG) Eligibility 77
TABLE 21 : Effects of PG Recognition Program on Giving with No Dynamics - 2002 to 2005 78
TABLE 22 : Effects of PG Recognition Program on Giving with No Dynamics - 2004 to 2005 79
TABLE 23 : Effects of ISS Recognition Program on Giving with No Dynamics - 2002 to 2007 80
TABLE 24 : Effects of ISS Recognition Program on Giving with No Dynamics - 2006 to 2007 81
TABLE 25 : Effects of ISS and PG Recognition Programs on Giving Accounting for Dynamics - Young Alumni 82
TABLE 26 : Effects of ISS Recognition Program on Giving Accounting for Dynamics - Non Young Alumni 83
TABLE 27 : Effects of PG Recognition Program on Bunching with No Dynamics 84
TABLE 28 : Effects of ISS Recognition Program on Bunching with No Dynamics 85
TABLE 29 : Effects of ISS and PG Recognition Programs on Bunching Accounting for Dynamics - Young Alumni 87
TABLE 30 : Effects of ISS Recognition Program on Bunching Accounting for Dynamics - Non Young Alumni 88
TABLE 31 : Effect of Introduction of Ivy Stone Society by Years of Consecutive Giving 89
TABLE 32 : Individual Level Summary Statistics 110
TABLE 33 : Team Level Summary Statistics 111
TABLE 34 : Impact of Player Performance on Spell Length 112
TABLE 35 : Impact of Player Performance on Spell End 113
TABLE 36 : Impact of Player Performance on Spell Length and Spell End Accounting for "Goal" Cutoffs 114

TABLE 37 : Impact of Team Performance on Spell Length 115
TABLE 38 : Impact of Team Performance on Spell End 116
TABLE 39 : Sample Restriction Details 119
TABLE A.2.1 The Timing of Three Paycheck Months from 1989-2010 125
TABLE A.3.1 Response to Extra Paychecks by Consumption Categories Using Log Changes in Consumption 128
TABLE A.3.2 Response to Extra Paychecks by Consumption Categories Using Level Changes in Income 129
TABLE A.4.1 Response to Extra Paychecks by Consumption Categories by Marital Status 130
TABLE A.4.2 Response to Extra Paychecks by Consumption Categories by Race 131
TABLE A.4.3 Response to Extra Paychecks by Consumption Categories by Gender 132
TABLE A.4.4 Response to Extra Paychecks by Consumption Categories by Housing Ownership 133

LIST OF ILLUSTRATIONS

FIGURE 1 : Total Income Distribution by Pay Frequency 32

FIGURE 2 : Timing of Durable and Strictly Non-Durable Spending 33

FIGURE 3 : Distribution of Placebo Estimates of Spending Effect 34

FIGURE 4 : Honor Roll Sample 68

FIGURE 5 : Timeline of Recognition Programs 69

FIGURE 6 : Donation Frequency for Amounts between $200 and $5000 69

FIGURE 7 : Eligibility of Alumni by Recognition Programs 70

FIGURE 8 : Mean Residuals for Giving over Time for Eligible and Ineligible . 71

FIGURE 9 : Effect of Introduction of Ivy Stone Society by Years of Consecutive Giving 72

FIGURE 10 : 4x4 Game Board 100

FIGURE 11 : Games Per Player 101

FIGURE 12 : Spells Per Player 102

FIGURE 13 : Length of Spell 103

FIGURE 14 : Days Between Spells 104

FIGURE 15 : Points Per Player in First Game 105

FIGURE 16 : Points Per Player in All Games 106

FIGURE 17 : Linear Prediction of Spell Length on Player Points in First Game of Spell 107

FIGURE 18 : Linear Prediction of Spell End on Player Points in First Game of Spell 108

FIGURE 19 : Learning Based on Average Player Points 109

FIGURE A.2.1 Simulated Bias 126

CHAPTER 1 : Consumption Responses to Pay Frequency: Evidence from 'Extra' Paychecks

1.1. Introduction

One of the hallmark implications of the life-cycle/permanent income hypothesis (LC-PIH) is that household consumption should not respond to anticipated income changes (Friedman 1957; Modigliani 1988). Whether this implication of the theory holds true has generated significant debate among economists and has given rise to an extensive empirical literature testing for excess sensitivity to anticipated income. However, results from these empirical studies vary substantially.[1] At present, no consensus exists in the literature on the extent to which households may respond to anticipated income changes or on the mechanisms underlying observed responses.

This paper contributes new evidence on household consumption responses to anticipated income by examining an important household financial problem: how to adjust the timing of consumption to the timing of income. For many households, the frequency for which they make their consumption decisions often differs from that at which they receive their pay. This misalignment can result in predictable variation in the amount of income received per consumption decision period. The LC-PIH predicts that consumption should not respond to such variation. I empirically test this prediction by exploiting variation in monthly income arising from the timing of bi-weekly pay schedules. Bi-weekly workers are paid on a regular two-week schedule and receive 26 paychecks over the course of the year. Because these 26 paychecks must be disbursed over 12 months, bi-weekly workers typically receive two paychecks per month with the exception of two months out of the year, during which they receive three.[2] As a result, the level of wage and salary income that a household receives during these atypical months with three paychecks is higher than during typical months with only two. The timing of bi-weekly pay schedules thus provides predictable variation in monthly income while holding constant both total lifetime income and the environment in which that income is received. This is in contrast to semi-monthly or monthly pay schedules through which workers receive the same income each month.

My focus on the variation generated by bi-weekly pay schedules is motivated by two main considerations. First, this variation in income (and the extent to which there is an "extra" third paycheck) is simply an artifact of evaluating income on a monthly basis. For a bi-

[1] See Browning and Lusardi (1996), Browning and Crossley (2001), and Jappelli and Pistaferri (2010) for surveys of the literature.

[2] On occasion, the calendar year is such that a bi-weekly worker will actually receive two paychecks per month with the exception of *three* months out of the year, during which they receive three. This occurs approximately once every eleven years and is accounted for in the analysis of this paper.

weekly worker there is predictable variation in *monthly* income but not in twice-weekly income. While it is not difficult to imagine bi-weekly households thinking or budgeting on a monthly basis, the fact that whether or not a particular paycheck can be considered "extra" is not a priori assumed makes the evidence I provide in this paper even more compelling.[3] Second, in contrast to much of the existing literature, the third paychecks are not referred to as a bonus, special payment, or in any other manner which might induce bi-weekly workers to differentially respond following their receipt. The concern with such special designations is that they may naturally lead households to categorize or evaluate the payments separately from other income. In my setting, the only way in which these third paychecks are "special" is if people choose to evaluate their income at a monthly frequency.

I test for excess sensitivity in months following three paycheck months using panel data from the Consumer Expenditure Survey (CEX). I first identify households whose heads are paid at a bi-weekly frequency and then determine the months during which they receive three paychecks. The empirical strategy takes advantage of the fact that the months during which bi-weekly workers receive three paychecks differ from year to year. For example, a bi-weekly worker paid in the first week of January in 2008 would have received three paychecks in February and August of that year. In 2009, that same worker would have received three paychecks in January and July. The causal effect of these third paychecks is identified using a difference-in-differences strategy that compares spending responses following a given calendar month in years during which there are three paychecks distributed during that month to responses in years during which there are only two paychecks during that month.[4]

Using this identification strategy, I establish two main empirical results. First, I find that total household spending increases by approximately \$262 (in 2010 dollars) on average in the month following a three paycheck month and that this effect on spending does not persist in subsequent months. Second, I find that this spending increase is due entirely to changes in durable spending, and specifically new vehicle purchases, with no corresponding response in non-durables. Conditional on purchasing a vehicle during the interview period, household spending on vehicles increases by roughly \$2205 dollars on average following three paycheck months. This represents a fairly sizable effect given that the average gross paycheck for a

[3] Many recurring expenditures, such as mortgage payments and utility bills, are due monthly, so consumption tends to be budgeted on a month-to-month basis (Thaler 2008). From the 2010 wave of the Survey of Consumer Finances, over 96 percent of surveyed individuals make mortgage, rent, education loan, and vehicle loan payments monthly (conditional on making payments).

[4] This empirical strategy uses variation within months but across households because the CEX only provides income and expenditure data for at most twelve consecutive months for any given household that is interviewed. Alternatively, I could compare responses within a household across months. Employing this procedure, however, would not allow me to control for seasonal variation in consumption across months. Estimates using variation within households and across months are slightly larger than those controlling instead for seasonal variation across months.

bi-weekly worker in my sample is approximately $1669. These results are consistent with several other papers in the literature on responses to anticipated income receipt which also find large responses in durable spending (Parker 1999; Souleles 1999; Adams et al. 2009; Parker et al. 2011; Aaronson et al. 2012). The results are robust to changes in sample composition as well as to variants of the main specification. These findings suggest that, contrary to the predictions of the LC-PIH, individuals do in fact respond to the variation induced by misalignment between the timing of consumption and the timing of income.

These results rely on the identification assumption that changes in consumption expenditures following a given calendar month in year in which there were three paychecks distributed during that month versus years in which there were only two would have evolved similarly were it not for the third paycheck. To assess the validity of this assumption, I employ a difference-in-difference-in-differences (triple difference) research design using households paid at other frequencies as controls. In contrast to bi-weekly workers, I find no corresponding effect for workers who are paid either weekly or monthly. To further evaluate the robustness of my main estimates, I also conduct a series of placebo tests where I re-estimate the main specification using randomly generated schedules of third paychecks. These tests show that the probability of finding an effect as large as I do by chance is extremely small.

I explore four potential explanations for why I find evidence of excess sensitivity: liquidity constraints, budgeting heuristics, time inconsistency with sophistication, and mental accounting. For each, I discuss whether the empirical findings can be reconciled with the predictions of the proposed model of behavior.

Liquidity constraints are perhaps the most often cited explanation for observed spending responses to anticipated income receipt. To the extent that these constraints are binding, spending responses may be due to the third paychecks providing additional liquidity upon receipt. Since a household's liquidity is unobserved, I test for evidence of binding constraints using two complementary strategies. First, I proxy for liquidity using four measures: liquid assets, before-tax income, age, and committed expenditures as a fraction of monthly wages. Using these measures, I estimate responses separately for constrained and unconstrained households. As a second empirical strategy, I classify households as constrained or unconstrained by the extent to which their observed choices are consistent with binding constraints. To do so, I focus on two types of observed behavior: consumption volatility and vehicle loan characteristics. These tests provide evidence that is at best weakly supportive of liquidity constraints as an explanation for the empirical findings. However, the imprecision of the estimates and the lack of response in non-durable spending following months with three paychecks suggests that liquidity constraints alone cannot explain the observed responses.

The second explanation – budgeting heuristics – is motivated by the fact that people often use heuristics to alleviate the cognitive burden of complicated decisions problems. Such complexity may arise for a number of reasons, including the presence of variation in income. With this in mind, I propose a specific heuristic in which households construct monthly budgets by naively extrapolating their current monthly income into the future. This behavior is similar in spirit to the model of projection bias developed by Loewenstein et al. (2003) and relates biased beliefs to the rule-of-thumb behavior of consumers in Campbell and Mankiw (1989). When the frequency of consumption and the frequency of income are sufficiently misaligned, this budgeting heuristic model predicts excess sensitivity of consumption in months with atypical income. Allowing for two types of consumption, durable and non-durable, the model also predicts that budgeting heuristics will lead some households to increase their durable consumption in months with atypical income while having an ambiguous effect on non-durable consumption.

The final two explanations that I consider are motivated by existing theories from behavioral economics. First, I discuss the possibility that individuals are sophisticated with (β, δ)-preferences (Strotz 1956; Phelps and Pollak 1968; Laibson 1997; O'Donoghue and Rabin 1999). Such individuals are time-inconsistent with present-biased preferences and exhibit problems with self-control. As a result, they may choose to invest their third paychecks in durable goods because the illiquid nature of these goods acts as a commitment mechanism against overconsumption (Laibson 1997). However, although this could explain the existence of an increase in durable spending following three paycheck months, it is difficult to reconcile with the fact that the observed response is driven primary by *debt-financed* purchases. Such purchases commit a household to a stream of installment payments which, given rational income expectations, they may not be able to afford in the future. Second, I discuss the possibility that individuals use a system of mental accounts to categorize and evaluate their income (Thaler and Shefrin 1981; Shefrin and Thaler 1988; Thaler 1990, 1999, 2008). Bi-weekly workers may choose to treat their typical income of two paychecks per month as their regular income while viewing third paychecks as a "bonus" or windfall gain. This type of mental accounting could explain consumption responses to third paychecks even in the absence of any change in total lifetime income and with correct beliefs about future income. However, given rational expectations about future income, it is still difficult to reconcile with the empirical results indicating an increase in debt-financed vehicle purchases. A related explanation is that bi-weekly workers with earmark third paychecks specifically for other uses, such as the purchase of a large durable. Mentally placing that income as otherwise off-limits would allow workers to ensure they save enough income in order to purchase a vehicle. Such behavior would be consistent with the observed spending responses in durable consumption following months with three paychecks.

This research builds upon and contributes to the extensive literature using household-level micro-data to examine consumption responses to various types of anticipated income receipt. The breadth of the literature reflects both the general interest in understanding and cleanly identifying consumption responses to predictable or transitory changes in income, and the importance of estimating the causal effect of various fiscal policies that provide payments to households. The sources of income receipt that are typically analyzed include both changes to permanent income (Wilcox 1989; Paxson 1993; Shapiro and Slemrod 1995; Shea 1995; Lusardi 1996; Parker 1999; Souleles 2002; Stephens, Jr. 2008; Aaronson et al. 2012) and predictable one-time payments such as tax refunds or stimulus payments (Souleles 1999; Browning and Collado 2001; Hsieh 2003; Johnson et al. 2006, 2009; Agarwal et al. 2007; Parker et al. 2011). I contribute to this literature by leveraging the variation in income arising from misalignment, which allows me to look at consumption responses to the timing of income flows using a source of anticipated income that is completely "designation-free." Using this variation, I am able to test whether households are excessively sensitive to anticipated income receipt in an environment where the income being received is not likely to be labeled or categorized in any special way, thus lending a high degree of external validity to my findings.

This paper also complements a related literature examining high-frequency responses to payments from a constant periodic income stream such as a paycheck or government transfer check (Stephens, Jr. 2003; Shapiro 2005; Huffman and Barenstein 2005; Stephens, Jr. 2006; Stephens, Jr. and Unayama 2011). While this paper also examines high-frequency responses, its focus is somewhat different. The papers in this literature are primarily concerned with the path of consumption over the course of a given pay period. Specifically, they look at whether households consumption smooth between anticipated payments and find that consumption significantly responds immediately following receipt before then declining. In contrast, I look at the path of consumption *across* pay periods and in response to periods with atypical income.

Finally, I contribute to the large and growing literature on heuristics and biases (Tversky and Kahneman 1974; Laibson 1997; Barberis et al. 1998; Rabin and Schrag 1999; Loewenstein et al. 2003; Gabaix et al. 2006; Pope and Schweitzer 2011; Lacetera et al. 2012; Hastings and Shapiro 2013; Kőszegi and Szeidl 2013).[5] Most of the literature on household budgeting and mental accounting has focused on the categories to which expenditures or funds are allocated

[5]See Gilovich et al. (2002) and DellaVigna (2009) for more examples.

while giving relatively little attention to the *frequency* at which accounts are evaluated. In contrast, I focus on the frequency of consumption decisions and specifically whether it differs from the frequency at which income arrives (i.e. whether there is misalignment). I show that when there is misalignment, adopting a budgeting heuristic can have important implications for consumption. To the best of my knowledge, this is the first paper to develop a formal model of a temporal budgeting heuristic and provide empirical evidence in its support.

The rest of the paper proceeds as follows. Section 1.2 introduces misalignment more formally. Section 1.3 provides a brief description of the Consumer Expenditure Survey data and is followed by a discussion in Section 1.4 of the empirical methodology. Section 1.5 presents the main results of the empirical analysis in this paper. Section 1.6 considers several potential explanations and finally, Section 1.7 concludes.

1.2. Misalignment in a Consumption-Savings Framework

This section provides a more formal description of how misalignment between the timing of consumption and the timing of income can lead to variation in income and discusses the implications of that variation for the path of consumption in a standard consumption and savings framework. While in later sections, I allow for additional refinements to the standard framework based on more recent work in consumption theory, for now, I adopt a standard model and assume perfect capital markets and no income uncertainty. This model is commonly used in the empirical literature and provides a useful benchmark with which to approach the empirical analysis.

To begin, consider an infinitely-lived individual, with discount factor δ, who chooses consumption, c_t, each period to maximize her expected total remaining lifetime utility

$$U_t(\{c_s\}_{s=t}^{\infty}) = E_t\left[\sum_{s=t}^{\infty} \delta^{s-t} u(c_s)\right], \tag{1.1}$$

where $u(c)$ satisfies the conditions: $u_c \to \infty$ as $c \to 0$ and $u_c \to 0$ as $c \to \infty$. For my purposes, it is easiest to think of time as a continuous measure that is divisible into periods of varying lengths. For example, if t indexes years, then we can think of a period t as being of "length" equal to one year and of the frequency for which consumption decisions are being made as yearly.

The individual receives income, y_τ, where I allow for the possibility that income is received at a different frequency than the frequency for which consumption decisions are made (i.e. a τ-period need not refer to the same length of time as a t-period). In general, individual income can be aggregated to the frequency for which consumption decisions are made with the simple transformation $y_t = \sum_{\tau \in [t,t+1)} y_\tau$. When income is received at the same frequency as the frequency for which consumption decisions are made, this transformation reduces to the expression $y_t = y_\tau$.

I assume that individuals receive a fixed payment, $y_\tau = \bar{y}$, as income each τ-period. This assumption reflects the fact that pay is often disbursed regularly and in fixed amounts. Using the previously mentioned transformation, the individual's t-period income, evaluated at the frequency at which she makes her consumption decisions, can then always be expressed as

$$y_t = \begin{cases} Y & \text{if } (t-1) \bmod n \neq 0 \\ (1+b)Y & \text{if } (t-1) \bmod n = 0 \end{cases}, \tag{1.2}$$

for some $n \geq 1$ and $b \geq 0$. That is, in most periods, the individual receives a fixed payment of Y ("typical" income); every n periods, however, she receives a payment of $(1+b)Y$ ("atypical" income) instead.[6] Under this framework, misalignment between the timing of consumption and the timing of pay exists when b is strictly greater than zero and n is strictly greater than one. That is, t-period income varies when the timing of consumption and the timing of pay are misaligned even though τ-period income (pay) is constant.

In the context of this paper, an individual who is paid bi-weekly but makes her consumption decisions monthly has misalignment between the timing of her consumption and the timing of her income. Specifically, her period t income is given by Equation 1.2, where each period t is a month, typical income is equal to two paychecks ($Y = 2\bar{y}$), and atypical income is a result of the third paycheck ($bY = \bar{y}$) that is received every $n = 6$ months.[7] The same bi-weekly worker would have no misalignment between the timing of her consumption and the timing of her pay if she were to instead make her consumption decisions at a bi-weekly frequency.

1.2.1. Basic Environment with Misalignment

The standard assumptions of the LC-PIH predict that misalignment should not matter. To illustrate this, consider once again an infinitely-lived individual with discount factor δ. The individual has CRRA preferences and maximizes expected utility given by Equation 1.1 with income given by Equation 1.2. Following Deaton (1991), I define cash-on-hand, X_t, as the sum of current income and assets. Let R be the gross interest rate which is assumed to

[6] While the above is intended to capture features of pay schedules, more generally, the income bY can be thought of as an infrequent recurring payment. Equation 1.2 can thus be extended to include contexts such as the receipt of a yearly bonus or tax refund.

[7] The frequency at which third paychecks arrive is not quite every six months, but $n = 6$ provides a good approximation and simplifies the exposition of the model substantially.

be constant over time. The present discounted value of the individual's remaining lifetime wealth at the beginning of time t can then be written as

$$\begin{aligned} W_t &= X_t + \sum_{i=1}^{\infty} \frac{y_{t+i}}{R^i} \\ &= X_t + \sum_{i=1}^{\infty} \frac{Y}{R^i} + \sum_{i=1}^{\infty} \frac{bY}{R^i} \cdot (\mathbb{1}_{\{(t-1+i) \bmod n=0\}}) \end{aligned} \tag{1.3}$$

where the indicator $\mathbb{1}_{\{(t-1+i) \bmod n=0\}}$ evaluates to one in periods with atypical income and zero otherwise. Assuming the no-Ponzi condition, $\lim_{t\to\infty} \frac{X_{t+1}-y_{t+1}}{R^t} \geq 0$, is satisfied, the usual Euler equation characterizing optimality is given by

$$u_c(c_t) = \delta R u_c(c_{t+1}).$$

Given the specified preferences, the above expression can be rewritten as

$$c_{t+1} = (\delta R)^{1/\rho} c_t. \tag{1.4}$$

Assuming $(\delta R)^{1/\rho} < 1$ (some minimum impatience), then there exists a well-defined solution where, in equilibrium, consumption is proportional to lifetime wealth for all periods t. Together with the Euler equation, this fact implies that optimal consumption is given by

$$c_t = (1 - (\delta R^{1-\rho})^{1/\rho}) W_t. \tag{1.5}$$

As is standard in such setups, consumption is independent of current income and depends only on lifetime wealth. Furthermore, lifetime wealth does not depend on the frequency at which income arrives or whether this frequency differs from the frequency at which consumption decisions are made.[8] Thus, any misalignment between the timing of consumption and the timing of income should be irrelevant for the path of consumption under the standard consumption and savings framework. I test this prediction in the following empirical analysis of this paper.

1.3. Data

The empirical analysis in this paper uses data between 1990 and 2010 from the quarterly interview portion of the Consumer Expenditure Survey (CEX). The CEX is a nationally representative rotating panel survey, conducted by the Bureau of Labor Statistics (BLS), which provides detailed information on household spending as well as income and house-

[8] To see this, note that $W_t = X_t + \sum_{i=1}^{\infty} \frac{y_{t+i}}{R^i} = X_t + \sum_{i=1}^{\infty} \frac{\sum_{\tau \in [t+i, t+i+1)} y_\tau}{R^i}$.

hold characteristics. Households are interviewed every three months over five consecutive quarters. The first introductory interview gathers basic information on demographics and the stock of major durables owned by the household. During the second to fifth interviews, households are asked to recall their consumption expenditures over the previous three months as well as the month in which each expenditure occurred. While households are interviewed every quarter, the design of the survey effectively provides monthly data on household spending.[9]

In addition to including extensive information on consumption expenditures, the CEX also collects income and employment information for each household member (aged 14 or older) during their second and fifth interview. Crucially, the survey indicates both the gross amount of each household member's last pay as well as the length of time this last gross pay covers. I use this information to identify the frequency with which household members are paid and in particular whether a member is paid bi-weekly.

For the baseline empirical analysis, I focus on five aggregated measures of consumption expenditures. As a first measure, I examine household expenditures on food, which includes both food consumed at home and food consumed away from home, as well as any consumption of alcoholic beverages. Expenditures on food have been the focus of a number of papers studying consumption behavior, in part because of the often limited information on non-food consumption in commonly used longitudinal data sets, such as the Panel Study of Income Dynamics.[10] However, focusing on food consumption alone has obvious limitations. For this reason, I also consider the broader category of non-durable goods and services as a second measure of consumption expenditure. Because the empirical analysis focuses primarily on consumption expenditures over relatively short time-horizons, I further include a subset of non-durable expenditures commonly referred to as "strictly non-durable" (following Lusardi 1996). This is the subset of non-durable goods and services excluding those which may be considered semi-durable in the short run, such as apparel or reading materials. The fourth consumption category is the set of expenditures on durable goods, which includes expenditures on major appliances, flooring, furniture, shelter, and vehicle purchases. The final measure of consumption expenditures I consider, total expenditures, is simply the sum of durable and non-durable goods and services. Changes in these five measures serve as the main outcomes of interest in the baseline empirical analysis.

[9]Because I am interested in estimating spending responses in *months* following three paycheck months, I use the information provided by households on the month of expenditure to construct expenditure data on the household-monthly level rather than the household-quarterly level. As such, time adjustment routines used by the BLS when constructing the CEX data have important consequences for estimation. These routines and their implications for my analysis are discussed in detail in Section 1.5.3 of the paper.

[10]See Hall and Mishkin (1982), Zeldes (1989), Runkle (1991), and Shea (1995) for some examples of papers focusing on food consumption in testing the LC-PIH.

Appendix A.1 provides a detailed description of the sample selection procedures which follow the existing literature closely. In addition to the standard restrictions imposed by the rest of the literature, I make two further restrictions that bear mention. First, I restrict the sample to only those households whose heads report working full time over the full past year (at least 50 weeks) and whose reported gross pay has not changed between the second and fifth interview. Because the CEX does not include direct information on job tenure, this restriction attempts to select for households who are likely to have been employed at the same job during all of the past year. While income information is only obtained in the second and fifth interview, it is plausible that individuals receive the same amount of pay in the other interview quarters during which income information is not collected. For workers paid bi-weekly, those who have held their job for the past year are more likely to be aware of the presence and timing of the third paychecks.

The second restriction excludes households where other employed members of the household are paid at the same frequency as the head of household. This restriction reflects an important limitation of the data. While the CEX provides information on the gross amount and frequency of the last pay, I am unable to observe the actual date on which this last pay occurred. As a result, it is not possible to distinguish which of the two possible alternate schedules by which a bi-weekly worker can be paid applies to any given member. For example, a bi-weekly worker paid during the first week of January in 2008 and then every two weeks afterwards is paid on an alternate schedule to that of a bi-weekly worker paid during the *second* week of January and then every two week afterwards. The first bi-weekly worker would receive three paychecks in February and August of that year whereas the second bi-weekly worker would receive three paychecks in May and October. Which schedule a bi-weekly worker follows therefore determines the months of the year during which he or she will receive three paychecks instead of two. The possibility for members of the same household to be on alternate bi-weekly schedules introduces potential noise, and thus I exclude such households from the main analysis.[11] This limitation of the data has further implications for the empirical strategy used in this paper. In particular, the inability to observe which schedule a bi-weekly worker follows leads to classification error in the designation of which months will have three paychecks for that worker. This classification error and the appropriate adjustments made to correct for it are discussed in further detail below in Section 1.4.

After sample restrictions, I am left with a full sample of 24,822 household-month observations for 7,776 households whose heads report being paid either weekly, bi-weekly, or monthly.[12] The main analysis sample used in estimation consists of the 4,316 households

[11] This restriction is later relaxed as a robustness check in Section A.3 and leads to similar results.

[12] In preliminary analysis, households with heads who report being paid semi-monthly were also consid-

from the full sample whose heads of household report being paid bi-weekly. Table 32 presents summary statistics for both the full sample of households and the main analysis sample of only bi-weekly households. All monetary values are in 2010 dollars in this and all subsequent tables. Average monthly consumption expenditures for each aggregated measure of interest are reported both in levels and in changes. Durable and non-durable expenditures each compose about half of total monthly expenditures on average. Monthly changes in expenditure are quite small relative to the mean level of expenditure for each aggregate measure.

The final rows of the table summarize income and assets for the households in each sample. The mean gross paycheck across all households in the full sample is \$1,524 while the mean annual before-tax income is \$53,853. Comparing across the two sample groups, bi-weekly households earn slightly more and hold slightly higher balances in their checking and savings accounts than the full sample of households but do not otherwise appear to significantly differ from the full sample across the key variables of interest. The lower income and assets for the full sample is primarily due to the inclusion of weekly households, who earn significantly lower wages and have significantly lower bank balances than bi-weekly and monthly paid workers. Figure 1 plots the distribution of household income by the three pay frequencies available in the full sample. As one might expect, the bi-weekly distribution is shifted slightly to the left of that of monthly workers and the weekly distribution is shifted slightly to the left of that of bi-weekly workers.

1.4. Empirical Methodology

The timing of bi-weekly pay schedules is such that the calendar months in which there are three paychecks changes from year to year. For instance, a bi-weekly worker paid in the first week of January in 2008 would receive three paychecks in February and August of that year. In 2009, she would instead receive three paychecks in January and July. Appendix Table A.2.1 lists the full schedule of months with three paychecks for each year in the sample. Each calendar month serves as a three paycheck month at least once during the sample period. To estimate responses to third paychecks, I compare changes in consumption expenditures following a given calendar month in years in which there were three paychecks distributed during that month to changes in consumption expenditures in years during

ered. Semi-monthly workers are paid twice a month, typically on the 15th and 30th of a month, and would be the ideal control group since they receive a similar number of paychecks per years as bi-weekly workers but do not experience variation in the number of paychecks received each month. However, the number of households with heads who are semi-monthly is small (less than 4 percent of the households in my sample over the time span considered) and so estimates of spending responses are not feasible. For this reason, I include in the main sample only households whose heads report being paid either weekly, bi-weekly, or monthly. The other excluded categories of pay frequency are quarterly, yearly, and other. These comprise an even smaller sample of households and were thus also excluded.

which there were only two paychecks. Spending responses are thus identified by variation in the timing of months with three paychecks across years.

As described in the previous section, there are two possible schedules by which a bi-weekly worker can be paid. The schedule a given worker is on then determines the set of months with three paychecks for that worker. Let S_1 be the set of three paycheck months under the first schedule and S_2 be the set of three paycheck months under the second schedule. Since I am unable to determine which of these two possible bi-weekly schedules applies to a given individual, I define $S = S_1 \cup S_2$ as the set of all three paycheck months under either schedule and let $\mathbf{1}_{\{t-1 \in S\}}$ be an indicator for whether the previous month is a three paycheck month under either schedule.

The empirical strategy I use follows from the previous literature (Zeldes 1989; Lusardi 1996; Parker 1999; Souleles 1999; Johnson et al. 2006, 2009). Specifically, I estimate variants of the following specification

$$\Delta c_{it} = \beta * \mathbf{1}_{\{t-1 \in S\}} + \theta'_{it}\alpha + \gamma_t + \epsilon_{it}, \tag{1.6}$$

where c_{it} denotes household consumption expenditures and $\mathbf{1}_{\{t-1 \in S\}}$ is the explanatory variable of interest. Since third paychecks are paid out in the last week of a month, I estimate responses in months *following* and not during three paycheck months. The vector θ'_{it} is a set of taste-shifters comprised of the following variables: the age of the head of household, changes in the number of children, and changes in the number of adults. These are included to reflect potential preference driven changes in consumption. Finally, I include a full set of month and year dummies, γ_t, to account for the possibility that changes in household spending reflect seasonal variation in spending, aggregate trends, or supply side movements.[13]

The coefficient of interest, β, is the difference-in-difference (DD) estimate of the expenditure response to third paychecks. The manner in which the indicator, $\mathbf{1}_{\{t-1 \in S\}}$, is defined, however, introduces noise into the estimate of β. In particular, for any given person there are only two months during which they receive three paychecks, but because the CEX does not allow me to observe which schedule a bi-weekly worker is on, the indicator is defined in a way that assumes four. Furthermore, not only am I unable to observe which schedule a particular bi-weekly worker follows, but I am also unable to observe the overall fraction of households in the sample that follow one schedule versus the other. Absent any correction,

[13] Adams et al. (2009) note that auto loan applications and sales spike precisely at the time of tax rebates, and that auto sales companies pay close attention to "tax rebate season." In discussions with a large auto company to assess whether they were likewise aware of third paychecks due to bi-weekly schedules, the company claimed no knowledge of these third paychecks, although they did respond to tax rebate season.

this mis-classification gives inconsistent estimates of β that are biased towards zero. To address this classification error, I formally derive and then estimate a correction factor using the fraction of three paycheck months in the sample that belong to S_1 versus S_2, allowing the true proportion of individuals on either schedule to vary between zero and one. The estimated correction factor ranges from 1.9 to 2.1 and is on average equal to 2.0. Using this estimate, it is then possible to to correct for the classification error by multiplying the estimate of β from Equation 1.6 by a correction factor of approximately 2.0. For simplicity, I present the uncorrected estimates in the results. Additional details on the classification correction are provided in Appendix A.2.

The central identification assumption underlying the difference-in-difference strategy outlined above is that changes in consumption expenditures following a given calendar month in years in which there were three paychecks distributed during that month versus years in which there were only two would have evolved similarly were it not for the third paycheck. This assumption would not hold if, for example, store promotions or sales happened to coincide with months following three paycheck months. To test the validity of the identification assumption, I employ a triple difference (DDD) research design using households whose heads were paid either weekly or monthly as controls. The results from this analysis are qualitatively similar, lending support to the claim that the estimates I find represent the true causal effect of the third paychecks rather than differential trends following months with third paychecks and months without third paychecks.

1.5. Results

1.5.1. Main Results

This section presents the main empirical findings from testing for excess sensitivity. I first present estimates of total household consumption expenditure responses following three paycheck months. I then decompose this response further to determine the subcategories of expenditure that may be driving any observed responses.

1.5.1.1 Aggregate Measures of Spending

Table 2 presents the main results from estimating the baseline specification given by Equation 1.6 for each of the aggregate measures of expenditure. Each column reports estimates from a separate OLS regression where the dependent variable is the dollar change from month $t-1$ to month t in the listed expenditure category. In these and all subsequent specifications, standard errors are clustered at the household level. The first row presents the estimates of the coefficient of interest, β, which measures the average spending response to extra paychecks. The remaining rows provide estimates for the taste-shifters included in the regression.

Column 1 of Table 2 shows that households increase their total expenditures by $262 on average in months following three paycheck months. This increase in total household spending is statistically significant at the 1 percent level and is quite sizable, representing roughly 15.7 percent of the average bi-weekly paycheck and 9.3 percent of average monthly spending by households. Adjusting for classification error, this translates to a $523 or 30.4 increase in total spending following three paycheck months. By comparison, an additional child increases monthly spending by $317, and an additional adult increases monthly spending by $365.

To shed further light on the types of expenditure driving the observed response to extra paychecks, Columns 2 through 5 decompose the total spending responses into spending on durable and non-durable goods. I find that the spending response to extra paychecks is driven almost entirely by changes in spending on durable goods. Household spending on durable goods increases by $257 on average in the months following three paycheck months. By contrast, the estimates for the non-durable measures are both small and statistically insignificant.

Because extra paychecks are a feature of the timing profile of bi-weekly pay and do not lead to a change in lifetime income, we might expect any spending responses to extra paychecks to decrease over time. Figure 2 plots the timing of the durable and strictly non-durable spending responses in months relative to the three paycheck month, denoted by $t = 0$. The consumption paths in this plot were generated by re-estimating Equation 1.6 with the inclusion of leads and lags of the month following a three paycheck month. The dashed lines represent 95 percent confidence intervals surrounding the point estimates.[14]

As shown in the first panel of Figure 2, durable spending spikes in the month immediately following a three paycheck month ($t = 1$) before returning back to prior levels of spending in the following month. This pattern indicates that the effect of extra paychecks does not persist beyond the first month following receipt. Additionally, there appears to be no pre-trend in durable spending leading up to the three paycheck month. This may reflect the existence of constraints on the part of households who are either unable or unwilling to borrow in anticipation of the extra paycheck. Alternatively, households may simply fail to anticipate the existence or timing of these third paychecks. In Section 1.6, I address which of these interpretations better explains the pre-trend behavior. The second panel of Figure 2

[14]Because I am unable to observe which of the two alternate bi-weekly schedules applies to a given individual, I am able to include at most two leads and one lag of the month following a three paycheck month.

shows the analogous timing for spending on strictly non-durable goods. Consistent with the results from Table 2, there is no discernible effect in strictly non-durable spending.[15]

Taken together, these results show that household spending responds significantly to the extra paychecks. Moreover, they indicate that the response is primarily in durables and does not persist beyond the first month following three paycheck months. Additionally, the lack of any discernible pre-trend in spending responses suggests that households do in fact attribute the income from the third paycheck to the months following rather than during three paycheck months. In Appendix A.3, I show that the estimates in Table 2 are robust to alternative specifications. In Appendix A.4, I also look at potential heterogeneity in spending responses by marital status, race, gender, and home ownership. I find no significant differences in estimated spending responses by these various measures. For the remainder of the paper, I therefore focus on the full set of individuals when estimating spending responses.

1.5.1.2 Durable Spending

Because the observed spending response appears primarily in durable goods, Table 3 next decomposes durable spending into several subcategories of expenditure. Specifically, I look at spending responses to extra paychecks for spending on vehicles, furniture, flooring, and major appliances such as refrigerators or stoves. Spending on vehicles refers to the value of the vehicle minus any trade-in allowance. Each column of Table 3 represents a separate regression run with the same covariates but different dependent variables. The first row of Column 1 presents the same estimate of durable spending responses as Column 2 of Table 2 and is listed for reference. These results show that not only is the spending response driven largely by spending on durables, but also this effect is concentrated almost entirely in spending on vehicles. There is a positive but economically insignificant response in spending on flooring. All other estimates are both statistically and economically insignificant.

I further decompose the vehicle spending response by whether the vehicle is a car or motor vehicle and by whether the vehicle is new or used.[16] The first column of Table 4 shows the previously estimated spending response for vehicles and is identical to Column 2 of Table 3. Column 2 shows that purchases of new cars account for nearly 80 percent of the total spending response in vehicles, with expenditures on new cars increasing by $201 on average following a three paycheck month. Columns 6 and 7 report estimates from a linear

[15] The higher precision in the estimated spending response for strictly non-durable goods than for durable goods reflects the higher variance in durable goods spending, which is often characterized by large and infrequent purchases.

[16] Cars include automobiles, trucks, and vans. Motor vehicles include motorcycles, motor scooters, and mopeds.

probability model and show that not only does spending on new cars increase following three paycheck months, but the probability of making a new car purchase increases significantly by 0.4 percentage points.[17]

While the preceding evidence suggests that spending on vehicles increases in response to the extra paychecks, it refers only to the full purchase price of the vehicle and does not consider actual out-of-pocket spending. Table 5 provides statistics on vehicle expenditures as well as three categories related to vehicle financing: out-of-pocket expenditures, debt financing, and down payments. Out-of-pocket expenditure refers to either the down payment amount if the vehicle is financed or the full purchase price (minus the trade-in allowance) if the vehicle was not financed. Statistics are presented both unconditional and conditional on ever purchasing a vehicle during the household's survey period.

Conditional on ever purchasing a vehicle, the mean expenditure on new and used vehicles is \$2723, while the unconditional mean is \$299. This difference reflects the fact that few households report a vehicle purchase during their survey period. However, even these statistics fail to provide a complete picture of vehicle spending and financing since households who ever purchase a vehicle do not purchase a vehicle every month. For comparison, the average value of a vehicle purchased is approximately \$12,497. The average out-of-pocket expenditure is \$3569 while the average amount of financing for vehicle purchases is \$8928. Conditional on making a down payment, households put down 21.2 percent of the purchase price of the vehicle on average, which corresponds to a payment of \$3531.

Table 6 presents estimates of the spending response on vehicles, taking into account vehicle financing. Spending on cars increases by \$247 on average following months with three paychecks. The spending response for debt financing is nearly twice the size of the spending response for out-of-pocket expenditures, with roughly 67 percent of the overall vehicle spending response from financing and 33 percent from out-of-pocket spending. The estimated response in down payments is smaller and imprecisely estimated.

The spending response in vehicles is driven by a small number of households. In particular, only 407 households in the sample report purchasing cars during their interview period. Because only a small number of households are responsible for vehicle spending, the magnitudes of the estimates are necessarily small. They measure the average increase in spending for all households, most of whom did not purchase a car during their survey period and therefore have little to no reported change in vehicle spending. Table 7 presents conditional estimates analogous to those in Table 6. Conditional on ever purchasing a car during the survey period, spending on cars increases by \$2205 on average following a three paycheck

[17] For the remainder of the paper, I refer to the terms vehicles and cars interchangeably.

month.

1.5.2. Placebo Tests

I interpret these findings as evidence that, contrary to the predictions of standard theory, individuals respond to the variation in their monthly income induced by their bi-weekly pay schedules. This interpretation relies on the identification assumption that changes in consumption expenditures following a given calendar month in years in which there were three paychecks distributed during that month versus years in which there were only two would have evolved similarly were it not for the third paycheck. The implication of the assumption is that there should be no corresponding response for workers paid at frequencies that do not induce such variation. Specifically, we would expect no response for workers paid monthly and a smaller response, if any, for workers paid weekly.[18] To evaluate the validity of the assumption, I employ a triple difference strategy to compare the estimated spending responses in the full sample using weekly and monthly households as controls. The results from this specification are presented in Table 8. In contrast to bi-weekly workers, weekly and monthly households do not respond in months following three paycheck months. The fact that the estimated spending responses for weekly and monthly households are not statistically different from zero suggests spending responses following a given month in years in which three paychecks were distributed during that month versus years in which there were only two would in fact have evolved similarly were there no third paycheck.

As a final check on the findings, I run a series of placebo tests to evaluate the extent to which the baseline estimate of β (Column 1 of Table 2) may be spurious. I first generate 1000 placebo paycheck schedules, each of which designates a subset of the months spanning my sample period from 1990-2010 as three paycheck months. To do this, I randomly select four different months for each year in my sample period. In doing this, I also impose the additional restriction that none of the four months that are randomly chosen in each representative year can be consecutive since the true set of months that have three paychecks never follow consecutively. This procedure is repeated 1000 times to create the set of 1000 placebo schedules. I then re-estimate the regression from Equation 1.6 to obtain estimates of β for each of these placebo schedules. Figure 3 plots the distribution of these placebo estimates as well as the true estimate, which is marked by the vertical dashed line. As the figure makes clear, the true estimate is to the far right of the distribution, which suggests that the probability of finding an effect as large as mine simply by chance is extremely

[18] The time profile of weekly workers is such that they receive four paychecks per month with the exception of four months out of the year during which they receive five. However, the extra paycheck is both smaller and occurs more frequently (four times a year instead of two) for weekly workers relative to bi-weekly workers, so we might expect consumption responses, if any, to be small.

small.

1.5.3. Time Adjustments in the CEX

Households are interviewed in the CEX at a quarterly frequency. In these interviews, households report expenditures for the three months prior to the interview month (the reference period), which allows for conversion of the data from a quarterly frequency to a monthly frequency. However, households on occasion report quarterly or annual expenditures over a reference period rather than their monthly expenditures. As a result, for a subset of expenditures, the Bureau of Labor Statistics uses pre-determined time adjustment routines when mapping expenditures to the associated month of purchase.[19] The conditions for whether a given expenditure category is time adjusted depends on both the type of expenditure and the information source for the expenditure. Three time adjustment methods in particular are worth noting as they may affect whether I observe household spending responses. Depending on the expenditure category, reported expenditures can either be divided by three and then assigned for each of the three months in the reference period, divided by twelve and then assigned for each of the three months in the reference period, or allocated to a random month in the reference period.

The effect of these time adjustment routines on my estimates is ambiguous. The first two methods described may bias towards finding evidence of consumption smoothing whereas the last method may bias towards finding evidence against smoothing. Given that the expenditures are allocated to a random month for the third time adjustment routine, there is no reason to expect that the assigned month is correlated with the indicator variable for following a three paycheck month. Nonetheless, to address potential concerns stemming from the use of these time adjustment routines, I re-estimate responses following three paycheck months for the original bi-weekly sample but excluding expenditures which undergo one of the three time adjustment methods.[20] Table 9 reports the results of of this re-estimation for the categories of spending. I find evidence of spending responses even after excluding expenditure categories that undergo time adjustment. These results are of similar magnitude and statistical significance to the original estimates in Table 2, which suggests that the time adjustment routines do not have a large effect on my analysis.

[19] See Hai et al. (2013) for additional description of the time adjustment routines used by the BLS.

[20] I use the parsing file provided by the Bureau of Labor Statistics which provides a list of the universal classification codes (UCCs) for expenditures which undergo time adjustment in addition to the type of time adjustment routine that is applied.

1.6. Potential Explanations

Taken together, the results presented above provide substantial evidence that household spending responds following three paycheck months. In this section, I explore several possible explanations that may account for these findings. First, households may face binding liquidity constraints and thus respond to the extra paychecks that provide additional liquidity upon arrival. Second, households may adopt budgeting heuristics where they deviate from the assumption of fully rational expectations and instead hold incorrect beliefs about future income. Third, individuals may be sophisticated with time-inconsistent preferences, thus choosing to invest their third paychecks in an illiquid asset as a commitment device. Finally, individuals may hold a system of mental accounts by which they treat third paychecks as a bonus or windfall gain distinct from their typical income of two paychecks each month.

1.6.1. Liquidity Constraints

Liquidity constraints are perhaps the most often cited explanation for observed household spending responses to predictable income changes or predictable income receipt. The intuition behind why liquidity constraints might matter is straightforward. Consider the same individual from the basic environment in Section 1.2.1, but now assume that she faces a borrowing constraint of the form $X_t - c_t \geq 0$ so that current period consumption can never exceed current cash-on-hand. The individual's modified Euler equation is then

$$u_c(c_t) = \max\{u_c(X_t), \delta R u_c(c_{t+1})\}.$$

Let period $t+1$ be an atypical income period (e.g. a month following three paychecks) and suppose that the individual's borrowing constraint is binding in the previous period. In other words, suppose that

$$c_t = X_t < (1 - (\delta R^{1-\rho})^{1/\rho})W_t,$$

where $(1 - (\delta R^{1-\rho})^{1/\rho})W_t$ is her optimal level of consumption under no constraints. The additional income bY in period $t+1$ relaxes the borrowing constraint, which allows the individual to increase her spending in response. Furthermore, she spends more in the atypical income period than she otherwise would have under the same model with no borrowing constraints because of the forced savings from the presence of the binding constraint. Given that the constraint was binding in period t, the forced savings is equal to the difference between her optimal level of consumption and her cash-on-hand in period t.

To see whether this intuition holds in the data, I evaluate the role of liquidity in the estimated spending responses to extra paychecks using two complementary strategies. I first estimate the spending response separately for constrained and unconstrained households using a series of proxies for the presence of liquidity constraints.[21] For each proxy, I classify households into two groups using the median value of the measure. I also implement a second empirical strategy that uses observed behavior to classify households as constrained or unconstrained. Specifically, I infer whether a household is constrained or unconstrained by the extent to which the household's choices are consistent with binding liquidity constraints. I consider two types of observed behavior: consumption volatility and vehicle loan characteristics. The first of these distinguishes households by their ability to smooth strictly nondurable consumption in months excluding those following three paycheck months (months with typical income). The use of vehicle loan characteristics is motivated by Attanasio et al. (2008), who find that constrained and unconstrained households differ in their responsiveness to changes in certain vehicle loan characteristics, specifically loan maturity and interest rates.

1.6.1.1 Tests of Liquidity Constraints using Proxy Measures

I first test for the presence of binding liquidity constraints using traditional methods of asset-based sample splitting (Zeldes 1989; Runkle 1991). The proxy I consider is the total reported balance in household checking and savings accounts, which I refer to as "liquid assets." While this is perhaps the most relevant way to proxy for liquidity, limitations of the CEX make this an especially difficult proxy to measure. Information on checking and savings account balances is collected only once over the survey period, during a household's final interview. Furthermore, few households report their savings and checking account balances, which in combination with the fact that the spending response is primarily driven by a small number of households, means that any tests of the role of liquidity constraints using this measure are significantly underpowered.[22] Nonetheless, since liquid assets are perhaps the most appropriate measure of liquidity available, I include estimates of the spending response for constrained and unconstrained households using liquid assets as a proxy. Table 10 shows these estimates for each of the aggregate categories of expenditure.[23]

[21] In general, it is not possible to actually determine whether a particular household faces liquidity constraints given the data available, and thus the presence of binding constraints must often be proxied with other measures. An exception to this is Agarwal et al. (2007) who use data from credit card accounts and find that spending rose most for individuals who were most likely to be liquidity constrained, which is defined using credit limits and credit card utilization rates. Their findings suggest that the presence of liquidity constraints may be important.

[22] If we restrict the sample to households with bi-weekly heads who report savings and checking account balances, the sample size drops to only 1479 households.

[23] In this and all remaining tables testing for liquidity constraints, I present estimates using the dollar change in monthly income ΔY_{it} as the explanatory variable of interest rather than the indicator $\mathbf{1}_{\{t-1 \in S\}}$. This is to account for the fact that unconstrained and constrained households likely have different sized third

The first row provides estimates for the constrained households (those with liquid assets below the sample median) while the second row provides the estimated difference in spending response for unconstrained households relative to constrained. Contrary to the predictions of a model with borrowing constraints, I find that households with higher levels of liquid assets exhibit *larger* spending responses than those with lower levels of liquid assets, though the difference is not statistically significant.

Because reports of liquid assets are underpopulated, I also proxy for constraints using household (before-tax) income. While income may not reflect liquidity in the same way in which liquid assets do, we might expect the two measures to be highly correlated. Results are presented in Table 11 and are similar to those using liquid asset holdings as a proxy: unconstrained households exhibit larger spending responses than constrained households, although the difference is not statistically significant.

I next classify households as constrained and unconstrained using the age of the head of household as a proxy measure. The relationship between age and liquidity is somewhat ambiguous. Labor income tends to be more concentrated later in life, implying less liquidity for younger households. Additionally, lenders may be reluctant to provide loans to younger households for whom they have little credit history to rely on. At the same time, younger households face a steep consumption profile, suggesting a smaller likelihood of being constrained. Jappelli (1990) finds that younger households are more likely to be liquidity constrained and suggests that the first two effects must dominate the third. Table 12 shows that while younger households exhibit stronger spending responses to extra paychecks, the difference is again not statistically significant.

The previous three proxies for liquidity – liquid assets, income, and age – are the standard measures of liquidity used in the literature. As a final proxy, I create a new measure using a household's committed consumption as a fraction of monthly wage income. Individuals who face large recurring expenditures that are difficult to adjust and that constitute a large share of their available income may have less disposable income and thus may exhibit greater sensitivity to cash-on-hand. To construct the measure of committed consumption, I first aggregate monthly expenditures on mortgage payments, rent, car loans, and utilities for each household and then divide the total level by the monthly wage income for that household.[24] While this measure likely understates the true level of committed consumption, it is composed of expenditures we might reasonably think would be difficult to adjust. As Table 13 shows, while I find larger durable spending responses for households with higher

paychecks.

[24]Monthly wage income here is calculated as the wages for the head of household in a two paycheck month and is thus equal to two times the reported gross pay.

committed consumption, the difference is not statistically significant.

1.6.1.2 Tests of Liquidity Constraints Using Observed Behavior

Liquidity constrained households have a limited ability to smooth consumption intertemporally. One potential implication of this limited ability is that constrained households may exhibit more volatile consumption patterns than unconstrained households. With this in mind, I construct a measure of consumption volatility for each household using the standard deviation of strictly non-durable spending in months excluding those following three paycheck months. Households are classified as liquidity constrained or unconstrained by comparing their consumption volatility to the median value of the measure. I then estimate spending responses to extra paychecks allowing the response to differ by whether the household is constrained or unconstrained. An important caveat to this approach is that consumption volatility defined in this manner is endogenous to the measured effect. As such, tests using consumption volatility to distinguish between constrained and unconstrained households are simply suggestive and are not meant to be taken as strong evidence for or against the presence of liquidity constraints. Table 14 presents results using the consumption volatility meaures. I find that spending responses to extra paychecks by unconstrained households are larger than estimated responses by constrained households, but the difference is not statistically significant.

The presence of liquidity constraints also has implications for the types of loans a household may choose. Since the observed spending responses to extra paychecks are primarily due to vehicle purchases, I focus on the implications of constraints for vehicle loans specifically. Furthermore, I consider loan characteristics whose values we might expect to differ between constrained and unconstrained households: maturity, monthly payment size, loan-to-value ratio, and down payment. To the extent that these extra paychecks provide additional liquidity to constrained households, we might expect loans associated with vehicles purchased following three paycheck months (atypical income months) to have characteristics more consistent with the presence of constraints relative to loans associated with vehicles purchased in other months (typical income months). I thus compare the loan characteristics of vehicles purchased in months following three paycheck months with the loan characteristics of vehicles purchased in other months to gauge whether constrained households may disproportionately be responding to extra paychecks.

Table 15 reports the results of this comparison using the subset of vehicles for which I observe loan characteristics. The table shows that the loan characteristics of interest do not differ significantly between vehicles purchased following three paycheck months and those purchased in other months. While this comparison does not necessarily imply that third paychecks do not help relieve binding constraints for households, it suggests that vehicles purchased following three paycheck months do not appear to be disproportionately purchased by constrained households.

Overall, these tests for liquidity constraints at best weakly suggest the relevance of liquidity constraints since any differences in response were not significant and several go in the opposite of what might be expected.[25] Moreover, if households were liquidity constrained, we would expect there to be a response in non-durables spending in addition to any response in durable spending. But as Table 2 shows, there is no corresponding response in non-durable spending following months with three paychecks. Taken together, this suggests that liquidity constraints alone cannot explain the findings in this paper.

Several other papers in the literature have also found that spending responses to predictable income receipts are not correlated with conventional measures of liquidity constraints (Parker 1999; Souleles 1999; Shapiro and Slemrod 1995; Stephens, Jr. 2008; Stephens, Jr. and Unayama 2011). The literature is largely undecided as to the extent to which such constraints may or may not play a role. At present, it remains unclear whether the lack of consensus stems from the acknowledged difficulties in measuring liquidity constraints or is due to heterogeneity in the contexts where liquidity constraints play a role.

1.6.2. The t-Budgeting Heuristic

An alternative explanation for the empirical findings is that bi-weekly workers adopt budgeting heuristics in response to the misalignment between the timing of income and the timing of consumption. The variation in income generated by this misalignment introduces additional complexity to the household financial decision problem. Empirical evidence has shown that individuals often have difficulty fully optimizing when facing such complexity.[26] One way in which they may respond is by adopting rules-of-thumb or *heuristics* that ease the cognitive burden associated with complex decision problems (Gilovich et al. 2002; Bernatzi and Thaler 2007; Lacetera et al. 2012). In this section, I propose a specific heuristic in which households construct monthly budgets by naively extrapolating their current income into the future. This heuristic is motivated in part by the fact that many sources of informal financial advice commonly suggest that households form monthly budgets based on their current cash flow.

[25] As Jappelli et al. (1998) note, proxy measures for liquidity are prone to mis-classification error, which may downward bias any estimated difference between constrained and unconstrained groups.

[26] For example, Choi et al. (2011) find that individuals sub-optimally invest in their 401(k) by contributing at a rate below the threshold matched by their employer, even when doing so is dominated by contributing at the match threshold. Abaluck and Gruber (2011) show that individuals choose Medicare Part D prescription drug plans that are strictly worse than other available plans.

With this motivation in mind, I develop a simple model in order to formalize the budgeting heuristic and investigate its implications for the path of consumption. Consider once again the individual from the standard consumption and savings framework from Section 1.2. Let $y^E_{t,t+i}$ denote the income the individual in period t believes will be received in period $t+i$ for $i \geq 1$. I parameterize beliefs according to

$$y^E_{t,t+i} = \alpha \cdot y_t + (1-\alpha) \cdot y_{t+i}, \tag{1.7}$$

where $0 \leq \alpha \leq 1$. This formulation of beliefs nests the rational model as individuals with $\alpha = 0$ hold fully rational beliefs about future income. More generally, the parameter α can be thought of as a measure of the extent to which individuals extrapolate their current income into the future. In other words, individuals with $\alpha > 0$ mistakenly believe that any deviation in income from their current period t income is a weighted average of a permanent and a transitory income shock, where α is the weight the individual places on the deviation $y_t - y_{t+i}$ being a permanent income shock.

This parameterization of beliefs is similar in spirit to the model of projection bias introduced in Loewenstein et al. (2003). With projection bias, individuals extrapolate their current tastes, rather than income, into the future. This parameterization can also help to explain the presence of "rule-of-thumb"consumers who consume out of current disposable income as discussed in Hall and Mishkin (1982) and Campbell and Mankiw (1989).

I define an individual to be following a *t-budgeting heuristic* if she makes her consumption decisions at a frequency t and has $\alpha > 0$. In the context of this paper, an individual thus adopts a monthly budgeting heuristic if she makes her consumption decisions at a monthly frequency and extrapolates her current income into the future.

1.6.2.1 The t-Budgeting Heuristic with Misalignment Between the Timing of Consumption and the Timing of Income

Section 1.2.1 demonstrates that, in the absence of a t-budgeting heuristic, misalignment between the timing of consumption and the timing of income is irrelevant in a basic environment with no borrowing constraints and no income uncertainty. However, when an individual does follow a t-budgeting heuristic ($\alpha > 0$), misalignment between the timing of consumption and the timing of income becomes quite important.

For an individual who follows a t-budgeting heuristic, expected remaining lifetime wealth at time t can be written as,

$$W_t^E = X_t + \sum_{i=1}^{\infty} \frac{y_{t,t+i}^E}{(1+r)^i}, \tag{1.8}$$

where $y_{t,t+i}^E$ is given by Equation 1.7. For this individual, the relationship between expected wealth and true wealth depends on whether the current period is a typical income period or an atypical income period.

Definition 1. An individual has *overly optimistic* beliefs about future income if her expected lifetime wealth is strictly greater than her true lifetime wealth ($W_t^E > W_t$) and has *overly pessimistic* beliefs about future income if her expected lifetime wealth is strictly less than her true lifetime wealth ($W_t^E < W_t$).

It is easy to show that when the timing of consumption and the timing of income are misaligned, individuals who follow a t-budgeting heuristic have overly optimistic beliefs in periods with atypical income and overly pessimistic beliefs in periods with typical income.[27] The intuition underlying this result is straightforward. An individual who extrapolates current income into the future places too much weight on her current income when forming her expectations about future income. Since her current period income is less than her average per-period income in periods with typical income, the individual will have expectations about her wealth that are too low. Likewise, in periods with atypical income, her current period income is higher than her average per-period income and thus her expectations about her wealth are too high.

Since consumption is proportional to expected wealth, over-pessimism and over-optimism in beliefs imply that consumption will be either too low in typical income periods or too high in atypical income periods relative to the optimal consumption level under the rational benchmark. The following proposition characterizes the marginal propensity to consume out of additional income bY in periods with atypical income and how the measure relates to α.

Proposition 1. *In the presence of a t-budgeting heuristic and misalignment between the timing of consumption and the timing of income, the marginal propensity to consume out of additional income bY in atypical income periods is $(1-(\delta R^{1-\rho})^{1/\rho})\frac{R}{R-1}\alpha$ and is increasing in α.*

Proof. See Appendix A.5. □

Proposition 1 follows directly from the intuition outlined previously for over-optimism in periods with atypical income. As α increases, the extent to which the individual extrapolates

[27] This statement is formalized in Lemmas 1 and 2 in Appendix A.5.

current income into the future, and hence the erroneously expected "permanence" of the additional income bY, increases. Thus α determines not only whether an individual is overly optimistic in periods with atypical income but also *the degree to which* she is so. As a result, the marginal propensity to consume out of the additional income bY in an atypical income period is increasing in α. Proposition 1 thus predicts excess sensitivity of consumption in months following three paycheck months (i.e. in periods with atypical income) for bi-weekly workers who follow a monthly budgeting heuristic.

Finally, it is important to recognize that the adoption of a t-budgeting heuristic is irrelevant for consumption patterns when there is no misalignment between the timing of consumption and the timing of income. With no misalignment, income is constant each t-period so that $y_t = y_{t+i}$. In this case, the individual's expectations about her future income are not only independent of α, but are also equivalent to her expectations if she were rational: $y^E_{t,t+i} = y_{t+i}$.

1.6.2.2 The t-Budgeting Hueristic with Borrowing Constraints

While the results in Section 1.6.1 suggest that liquidity constraints alone cannot explain the findings in this paper, this does not mean they do not play any role. It is therefore important to consider how the presence of liquidity constraints might interact with the adoption of a t-budgeting heuristic. Incorporating the simple borrowing constraint, $X_t - c_t \geq 0$, from before gives the following proposition.

Proposition 2. *For all α, the presence of borrowing constraints of the form $X_t - c_t \geq 0$ leads to weakly larger increases in consumption in periods with atypical income than in the absence of borrowing constraints.*

Proof. See Appendix A.5 □

As previously mentioned, the presence of binding borrowing constraints results in forced savings which leads to greater consumption in atypical income periods when the constraint is relaxed by the additional income bY. This amplifies any effects on consumption from individuals extrapolating current income into the future, and thus consumption in atypical income periods is higher than would otherwise by implied by in the absence of borrowing constraints.

1.6.2.3 The t-Budgeting Heuristic with Committed Consumption

The consumption model thus far has assumed a single, composite consumption good, c. An important first generalization is to allow for both durable and non-durable goods. As Table 32 shows, durable goods constitute nearly half of total expenditure. Furthermore, features of durable goods often make adjustment costly in ways that are distinct from the adjustment of non-durable goods. Following a t-budgeting heuristic may therefore have very different implications for the consumption of non-durable versus durable goods. With this in mind, I follow the approach taken by Chetty and Szeidl (2007) and extend the current model to allow for two types of consumption goods: adjustable (c) and committed (d). Committed consumption is differentiated from adjustable consumption through the presence of a fixed cost to adjustment, $\psi(d) = k \cdot d > 0$. In order for the individual to adjust her committed consumption, she must pay this fixed cost. Expected remaining lifetime utility is now defined as

$$U_t(\{c_\tau, d_\tau\}_{\tau=t}^{\infty}) = E_t\left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} u(c_\tau, d_\tau)\right], \tag{1.9}$$

where $u(c,d)$ satisfies the conditions: (i) $\lim_{c\to\infty} u_c(c,d) = \lim_{d\to\infty} u_d(c,d) = 0$ and (ii) $\lim_{c\to 0} u(c,d) = \inf_{c',d'} u(c',d')$ for all d. Within each period, the individual has Cobb-Douglas preferences over adjustable and committed consumption: $u(c,d) = \frac{(c^{1-\gamma}d^{\gamma})^{1-\rho}}{1-\rho}$ with $0 < \gamma < 1$. I assume $\rho < 1$ to ensure that the marginal utility of adjustable consumption is always increasing in committed consumption ($u_{c,d}(c,d) > 0$).

Individuals enter period 0 with some level of assets X_0 which is distributed according to the density $f(X_0)$. Since we are interested in the implications of following a t-budgeting heuristic, assume that there is misalignment between the timing of consumption and the timing of income, with period t income given by Equation 1.2. The individual's dynamic budget constraint is given by

$$X_{t+1} = R(X_t - c_t - d_t - kd_{t-1} \cdot \mathbb{1}\{d_{t-1} \neq d_t\}) \quad \text{for all } t, t+1.$$

For simplicity, I assume that that there is no heterogeneity in income y_t. Then, individuals can be uniquely identified by their initial level of assets, their coefficient of relative risk aversion, and the extent to which they extrapolate current income into the future (X_0, ρ, α).

Conditional on entering period 0 with assets X_0 and income $y_0 = Y$, the individual chooses her optimal consumption path, $\{c_\tau, d_\tau\}_{\tau=0}^{\infty}$, in order to maximize her expected lifetime utility. Denote the individual's optimal consumption in period t by $\{c_t, d_t\}$. Since the main goal is to understand the choice of adjustable and committed consumption in periods with atypical income, I focus on two periods: period 0 when individuals receive typical income Y and chooses their initial level of committed consumption and period 1 when individuals first receives additional income bY.

Given the preferences specified above, Chetty and Szeidl (2007) show that for each initial level of committed consumption d_{t-1} in period $t-1$, there exists $s_t < S_t$ such that the optimal policy in period t is to (i) maintain the current level of committed consumption when $W_t^E \in (s_t, S_t)$ and (ii) adjust committed consumption when $W_t^E \notin (s_t, S_t)$. In other words, at time t, individuals with expected remaining lifetime wealth W_t^E that lies outside the band (s_t, S_t) will discretely adjust their committed consumption. The following proposition characterizes how the probability of adjusting committed (durable) consumption depends on α.

Proposition 3. *Let wealth entering period 0 be distributed according to the CDF $G(W_0^E)$ and density $g(W_0^E)$, where the density $g(\cdot)$ depends on the joint distribution of (X_0, α). For a given initial level of committed consumption d_0, the probability that an individual discretely adjusts her committed consumption in period 1 ($d_1 \neq d_0$) is increasing in α.*

Proof. See Appendix A.5. □

To understand the basic intuition for this result, recall that individuals who follow a t-period budgeting heuristic have over-optimistic beliefs about future income in periods with atypical income. Then for some of these individuals (those with large α), the perceived positive shock to permanent income in periods with atypical income is sufficiently large that it is preferable to pay the adjustment cost and modify committed (durable) consumption rather than maintain the current level of committed consumption and allocate the additional wealth to adjustable consumption.

For adjustable (non-durable) goods, the predicted effect is ambiguous. For those individuals whose perceived positive shock to permanent income is sufficiently large, adjustable consumption may decrease as individuals adjust upwards their levels of committed consumption. However, this decrease will be offset by increases in adjustable consumption for those individuals whose perceived positive shock was *not* sufficiently large to lead to discrete adjustment of committed consumption.

Taken together, Proposition 3 makes the following predictions about durable and non-durable consumption: assuming bi-weekly workers follow a monthly budgeting heuristic, durable consumption will weakly increase on average in months following three paycheck months, while the response for non-durable consumption is ambiguous. While these predictions of the monthly budgeting heuristic model are consistent with the empirical findings in the paper, it is important to note that the ability to empirically distinguish this explanation from the other proposed explanations is limited by the information available in the CEX. However, the budgeting heuristic model does give rise to predictions that may be testable with other data. For example, one such prediction is that, all else equal, we would expect an individual paid bi-weekly to have lower levels of committed consumption expenditures (e.g. rent, mortgage payment, etc.) than an individual paid monthly. This prediction is a direct implication of the incorrect beliefs about future income: a bi-weekly worker will typically be overly-pessimistic about her future income relative to a monthly worker. Testing this prediction relies crucially on the ability to observe the pay frequency at the time of commitment (e.g. signing a rental lease), which the CEX does not allow for. However, the increasing availability of high-frequency data holds promise for future research in this area.

1.6.3. Time Inconsistency with Sophistication

The presence of sophisticated individuals with (β, δ)-preferences could also potentially explain the empirical findings (Strotz 1956; Phelps and Pollak 1968; Laibson 1997; O'Donoghue and Rabin 1999). These individuals are quasi-hyperbolic discounters with preferences given by

$$U_t(\{c_s\}_{s=t}^{\infty}) = E_t \left[u(c_t) + \beta \sum_{s=t+1}^{\infty} \delta^{s-t} u(c_s) \right], \tag{1.10}$$

where δ represents the standard exponential discount factor and β captures the relative trade-off between the present and the future. When $\beta < 1$, preferences are not only time-inconsistent but also *present-biased*. That is, individuals overweight the utility of present consumption relative to future periods of consumption, preferring immediate gratification to future gratification.

This present-bias can lead to future self-control problems. The implications of such problems for individual behavior depend on the extent to which individuals are aware of their self-control problems (O'Donoghue and Rabin 1999). In particular, individuals who are fully aware of their self-control problems (sophisticates) may in fact demand various forms of commitment in anticipation of these problems. For bi-weekly workers, sophisticated individuals with present-biased preferences may choose to invest their third paychecks in durable goods because the illiquid nature of these goods act as a commitment mechanism against overconsumption (Laibson 1997). Such behavior could explain the observed spending responses in durables following three paycheck months. However, the estimated durable spending responses were primarily due to *debt-financed* vehicles, the purchase of which commit a household to a stream of future installment payments which, with rational income expectations, they may not be able to afford. Moreover, time inconsistency with sophistication fails to explain why there is no effect of third paychecks on non-durable consumption.

1.6.4. Mental Accounting

An alternative explanation is that individuals engage in mental accounting behavior. Under a mental accounting framework, individuals no longer treat money as fungible and instead use a system of mental accounts to categorize and evaluate their income (Thaler and Shefrin 1981; Shefrin and Thaler 1988; Thaler 1990, 1999, 2008). While third paychecks are not designated any differently from non-third paychecks, bi-weekly workers may nonetheless choose to treat their typical income of two paychecks per month as their "regular" income and to view third paychecks as a bonus or windfall gain distinct from regular income. Mental accounting behavior in this manner can lead to different marginal propensities to consume out of different mental accounts (Shefrin and Thaler 1988). In the context of bi-weekly workers, mental accounting can predict responses to third paychecks even in the absence of any change in total lifetime income or binding liquidity constraints and with correct beliefs about future income. However, given rational expectations about future income, it is once again surprising that the spending responses are driven primarily by the purchase of debt-financed vehicles that commit a household to a stream of future payments.

Mental accounting not only captures behavior in which individuals label the source of income but also behavior in which individuals categorize the *use* of the income. Rather than simply treating third paychecks as a bonus or windfall gain, a related possibility is that bi-weekly workers earmark third paychecks specifically for other uses, such as the purchase of large durables. Such behavior could potentially explain the observed spending responses in debt-financed vehicles following three paycheck months.

1.7. Conclusion

For many households, the frequency for which they make their consumption decisions often differs from the frequency at which their income arrives. This misalignment can result in predictable variation in the amount of income received per consumption decision period. Under standard assumptions, such variation should not matter since its timing is fully anticipated by households. I test this prediction by leveraging variation in monthly income arising from the timing of of bi-weekly pay schedules. Bi-weekly workers typically receive two paychecks per month with the exception of two months out of the year, during which they receive three. I find that households with bi-weekly heads increase their spending significantly in response to third paychecks: total household spending increases by $262 on average following three paycheck months. Furthermore, I find that this spending response is due entirely to spending on durables, and specifically new car purchases, with no corresponding response in non-durables.

Contrary to standard consumption theory, the empirical findings suggest that differences in the timing of consumption and the timing of income can in fact have large effects on household consumption patterns. I present a number of explanations for these findings. Specifically, I consider the standard extension of liquidity constraints as well as three behavioral explanations: time inconsistency with sophistication, mental accounting, and a new model of budgeting heuristics in which individuals naively extrapolate their current income into the future. While this paper attempts to address some of the alternative explanations for the empirical findings, further research is warranted to fully distinguish between the alternative explanations and understand the mechanisms driving the response. For instance, while the empirical evidence suggests that binding liquidity constraints cannot fully explain the results, the ability to determine the extent to which individuals face such constraints is limited by the available data. As high-frequency data linking individual income, debt, and asset balances becomes increasingly available, future research will allow for a fuller picture of how individual behavior responds to the predictable variation induced by misalignment.

A natural extension to the results presented in this paper is to explore whether households respond to predictable variation in (committed) consumption expenditures arising from misalignment between the timing of income and timing of consumption. Rather than considering variation characterized by infrequent but recurring periods with increased income (e.g. three paychecks instead of two), this alternative variation is characterized by infrequent but recurring periods with increased committed expenditure and hence decreased (disposable) income. Studying anticipated *declines* in income has the advantage of avoiding the common empirical difficulty in testing for the presence of binding liquidity constraints. Specifically, any evidence of excess sensitivity cannot be attributed to the presence of liquidity constraints since households should save, not borrow, in anticipation of declines in disposable income. I plan to explore this further in future research using individuals who are paid monthly but make bi-weekly mortgage payments. These individuals typically make two mortgage payments each month, with the exception of two months out of the year when they make three. As a result, these individuals have less disposable income following three payment months. Studying responses to this source of anticipated would complement the findings in this paper.

Fig. 1.—Total Income Distribution by Pay Frequency

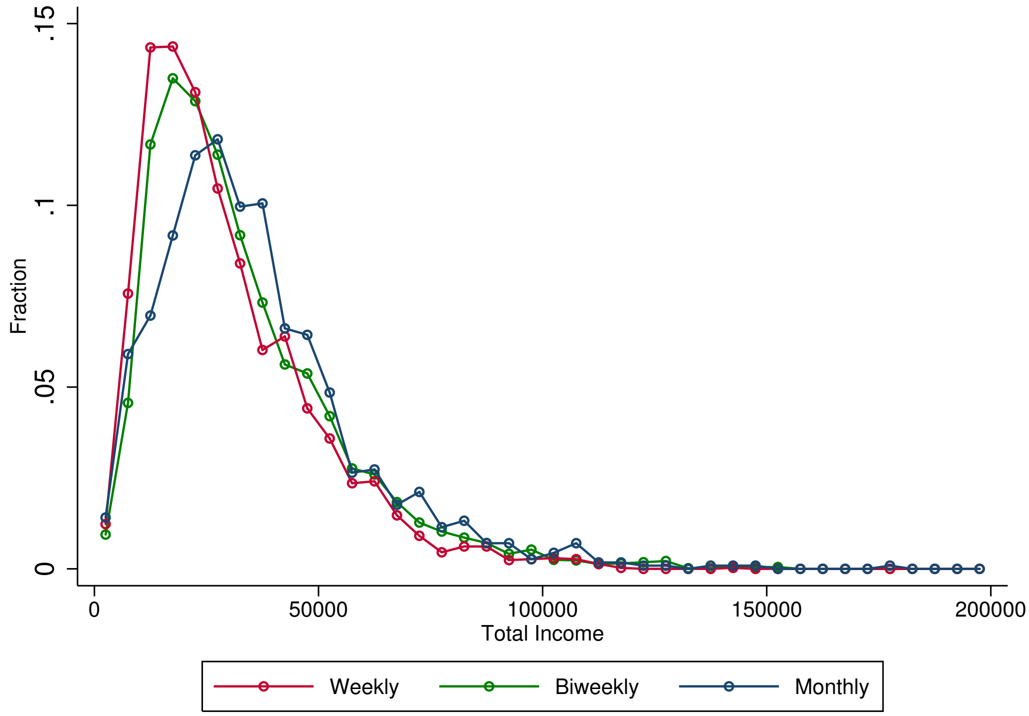

Note.— This figure plots the distributions of total before-tax income for households in the full sample by three different pay frequencies for the head of household - weekly, biweekly, and monthly. The figure was created using the interview portion of the Consumer Expenditure Survey from 1990-2010.

Fig. 2.—Timing of Durable and Strictly Non-Durable Spending

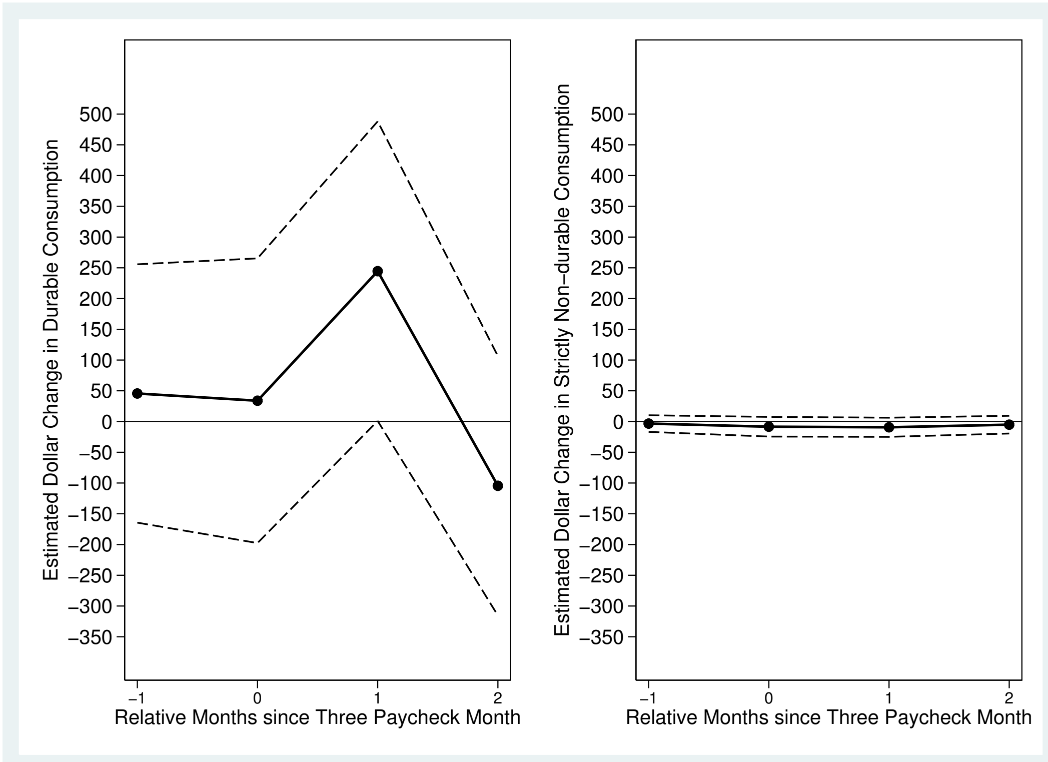

Note.— Each figure plots distributed lag estimates from a regression of the dollar change in spending on indicators for the month of expenditure relative to three paycheck months (t=0). The estimates control for changes in taste and seasonal variation. The dashed lines are 95 percent confidence intervals.

Fig. 3.—Distribution of Placebo Estimates of Spending Effect

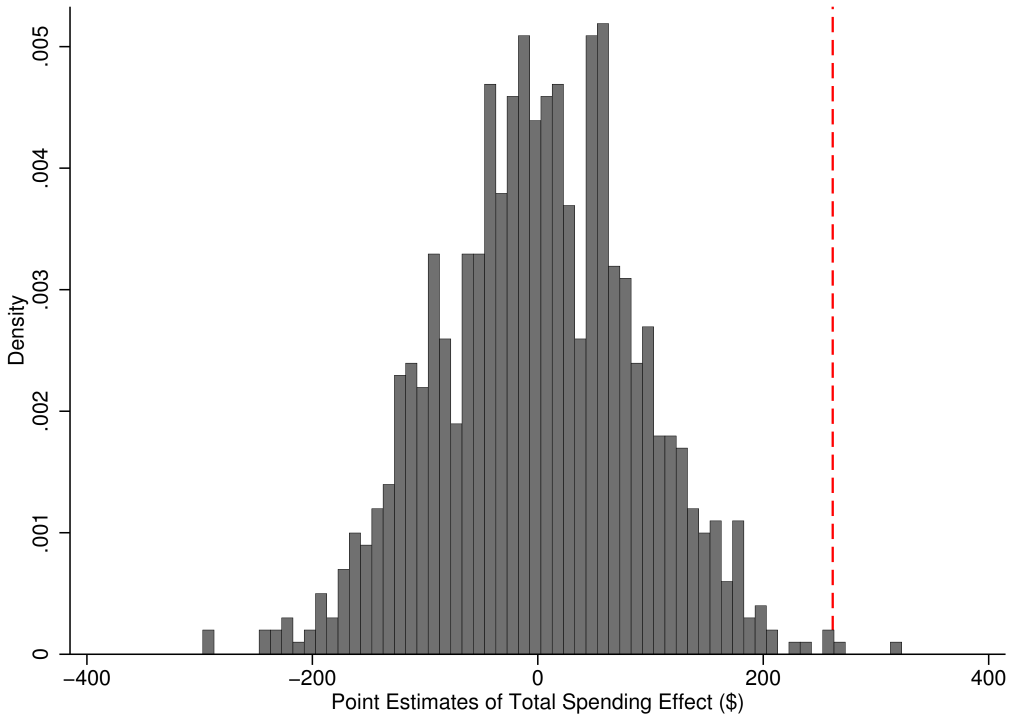

Note.— This figure shows the distribution of 1000 placebo estimates of the total spending effect following three paycheck months. These effects were estimated using randomly generated placebo schedules of three paycheck months.

TABLE 1
SUMMARY STATISTICS

	Full Sample		Biweekly	
Expenditure in Levels ($):				
Durables	1,378.9	(2,788.7)	1,410.2	(2,775.9)
Non-durables	1,429.6	(824.8)	1,432.6	(830.5)
Strictly Non-durables	1,050.7	(530.3)	1,044.2	(535.6)
Food	514.0	(320.1)	510.7	(331.4)
Total	2,808.5	(3,037.2)	2,842.8	(3,020.6)
Changes in Expenditure ($):				
Durables	16.9	(3,803.0)	12.9	(3,773.7)
Non-durables	26.2	(456.4)	26.3	(458.1)
Strictly Non-durables	8.1	(251.4)	7.2	(257.2)
Food	3.4	(128.3)	3.5	(127.4)
Total	43.1	(3,846.4)	39.1	(3,815.0)
Taste Shifters:				
Age	39.8	(10.4)	39.9	(10.4)
Δ Children	0.0	(0.1)	0.0	(0.1)
Δ Adults	0.0	(0.1)	0.0	(0.1)
Income and Assets ($):				
Paycheck	1,523.9	(1,274.0)	1,668.8	(1,035.6)
Annual Income	53,852.6	(34,582.0)	55,254.3	(35,467.8)
Liquid Assets (N = 8887)	7,639.2	(28,587.6)	8,698.4	(33,115.4)
N	24,822		13,707	
Num. of Households	7776		4316	

NOTE.—Table entries are means and standard deviations. Observations are at the household-month level. Columns 1 and 2 are based on the full sample of households with weekly, bi-weekly, and monthly heads. These households include both single and married heads of households but exclude households for whom other members of the family are paid at the same frequency as the head of household. Additional details on sample restrictions can be found in Appendix A.1. Columns 3 and 4 include only the subset of households from the full sample whose heads are paid bi-weekly. All monetary values are in 2010 U.S. Dollars. Total expenditures are composed of durable and non-durable expenditure. Strictly non-durable expenditures are a subset of non-durables expenditures, and food expenditures are a subset of strictly non-durable expenditures. Age refers to the head of household only. Changes in the number of children include only children younger than 18. Paycheck amounts refer to the gross amount of the head of household's last pay. Annual income refers total before-tax income received by the households in the past year. Liquid assets are composed of savings and checking account balances and are only available for N=8887 household-month observations out of the full sample of 24,822 observations.

TABLE 2

Response to Extra Paychecks by Consumption Categories

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
$\mathbf{1}_{\{t-1 \in S\}}$	261.641***	256.964***	4.668	-2.434	-0.955
	(95.240)	(93.976)	(9.763)	(5.505)	(2.613)
Age	2.182	2.208	-0.026	-0.074	0.061
	(1.565)	(1.537)	(0.271)	(0.151)	(0.077)
Δ Children	317.430	327.762	-10.337	38.742	30.751
	(650.112)	(612.642)	(70.371)	(26.738)	(19.572)
Δ Adults	365.079	381.360	-16.288	17.602	51.522
	(659.013)	(620.910)	(78.352)	(57.922)	(42.160)
R-squared	0.002	0.002	0.034	0.006	0.005
N	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

Note.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The indicator $\mathbb{1}_{\{t-1 \in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 3

Response to Extra Paychecks by Durable Consumption Categories

	(1) Durable	(2) Vehicle Purchases	(3) Furniture	(4) Flooring	(5) Major Appliances
Dependent variable: ΔC_t					
$\mathbf{1}_{\{t-1 \in S\}}$	256.964***	257.101***	-5.149	3.545***	-3.658
	(93.976)	(91.984)	(10.106)	(1.269)	(4.619)
Age	2.208	2.181	-0.186	-0.072*	0.045
	(1.537)	(1.442)	(0.213)	(0.043)	(0.137)
Δ Children	327.762	545.596	1.869	0.403	1.378
	(612.642)	(494.669)	(5.614)	(0.905)	(3.903)
Δ Adults	381.360	549.231	-15.418	0.562	-4.805
	(620.910)	(595.488)	(22.359)	(0.517)	(17.788)
R-squared	0.002	0.002	0.002	0.002	0.003
N	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

Note.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column. Column 1 specifically uses changes in durable consumption and presents the same estimates as Column 2 of Table 2. Columns 2-4 decompose durable consumption and include the following sub-categories of consumption: vehicles purchases, furniture, flooring, and major appliances. The indicator $\mathbb{1}_{\{t-1 \in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 4

Response to Extra Paychecks by Vehicle Consumption Categories

	Dependent variable: ΔC_t					Probability of Purchase	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Vehicle Purchases	New Cars	Used Cars	New Motor Vehicles	Used Motor Vehicles	New Cars	Used Cars
$\mathbf{1}_{\{t-1 \in S\}}$	257.101***	201.398**	45.408	8.086	2.209	0.004**	0.002
	(91.984)	(78.243)	(47.997)	(7.015)	(4.833)	(0.002)	(0.003)
Age	2.181	0.982	1.320	-0.128	0.007	0.000	0.000
	(1.442)	(1.039)	(0.987)	(0.132)	(0.052)	(0.000)	(0.000)
Δ Children	545.596	459.922	88.801	-2.694	-0.433	0.015	-0.003
	(494.669)	(478.736)	(111.384)	(2.026)	(0.794)	(0.016)	(0.007)
Δ Adults	549.231	573.234	-26.894	2.698	0.193	0.019	0.013
	(595.488)	(527.437)	(264.625)	(2.371)	(1.229)	(0.018)	(0.024)
R-squared	0.002	0.003	0.001	0.002	0.001	0.003	0.004
N	13,707	13,707	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y	Y	Y

Note.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is listed at the head of each column. Column 1 presents the same estimates as Column 2 of Table 3. Columns 2-5 use (2010) month-to-month dollar changes in consumption as the dependent variable and decompose vehicle purchases into the following sub-categories of consumption: new cars, used cars, new motor vehicles, and used motor vehicles. Columns 5 and 6 report estimates from a linear probability model with indicators for whether or not a new car or a used car was bought as the dependent variable. The indicator $\mathbb{1}_{\{t-1 \in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). In addition to month and year fixed effects, all seven specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 5
Summary Statistics for Car Purchases

	Full Sample		Conditional on Ever Purchasing	
Expenditure in levels ($):				
New and Used Car Purchases	299.1	(2,529.1)	2,723.4	(7,187.1)
Out-of-pocket Expenditures	83.4	(1,112.9)	757.5	(3,281.9)
Financed Amount	215.8	(2,115.3)	1,965.9	(6,108.3)
Downpayment	24.0	(400.3)	217.2	(1,190.3)
Changes in expenditure ($):				
New and Used Car Purchases	-5.6	(3,636.2)	-49.1	(10,736.8)
Out-of-pocket Expenditures	3.4	(1,635.9)	29.9	(4,830.2)
Financed Amount	-9.1	(3,032.8)	-79.0	(8,954.9)
Downpayment	-0.3	(581.1)	-2.4	(1,715.8)
N	13707		1573	
Num. of Households	4316		407	

Note.—Table entries are means and standard deviations. Observations are at the household-month level. Columns 1 and 2 are based on the subset of households from the full sample whose heads are paid bi-weekly. Columns 3 and 4 includes only the subset of households from the bi-weekly sample who ever purchased a new or used car during the sample period (1990-2010). All monetary values are in 2010 U.S. Dollars. Expenditure categories in levels and in changes are composed of expenditures on the following: new and used car purchases, out-of-pocket expenditures, financed amounts, and downpayments. Here, new and used car purchases refers to the total dollar value of the car. Downpayments are *not* conditional on ever having made a downpayment.

TABLE 6

RESPONSE TO EXTRA PAYCHECKS BY VEHICLE FINANCING CATEGORIES

	(1) Car Purchases	(2) Out of Pocket Expenditure	(3) Financing	(4) Downpayment for Financed Purchases
Dependent variable: ΔC_t				
$\mathbf{1}_{\{t-1 \in S\}}$	246.806***	82.726*	164.080**	22.591
	(91.567)	(47.757)	(72.758)	(15.284)
Age	2.302	0.474	1.828	0.119
	(1.435)	(0.544)	(1.241)	(0.200)
Δ Children	548.723	-0.577	549.300	29.308
	(494.879)	(53.869)	(466.521)	(24.728)
Δ Adults	546.340	196.834	349.506	18.956
	(595.621)	(210.163)	(521.810)	(29.667)
R-squared	0.002	0.002	0.002	0.001
N	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. The category of consumption for each specification is listed at the head of each column and includes: car purchases, out-of-pocket expenditure, financing amounts, and downpayments. The indicator $\mathbb{1}_{\{t-1 \in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). In addition to month and year fixed effects, all seven specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 7

Response to Extra Paychecks by Vehicle Financing Categories Conditional on Ever Purchasing a Car

	(1) Car Purchases	(2) Out of Pocket Expenditure	(3) Financing	(4) Downpayment for Financed Purchases
Dependent variable: ΔC_t				
$\mathbf{1}_{\{t-1 \in S\}}$	2204.625 ***	733.266*	1471.358 **	197.366
	(808.612)	(414.674)	(653.034)	(128.901)
Age	22.248	4.730	17.518	1.885
	(15.306)	(5.568)	(13.225)	(2.394)
Δ Children	1729.798	-47.820	1777.618	110.083
	(1517.104)	(166.779)	(1428.327)	(94.184)
Δ Adults	2266.511	615.268	1651.243	46.771
	(1583.493)	(604.529)	(1396.793)	(101.867)
R-squared	0.019	0.014	0.019	0.014
N	1,573	1,573	1,573	1,573
Month and Year FEs	Y	Y	Y	Y

Note.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the subset of households from the bi-weekly sample who ever purchase a car during their survey period. The category of consumption for each specification is listed at the head of each column and includes: car purchases, out-of-pocket expenditure, financing amounts, and downpayments. The indicator $\mathbb{1}_{\{t-1 \in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). In addition to month and year fixed effects, all seven specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 8

RESPONSE TO EXTRA PAYCHECKS BY PAY FREQUENCY ACROSS CONSUMPTION CATEGORIES

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
$\mathbf{1}_{\{t-1 \in S\}}$	256.902***	248.258***	8.625	-0.289	-0.754
	(92.034)	(91.063)	(9.483)	(5.414)	(2.601)
$\mathbf{1}_{\{t-1 \in S\}} * \mathbf{1}_{\{p=Week\}}$	-265.934 **	-276.803 **	10.876	6.957	5.707
	(134.243)	(132.384)	(15.514)	(8.394)	(4.215)
$\mathbf{1}_{\{t-1 \in S\}} * \mathbf{1}_{\{p=Month\}}$	-596.749 **	-586.945 **	-9.686	-7.315	-5.845
	(240.331)	(238.034)	(25.639)	(15.006)	(8.116)
Age	2.252*	2.370*	-0.118	-0.048	0.082
	(1.232)	(1.213)	(0.195)	(0.111)	(0.062)
Δ Children	361.470	305.765	55.702	68.100**	42.416**
	(361.351)	(333.372)	(52.039)	(27.202)	(16.788)
Δ Adults	129.569	115.668	13.905	55.044	50.341**
	(403.367)	(376.980)	(52.182)	(35.842)	(25.315)
R-squared	0.002	0.002	0.036	0.006	0.005
N	24,822	24,822	24,822	24,822	24,822
Month and Year FEs	Y	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The indicator $\mathbb{1}_{\{t-1 \in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). The estimate for this indicator represents the response to extra paychecks for households whose head is paid bi-weekly (the omitted category). The two indicators $\mathbb{1}_{\{t-1 \in S\}}\mathbb{1}_{\{p=Week\}}$ and $\mathbb{1}_{\{t-1 \in S\}}\mathbb{1}_{\{p=Month\}}$ give the response to extra paychecks for households with weekly and monthly heads relative to the response for households with heads who are paid bi-weekly. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 9
RESPONSE TO EXTRA PAYCHECKS BY CONSUMPTION CATEGORIES WITH TIME ADJUSTMENT

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
$\mathbf{1}_{\{t-1 \in S\}}$	250.788**	252.892**	-2.104	0.001	2.743*
	(112.728)	(112.024)	(7.098)	(2.341)	(1.521)
Age	-0.728	-0.682	-0.046	0.045	0.099*
	(2.117)	(2.111)	(0.222)	(0.079)	(0.058)
Δ Children	231.268	181.908	49.360	39.885	36.254*
	(162.843)	(140.276)	(38.194)	(25.726)	(20.804)
Δ Adults	189.639	90.916	98.723**	44.619	65.899**
	(471.797)	(457.168)	(43.439)	(36.764)	(30.756)
R-squared	0.003	0.003	0.037	0.008	0.010
N	8,466	8,466	8,466	8,466	8,466
Month and Year FEs	Y	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. Measures of levels of consumption exclude any expenditures which underwent time adjustment routines by the Bureau of Labor Statistics as detailed in Section 1.5.3. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The indicator $\mathbb{1}_{\{t-1 \in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 10

RESPONSE TO EXTRA PAYCHECKS BY CONSUMPTION CATEGORIES BY LIQUID ASSET HOLDINGS

	(1)	(2)	(3)	(4)	(5)
	Total	Durable	Non-durable	Strictly ND	Food
Dependent variable: ΔC_t					
ΔY_t	0.136	0.142	-0.006	-0.004	-0.003
	(0.116)	(0.113)	(0.016)	(0.009)	(0.004)
$\Delta Y_t * \mathbf{1}_{\{Unconstrained\}}$	0.009	-0.006	0.016	0.012	0.008^{*}
	(0.157)	(0.154)	(0.020)	(0.011)	(0.005)
Age	6.637^{**}	6.581^{**}	0.056	0.135	0.221^{*}
	(3.207)	(3.156)	(0.504)	(0.251)	(0.127)
Δ Children	398.969	55.238	343.731^{**}	54.895	37.300
	(463.297)	(403.844)	(172.947)	(86.396)	(38.783)
Δ Adults	-786.871	-1.0×10^{3}	241.848	-28.817	17.681
	(817.497)	(769.647)	(188.882)	(96.221)	(48.714)
R-squared	0.005	0.003	0.045	0.012	0.010
N	4,863	4,863	4,863	4,863	4,863
Month and Year FEs	Y	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate 2SLS regression run at the household-month level, instrumenting for the change in income ΔY_t using the indicator for the previous month being a three paycheck month. The variable ΔY_t equals the 2010 dollar amount of the head of household's last gross pay if the previous month was a three paycheck month (if three paychecks of income are available in the present month t) and zero otherwise. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The estimate in the first row represents the response to extra paychecks for households who are constrained (the omitted category). Constrained households are defined as those for whom their liquid asset holdings are less than the median level of liquid assets; likewise, unconstrained households are those for whom liquid asset holdings of the household is greater than or equal to the median level of liquid assets. The estimate for $\Delta Y_t * \mathbb{1}_{\{Unconstrained\}}$ gives the response to extra paychecks for households that are unconstrained relative to the response for households who are constrained. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 11

RESPONSE TO EXTRA PAYCHECKS BY CONSUMPTION CATEGORIES BY TOTAL BEFORE-TAX INCOME

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
ΔY_t	0.083	0.075	0.008	-0.002	-0.002
	(0.082)	(0.081)	(0.009)	(0.005)	(0.002)
$\Delta Y_t * \mathbf{1}_{\{Unconstrained\}}$	0.115	0.123	-0.008	0.001	0.002
	(0.105)	(0.104)	(0.011)	(0.006)	(0.003)
Age	1.736	1.871	-0.134	-0.112	0.047
	(1.570)	(1.544)	(0.269)	(0.150)	(0.077)
Δ Children	325.758	339.688	-13.936	37.461	30.398
	(649.978)	(612.410)	(70.386)	(26.667)	(19.598)
Δ Adults	373.275	389.670	-16.402	17.570	51.559
	(658.173)	(620.983)	(78.393)	(57.829)	(42.123)
R-squared	0.004	0.003	0.035	0.006	0.006
N	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate 2SLS regression run at the household-month level, instrumenting for the change in income ΔY_t using the indicator for the previous month being a three paycheck month. The variable ΔY_t equals the 2010 dollar amount of the head of household's last gross pay if the previous month was a three paycheck month (if three paychecks of income are available in the present month t) and zero otherwise. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The estimate in the first row represents the response to extra paychecks for households who are constrained (the omitted category). Constrained households are defined as those for whom their total before-tax income is less than the median income; likewise, unconstrained households are those for whom the total before-tax income of the household is greater than or equal to the median income. The estimate for $\Delta Y_t * \mathbb{1}_{\{Unconstrained\}}$ gives the response to extra paychecks for households that are unconstrained relative to the response for households who are constrained. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 12

Response to Extra Paychecks by Consumption Categories by Age

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
ΔY_t	0.191**	0.192**	-0.001	0.000	0.000
	(0.081)	(0.081)	(0.008)	(0.004)	(0.002)
$\Delta Y_t * \mathbf{1}_{\{Unconstrained\}}$	-0.063	-0.071	0.007	-0.003	-0.001
	(0.108)	(0.107)	(0.011)	(0.006)	(0.003)
Age	-0.269	0.321	-0.589	-0.215	0.063
	(2.596)	(2.513)	(0.522)	(0.276)	(0.139)
Δ Children	324.036	334.088	-10.056	39.033	30.780
	(649.934)	(612.337)	(70.514)	(26.807)	(19.548)
Δ Adults	363.103	379.121	-16.026	17.149	51.438
	(654.937)	(616.875)	(78.283)	(57.820)	(42.044)
R-squared	0.003	0.003	0.034	0.006	0.005
N	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

Note.—Each column reports estimates from a separate 2SLS regression run at the household-month level, instrumenting for the change in income ΔY_t using the indicator for the previous month being a three paycheck month. The variable ΔY_t equals the 2010 dollar amount of the head of household's last gross pay if the previous month was a three paycheck month (if three paychecks of income are available in the present month t) and zero otherwise. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The estimate in the first row represents the response to extra paychecks for households who are constrained (the omitted category). Constrained households are defined as those for whom the age of the household is less than the median age for households heads (age 39); likewise, unconstrained households are those for whom the age of the household is greater than or equal to the median age for households heads. The estimate for $\Delta Y_t * \mathbb{1}_{\{Unconstrained\}}$ gives the response to extra paychecks for households that are unconstrained relative to the response for households who are constrained. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 13

Response of Aggregate Consumption Measures by Committed Consumption as Fraction of Monthly Wages

	Total	Durable	Non-durable	Strictly ND	Food
Dependent variable: ΔC_t					
ΔY_t	0.174	0.192	-0.018	-0.009	-0.003
	(0.123)	(0.122)	(0.012)	(0.007)	(0.003)
$\Delta Y_t * \mathbf{1}_{\{Unconstrained\}}$	-0.022	-0.056	0.034^{**}	0.012	0.005
	(0.154)	(0.154)	(0.014)	(0.008)	(0.004)
Age	1.551	1.668	-0.117	-0.058	0.051
	(1.563)	(1.538)	(0.275)	(0.155)	(0.078)
Δ Children	315.134	324.361	-9.230	39.659	30.881
	(647.718)	(609.701)	(70.487)	(26.676)	(19.525)
Δ Adults	370.191	386.478	-16.294	17.562	51.489
	(656.543)	(618.500)	(78.232)	(57.994)	(42.137)
R-squared	0.003	0.003	0.035	0.006	0.005
N	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

Note.—Each column reports estimates from a separate 2SLS regression run at the household-month level, instrumenting for the change in income ΔY_t using the indicator for the previous month being a three paycheck month. The variable ΔY_t equals the 2010 dollar amount of the head of household's last gross pay if the previous month was a three paycheck month (if three paychecks of income are available in the present month t) and zero otherwise. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The estimate in the first row represents the response to extra paychecks for households who are constrained (the omitted category). Constrained households are defined as those for whom the level of committed consumption as a fraction of monthly wages is less than the median; likewise, unconstrained households are those for whom the level of committed consumption as a fraction of monthly wages is greater than or equal to the median for households heads. Committed consumption is the sum of household expenditures on mortgage payments, rental payments, vehicle loan payments, and utilities payments for a given month. Monthly wages is based on typical wage income and is constructed using the amount of the head of household's last gross pay. The estimate for $\Delta Y_t * \mathbb{1}_{\{Unconstrained\}}$ gives the response to extra paychecks for households that are unconstrained relative to the response for households who are constrained. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 14
Response to Extra Paychecks by Consumption Categories by Consumption Volatility

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
ΔY_t	0.102	0.109	-0.007	-0.017***	-0.007***
	(0.081)	(0.080)	(0.009)	(0.005)	(0.003)
$\Delta Y_t * \mathbf{1}_{\{Unconstrained\}}$	0.132	0.110	0.022**	0.035***	0.015***
	(0.105)	(0.104)	(0.011)	(0.006)	(0.003)
Age	1.821	1.859	-0.037	-0.075	0.054
	(1.557)	(1.531)	(0.270)	(0.150)	(0.076)
Δ Children	314.686	325.873	-11.193	37.338	30.096
	(649.074)	(611.445)	(70.266)	(27.268)	(19.538)
Δ Adults	362.960	380.611	-17.658	15.243	50.448
	(655.495)	(617.540)	(78.421)	(57.750)	(41.945)
R-squared	0.004	0.003	0.034	0.008	0.008
N	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

Note.—Each column reports estimates from a separate 2SLS regression run at the household-month level, instrumenting for the change in income ΔY_t using the indicator for the previous month being a three paycheck month. The variable ΔY_t equals the 2010 dollar amount of the head of household's last gross pay if the previous month was a three paycheck month (if three paychecks of income are available in the present month t) and zero otherwise. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The estimate in the first row represents the response to extra paychecks for households who are constrained (the omitted category). Constrained households are defined as those for whom the level of consumption volatility is less than the median consumption volatility for households; likewise, unconstrained households are those for whom the level of consumption volatility is greater than or equal to the median age for households . The estimate for $\Delta Y_t * \mathbb{1}_{\{Unconstrained\}}$ gives the response to extra paychecks for households that are unconstrained relative to the response for households who are constrained. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE 15

COMPARISON OF VEHICLE LOAN CHARACTERISTICS BY TIMING OF PURCHASE

	Cars Purchased Following Three Paycheck Months	Cars Purchased Other Months	t-stat. of Test of Equality
Maturity	48.10	49.46	-0.51
Monthly Payment	380.61	413.21	-1.42
Loan-to-Value Ratio	0.87	0.87	-0.74
(Conditional) Downpayment	2900.50	3837.55	-1.53
N	132	298	
Num. Financed	67	141	

NOTE.—The table compares vehicle loan characteristics for cars based on the timing of their purchase. The four vehicle loan characteristics of interest are loan maturity, monthly payment, loan-to-value ratio, and the dollar amount of downpyament (conditional on making a downpayment). The first column presents the mean of each loan characteristic for cars purchased following three paycheck months; the second column presents mean values for cars purchased in other months. The third column provides the t-statistic from testing for equality of the means in Columns 1 and 2. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

CHAPTER 2 : Social Recognition in Charitable Giving: In Pursuit of Perfection

2.1. Introduction

People care deeply about receiving public recognition for their accomplishments and good deeds. For evidence of this in the marketplace, we can look to the prevalence of conspicuous consumption, the sale of trophies, and the popularity of tiered donor recognition lists, which many not-for-profit organizations display prominently on their building walls, websites, and newsletters. Past research by psychologists and economists has empirically demonstrated that people find public recognition appealing and motivating (see Heffetz and Frank 2011 for a review). Individuals are more likely to engage in prosocial behavior, such as exerting effort for a charity (Ariely et al. 2009), volunteering (Linardi and McConell 2011), and making religious offerings (Soetevent 2005), when their actions are made visible rather than kept private.[1]

Charitable organizations often use various forms of public recognition as an incentive for giving: donor lists are circulated in newsletters and annual reports, naming opportunities are provided for new facilities or scholarships, and donor plaques are displayed on walls. The prevalence of charitable organizations using public recognition to encourage giving and the variety of ways in which they do so suggests that charitable organizations well understand that prosocial behavior, and in particular charitable giving, may be motivated by a desire for social recognition or acclaim. While many charitable organizations use public recognition to elicit contributions, the empirical literature on social recognition and charitable giving is relatively scarce. Past work on this topic has been primarily conducted in laboratory settings (see for example: Andreoni and Bernheim 2009; Ariely et al. 2009; Li and Riyanto 2009; Bracha and Vesterlund 2013). Notable exceptions include Harbaugh (1998a), who uses data on alumni giving to estimate the prestige benefits of receiving social recognition for making charitable contributions, and Karlan and McConnell (2013), who conduct a lab and field experiment to determine whether social image concerns or the desire to encourage others to give is the dominant driver of giving when donors are publicly recognized.

Better understanding the motivations behind charitable giving and the role that public recognition plays in encouraging giving is of great interest to both researchers and fundraisers, especially given the considerable number of households who engage in charitable giving. Over two-thirds of households in the United States make a charitable contribution of some form annually. These contributions are significant, with charitable giving to not-for-profit

[1] There is a large empirical literature demonstrating that people engage in more pro-social behavior when their actions are visible to an audience rather than private. See for example, Bohnet and Frey (1999), Andreoni and Petrie (2004), Rege and Telle (2004), Soetevent (2005), Dana et al. (2006), Ariely et al. (2009), Linardi and McConell 2011, Samak and Sheremeta (forthcoming).

organizations in the United States totaling over $316 billion in 2012. Individual donations represent by far the largest share of this total, comprising 72 percent of charitable contributions altogether (Giving USA 2013).

We build on past research on the power of public recognition as a motivating force for charitable giving and present an archival field study of how public recognition affects donations to one large, private university in the Northeastern United States (the "University"). In our study, we examine the effects of public recognition for consecutive giving on donation behavior. To do so, we consider two programs that were introduced by the University during our sample period. The first program recognized young alumni for having made consecutive gifts since their senior year (i.e. for having a perfect giving record); the second program recognized alumni for having made consecutive gifts for three or more years in a row. We identify the causal effect of public recognition for consecutive giving using a difference-in-differences strategy that exploits variation in alumni eligibility for recognition under the two programs and the timing of when the programs were and were not in effect.

Our analysis uses panel data on alumni giving over a ten-year period from 2002 to 2011. Higher education institutions depend heavily on alumni for charitable support. In 2012, alumni giving totaled approximately $7.7 billion and accounted for nearly 25 percent of contributions made to colleges and universities in the U.S. (Council for Aid to Education 2014).[2] Fundraising efforts by higher education institutions can provide a particularly useful setting for studying donation behavior. In addition to having a large and growing base of alumni as potential donors, higher education institutions often solicit alumni on a recurring basis after graduation, thus allowing donor behavior to be tracked over time. Accordingly, a small but growing literature has begun to use fundraising efforts by colleges and universities to study different motivations behind donation behavior (Harbaugh 1998a; Meer and Rosen 2008a, 2009, 2011; Holmes 2009; Meer 2011).

In the main analysis, we focus on both the probability that an individual made a donation and the (log) dollar amount donated. Alumni donations can be made either directly to the University's main unrestricted giving program or indirectly to other University priorities. While alumni receive general public recognition for donations made to either destination, only consecutive contributions made directly to the University's main giving program are eligible for additional recognition under the two programs we study in this paper. As it turns out, this distinction proves useful. One of the central limitations of many empirical studies on charitable giving is that it is generally not possible to observe giving by the

[2] The Council for Aid to Education conducts an annual survey to collect data on contributions (by source) made to over 1,000 higher education institutions. See http://cae.org/fundraising-in-education/category/home/ for additional details about the survey.

same individuals to multiple domains, making it difficult to determine the overall effect of a program or treatment on giving.[3] While we are not able to observe overall giving in all domains by individuals in our sample, the distinction in types of donations that are recognized under the two recognition programs of interest allows us to not only study the effect of recognition for consecutive giving on donations made directly to the University's main giving program, but also to estimate any potential crowd out effect on donations made to other University priorities.

Using the identification strategy outlined above, we find that public recognition for consecutive giving has a strong positive effect on donation behavior. The introduction of recognition programs for consecutive giving significantly increased the probability of giving by between 15.8 to 19.3 percent of the baseline probability of giving for eligible alumni. There is a significant corresponding increase in the dollar amount donated directly to the University's main giving program. Looking at donations made to other University priorities, we find that the introduction of the recognition programs for consecutive giving has a *crowd in* effect and significantly increased the probability of giving and dollar amount donated to other University priorities. Rather than leading individuals to substitute donations away from other University priorities toward the University's main giving program, our results suggests that the recognition program induced individuals to increase giving to the University as a whole.[4]

We examine two additional questions of interest. First, we consider whether individuals exhibit strategic behavior in the amount they choose to donate. As previously mentioned, alumni in our sample are publicly recognized by name in an annual newsletter for their donation. The newsletter lists donation amounts by discrete categories of support or tiers (e.g. \$500 to \$999).[5] Consistent with Harbaugh (1998a,b), Cartwright and Patel (2013), and Barbieri and Malueg (forthcoming), we observe evidence that individuals "bunch" at the lower ends of the categories of support. We estimate the effect of public recognition for consecutive giving on the probability of bunching, conditional on haven made a donation in the previous year, and whether individuals are more likely to bunch by increasing, decreasing, or keeping constant their donation relative to their donation in the previous year. We find suggestive evidence that the recognition programs increased bunching. The introduction of the PG programs had a significant and positive effect on the probability of bunching for young alumni, with the response primarily coming from alumni who bunch at

[3] There exists a rich literature studying whether government grants crowd out private contributions. See Steinberg (1991), Andreoni (2006), and Vesterlund (2006) for an overview of this literature. However, little work has been done on intrapersonal crowd out effects.

[4] It is also possible that individuals are substituting donations over time.

[5] Categories of support are also sometimes referred to as giving circles.

the same or lower support category. In contrast, the introduction of the ISS program had a significant and positive effect on the probability for non-young alumni, with the response primarily coming from alumni who bunch at the same or higher support category.

Finally, we look at how the effect of public recognition for consecutive giving varies by whether the recognition does or does not convey additional information regarding the individual's past donation history. Previous models of behavior have proposed that through their visible actions, individuals are able to signal their personal traits to others, such as being generous or pro-social versus stingy and selfish (e.g., Bénabou and Tirole 1996; Ellingsen and Johannesson 2008; Andreoni and Bernheim 2009). Such signaling models help explain the propensity for social recognition opportunities to increase pro-social behavior. In a charitable giving setting, individuals may use their donations to signal generosity or altruism (Hollander 1990; Cartwright and Patel 2013), wealth or status (Ireland 1994; Glazer and Konrad 1996; Harbaugh 1998a,b), or the quality of the charitable organization receiving their donations (Vesterlund 2003; Potters et al. 2005; Rondeau and List 2008). A signaling motive depends on whether the public recognition the donor receives provides other individuals with information that would otherwise be unobservable. For example, an individual whose wealth was publicly known would likely not find public recognition of his donation to be a very useful signal of his wealth.[6] Simply put, we look to see whether individuals value recognition more when the recognition conveys information that can be used to signal a personality trait — whether it be his generosity, wealth, or intelligence — than when the recognition does not.

To determine whether alumni exhibit differential donation behavior depending on whether the public recognition they receive for consecutive giving conveys additional information, we now consider a third recognition program that was in effect during the entirety of our sample period. This program recognized alumni for having made consecutive gifts for five or more years in a row. We next classify alumni into two groups: the first group of alumni received recognition under the program recognizing five or more years of consecutive giving while the second group of alumni does not (fewer than five years of consecutive giving).

Public recognition for three or more years of consecutive giving for alumni does not convey additional information for alumni in the first group since it is already publicly observable that they have made consecutive donations for five or more years; however, recognition for three or more years of consecutive giving *does* convey additional information for alumni in the second group by making public information regarding their past donation history that was previously unknown. We exploit this difference and find that while the introduction of

[6]This does not preclude, however, the donation from being a useful signal of other personal traits such as generosity or altruism.

the program recognizing three or more years of consecutive donations significantly increased the probability of giving and dollar amount donated for both alumni groups, the effects are much larger for alumni in the second group, for whom the recognition conveyed additional information, than for alumni in the first group.

The remainder of this paper is organized as follows. Section 2.2 describes the data while Section 2.3 presents the empirical framework. Section 2.4 describes the main results. Then Section 2.5 discusses strategic bunching behavior by donors. Section 2.6 evaluates the role of recognition versus information. Finally, Section 2.7 concludes.

2.2. Data Description

The data used in this paper is the history of undergraduate alumni donations to the University of Pennsylvania. All alumni names were replaced with anonymous unique identifiers prior to being shared. Donations can be made to a variety of programs and funds within the University but are designated more generally in the data as either "Penn Fund" or "Non-Penn Fund" gifts. These two designations distinguish between donations made specifically to the Penn Fund, the University's unrestricted undergraduate alumni giving program, and donations made to other University priorities.

While the data provides the entire history of donations, we focus on donations made between 2002 and 2011 when the changes in giving recognition were enacted and the University considers the data to be most accurate. The year of donation refers to the University's fiscal year, which runs from July 1 to June 30, rather than to the calendar year. We aggregate observed donations to the individual-year-designation level where the designation distinguishes between Penn Fund and Non-Penn Fund donations. The donation level for an alumnus in our final estimation sample is thus the total amount he gave in that fiscal year. If no donation was made to a given designation in a particular year, the alumnus is included in the data as having made a zero dollar donation. For a given year, each alumnus appears twice in the data – once for his donation to the Penn Fund and once for his donation to other University priorities. In addition to donation behavior, the data includes information on gender, marital status, and the year in which the alumnus graduated from the University.

We begin with 2,795,340 individual-year-designation observations for 139,767 undergraduate alumni. We drop any individuals who are flagged as being deceased.[7] Since the recognition programs we examine are generally only applicable to alumni and not to current students, we further drop any observations that occur in years prior to an individual having graduated

[7] Several variables appear time-invariant in our data and represent the value of that variable for an alumnus at the time the data was extracted by the University. These variables include whether an individual is deceased and marital status.

from the University. The number of years an alumnus appears in our sample therefore depends on when he graduated from the University. Specifically, alumni who graduated in or after 2002 will appear fewer times in the data than alumni who graduated prior to 2002. Finally, we exclude observations where alumni are one year out from graduation (i.e. were seniors in the previous year). Seniors are differentially encouraged to donate by the University, so giving behavior in the previous year for the group of first-year graduates may differ greatly from that of non-first-year graduates (two or more years since graduation).[8] We are left with 2,149,234 observations for 117,937 undergraduate alumni.

Table 32 reports summary statistics for the estimation sample. Reported monetary values are in real 2010 dollars in this and all tables that follow unless otherwise noted. The alumni in our sample are 57 percent male, and 49 percent are married. The average number of years since graduation for a donation is 27.5 years. Roughly 17 percent of opportunities to give to the Penn Fund and 9 percent of opportunities to give to other University priorities lead to an actual donation. Conditional on giving, the mean donation amount to the Penn Fund ($1,723.64) is smaller than the mean donation to other University priorities ($11,759.74), but the median donation amount for the two designations do not greatly differ ($108.16 for the PennFund versus $111.48 for other University priorities).

2.2.1. Honor Roll Recognition

Donors to the University are recognized on the Honor Roll, an annual publication of the names of all donors and their associated donation level for the year. Figure 4 provides an example of a page in the Honor Roll recognizing alumni for their donations. Each donor has the opportunity to be listed twice, once for any donation to the Penn Fund and once for any donation to other University priorities. Rather than report exact dollar donations, the Honor Roll reports donations in bins or "categories of support." For example, an alumnus who donates $1000 in a given year and an alumnus who donates $2000 in that same year will both be listed in the $1000 to $2499 donation support category. While the true donation amount between the two alumni differs, they would appear no differently in the Honor Roll. Both Penn Fund and Non-Penn Fund donations are reported in this manner. The categories of donation support depend on the number of years since graduation for a given alumnus as well as the designation type of the donation and are listed in Table 17. The dollar amounts bracketing each category of support are well-publicized and do not change during our sample period.

[8] The University makes a strong effort to encourage giving during senior year by appealing to students' upcoming graduation. For instance, students can volunteer as part of a group known as "Seniors for the Penn Fund," a program of Penn Traditions that encourages alumni engagement. One of the main goals of Seniors of the Penn Fund is to secure participation in a yearlong fundraising campaign.

Consistent with prior literature, our data shows that when donations are reported using categories of support, individuals often "bunch" at the lower bound by donating the minimum dollar amount necessary to fall within a given category (Glazer and Konrad 1996; Harbaugh 1998a, 1998b). For each category of support, Table 18 provides statistics on the percentage of donations that were made at exactly the lower bound.[9] Categories for which the lower bound is $1 have very little bunching, reflecting the fact that conditional on giving, few people donate only one dollar. In contrast, the top two support categories, $10,000 to $24,999 and $25,000 and above, have between 10 and 30 percent of donations at the lower bound. For the remaining support categories, the percent of donations made at the lower bound ranges from 60 to as high as 85 percent.

There are two potential channels which might simultaneously explain this behavior. One explanation is that donors are motivated by prestige and receive utility from having their donation publicly recognized (Harbaugh 1998a, 1998b). To the extent that this is true, strategically donating at the lower bound of a support category is the lowest cost way to achieve the utility provided by recognition in that category. An alternative explanation is simply that individuals tend to give at round numbers.[10] However, descriptive evidence from the data suggests that a tendency to give at round numbers cannot entirely explain the bunching behavior we observe. Figure 6 examines the frequency of non-zero donation amounts in our data between $200 and $5000 in increasing increments of 50 dollars. Together, these donations account for over 92.7 percent of all non-zero donations made between $200 and $500. Donation amounts that correspond to lower bounds of support categories are assigned a different color (red) than donation amounts that do not (blue). Figure 6 illustrates that while many donors give at round numbers that do not correspond to the lower bounds of support categories, these donations are infrequent relative to donations made at the lower bounds. This suggests that at least some of the bunching behavior is strategic and motivated by prestige. In later sections, we discuss the implications of these two explanations for our results.

2.2.2. Consecutive Giving Recognition

In addition to appearing on the Honor Roll, alumni donors can receive further recognition based on their history of donation behavior to the Penn Fund. Three programs recognizing consecutive giving were in place during our sample period. Figure 5 shows the relative timing of each of these recognition programs. The first program was in effect during the entirety of our sample period from 2002 to 2011 and recognized donors for five or more

[9] Statistics on bunching were calculated using nominal donation amounts.

[10] See Pope and Simonsohn (2011) and Rosch (1975) for examples of individuals using round numbers as goals or reference points.

consecutive years of giving to the Penn Fund (5YCG). Eligible donations for consideration included donations made in an alumnus's graduation (senior) year. Thus the earliest point at which an alumnus could be potentially be eligible for this program was beginning in the fourth year following graduation. The recognition for this program took the form of an asterisk placed next to the donors name in the Honor Roll.

The second program, the Ivy Stone Society (ISS), was enacted in 2007 and recognized donors who made consecutive gifts to the Penn Fund for at least three years. For alumni who graduated less than three years prior, the program recognized individuals with consecutive giving since their graduation year. The ISS was meant to (and did in 2012) eventually replace the program recognizing five or more years of consecutive giving but overlapped with that program for several years during 2007 to 2011. The ISS recognition program continued through to the end of our sample period and took the form of a carat placed next to the donor's name in the Honor Roll.[11]

The third program, the Perfect Giving (PG) program, recognized alumni with perfect giving donation records to the Penn Fund. It was in effect between 2005 and 2008 and was only applicable to young alumni who were between one and four years since graduation. Donations under consideration included donations made in the graduation year for an alumnus, so to have a perfect giving record, an alumnus had to have given every year since his senior year. Alumni who were recognized under this program had their name listed twice on the Honor Roll: in addition to being listed in the appropriate donation support category, the alumni names were also in a separate list in the Honor Roll for perfect giving.

For each of these programs, we consider an alumnus to be *eligible* for recognition in a given year if making a positive donation that year would qualify for recognition under that program. Eligibility for recognition under these programs depends only on an alumnus's past donation history and on the number of years since his graduation and does not depend on his donation behavior in the current year. Figure 7 illustrates which alumni are considered eligible for each of the three recognition programs as a function of the number of years since graduation and the number of years of consecutive giving. Importantly, eligibility in a given year as defined in this paper is independent of whether the recognition program is actually in effect in that year.

For our empirical analysis, we focus on the enactment of two of the recognition programs in particular – the Ivy Stone Society and the Perfect Giving program – which were both

[11]Donors who receive ISS recognition in the Honor Roll also received an ISS newsletter as well as a bumper sticker. We consider these to be of negligible monetary value; however, to the extent that individuals value the newsletter and bumper sticker independently of their valuation of the additional recognition by their name in the Honor Roll, the interpretation of our results changes slightly.

introduced in the middle of our sample period. Tables 19 and 20 provide summary statistics for the estimation sample by ISS eligibility and by PG eligibility respectively.

2.3. Empirical Setup

We are interested in estimating whether individual donation behavior responds to changes in the recognition of consecutive giving. We do this using a difference-in-difference strategy that exploits two sources of variation: (i) variation in whether a recognition program was in effect and (ii) variation in whether a potential donor would be eligible to receive the recognition (i.e. variation in their donation history). In other words, we compare how donation behavior changes for individuals when they are and are not eligible for a recognition program when the program is and is not in effect. The identifying assumption underlying this approach is that any differences in donation behavior when the program is and is not in effect can be solely attributed to the presence of the program.

The ISS and PG programs provide additional recognition for consecutive giving specifically to the Penn Fund. However, the literature on charitable giving and public goods in general has long recognized the possibility that increased giving in one domain may crowd-out giving in other domains (CITES). For this reason, we focus on both the direct effect of these recognition programs on giving to the Penn-Fund and the indirect (crowd-out) effect on giving to other University priorities. We look at both the extensive and intensive margin of giving for the direct and crowd-out effects. For each of these outcomes of interest, Y_{it}, we estimate variants of the following model:

$$Y_{it} = \alpha_i + \beta_{1,j}(\textsc{Eligible})_{ijt} + \beta_{2,j}(\textsc{Eligible})_{ijt} * (\textsc{InEffect})_{jt} + \beta_3 X_{it} + \eta_t + \epsilon_{it} \quad (2.1)$$

where i indexes an alumnus, j is the recognition program, and t denotes the fiscal year. There are two key indicators in the estimating equation: $\textsc{Eligible}_{ijt}$ equals one if a positive gift by alumnus i in fiscal year t would receive recognition under program j; and $\textsc{InEffect}_{jt}$ equals one if recognition program j was in effect in year t. The vector X_{it} represents the fraction of years since graduation in which an alumnus donated and controls for variation in past donation behavior. We include individual fixed effects, γ_i, to account for potential time-invariant heterogeneity in preferences for giving. We also include both fiscal year and years since graduation fixed effects, η_t, to control for any time or age trends in donation behavior. Years since graduation fixed effects are further interacted with the control variable X_{it}. The main parameter of interest, $\beta_{2,j}$, is the difference-in-difference estimator of the effect of recognition program $j \in \{\text{ISS,PG}\}$ on our outcomes of interest.

2.4. Effect of Recognition for Consecutive Giving on Donation Behavior

In this section, we identify the effects of recognition for consecutive giving on donation behavior looking at measures of giving on the extensive and intensive margin both directly to the Penn Fund and indirectly to other University priorities. A challenge of cleanly identifying the causal effects of interest is the dynamic nature of our setting. In particular, eligibility for recognition by the PG and ISS programs is a function of an alumnus's donation behavior in the past year. We therefore present the results in two steps. To avoid potential feedback effects, we first restrict our sample for each program to include years through only the first year in which the program was in effect. Thus, for all alumni, $\textsc{InEffect}_{jt}$ will equal one for the first year of program j and zero for all years prior. Once we have established the effects of the PG and ISS programs, we then return to the full sample from 2002-2011 and re-estimate the effects of the two programs accounting for dynamics.

2.4.1. Effect of Recognition for Consecutive Giving with No Dynamics

Table 21 reports results from estimating (2.1) for the PG program. The sample spans 2002 through 2005 and includes only young alumni four or less years out since graduation.[12] Each column reports estimates of $\beta_{2,j}$ from a separate OLS regression with the outcome of interest listed at the head of the column. In these and all subsequent specifications, standard errors are clustered by individual to account for potential serial correlation in errors. The first two columns report the difference-in-difference estimates of the direct effect of the PG program on the probability of giving and the (log) dollar amount donated to the Penn Fund. The third and fourth column present analogous estimates for giving to other University priorities (Non-Penn Fund).

The table shows two main results. First, the enactment of the PG program significantly increased the probability of giving directly to the Penn Fund by 11.73 percentage points. This estimate is quite large and represents a 15.8 percent increase in the baseline probability of giving for PG-eligible alumni. Donations by alumni to the Penn Fund increase by 41.2 percent on average when the PG program was in place. The magnitude of this response is due in part to the fact that these estimates include donations by individuals induced to give on the extensive margin by the recognition programs. Second, we find that the enactment of the PG program led to *crowd in* of donations made to other University priorities. The probability of giving increased significantly by 8.17 percentage points or 81.9 percent of the baseline probability of giving to other University priorities by PG-eligible alumni, while donations increased on average by 40.9 percent.

[12]Because the PG provides recognition only to young alumni, graduates more than four years out are ineligible by definition. We thus exclude non-young alumni from this sample.

The sample used in Table 21 excludes all years after the first year of the PG program. We further restrict the sample to only the first year of the program and the year prior (2004 to 2005) and re-estimate the difference-in-difference estimates as before. These results are reported in Table 22 and are qualitatively similar to the estimates from the previous table. The estimated effects on donation behavior directly to the Penn Fund are larger in magnitude than the previous estimates but still highly significant and positive while the estimated effect on giving to other University priorities is not significantly different.

We next examine the effects of recognition under the Ivy Stone Society program. Here, the sample spans 2002 through 2007 and includes only alumni seven or more years out since graduation. The ISS program came into effect in 2007 when the PG program was also in effect and had been so since 2005. We exclude alumni less than seven years out since graduation to avoid potentially confounding the effects of the ISS program with the effects of the PG program. Table 23 presents the difference-in-difference estimates for the ISS program. The introduction of the ISS program significantly increased direct giving to the Penn Fund by 15.53 percentage points or by 19.3 percent of the baseline probability of giving for ISS-eligible alumni. Donations also significantly increased by 73.0 percent. As with the PG program, we find that recognition from the ISS program crowds in giving to other University priorities. The probability of giving increased by 4.41 percentage points or 15.7 percent of the baseline probability of giving while donations increased by 24.3 percent.

Table 24 shows that the estimated effect remains when the sample is restricted to include only years 2006 to 2007. The results are once again qualitatively similar to the estimates from the previous table though they are larger in magnitude. When restricting the sample to only first year of the program and one year prior, the enactment of the ISS program increased the probability of giving directly to the Penn Fund by 29.2 percent of the baseline and indirectly to other University priorities by 26.3 percent of the baseline. There was a corresponding increase in the dollar amount donated to both the Penn Fund and to other University priorities.

Taken together, these results suggest important benefits to charitable organizations of providing public recognition for consecutive giving. Both recognition for a perfect giving record and for consecutive giving three or more years in a row led to increased donations to the University as a whole.

2.4.2. Effect of Recognition for Consecutive Giving Accounting for Dynamics

We next consider the potential dynamics of the effect of the recognition programs on donation behavior. To do so, we estimate variants of the specification from 2.1, with the inclusion of additional indicators for eligibility in future years. Throughout our analyses, we make the assumption that once a recognition program has been introduced, alumni believe that the program will continue to exist indefinitely into the future.[13] Thus, individuals who are not currently eligible for recognition under the ISS program (i.e. for whom a positive donation in the current year would not qualify for recognition) might still choose to give in the current year in anticipation of becoming eligible for recognition under the ISS program in the following years. Because an individual is considered ISS eligible in the current year if he has made at least two consecutive gifts in a row as of the beginning of the current year, we can include at most two indicators for ISS eligibility in the future: $\text{ISSEligibleInTwoYears}_{it}$ which equals to one for anyone who did not give in the previous year and therefore must give consecutively in the next two years to become ISS eligible and $\text{ISSEligibleInOneYear}_{it}$ which equals to one for anyone who gave in the previous year but not for the past two years prior and therefore must give again in the next year to become ISS eligible. Individuals who are not currently eligible for recognition under the PG program cannot become eligible in future years since the program recognizes individuals with perfect giving records since graduation. Accordingly, we do not include indicators for PG eligibility in the future.

We now include the full sample period from 2002 through 2011 during which both the PG and ISS programs were in effect. The analysis is preformed separately for young alumni and non-young alumni, as it is only young alumni who are eligible for recognition under the PG program. Table 25 presents estimates for the sample of young alumni. Alumni who are eligible for ISS recognition in two years (i.e. who did not give in the previous year) are the omitted group, so estimated effects can be interpreted as relative to the effect for those who did not give directly to the Penn Fund in the previous year. Our analysis here leads to three main findings. First, consistent with our previous findings, the enactment of the PG recognition program significantly increasing giving both directly to the Penn Fund and indirectly to other university priorities. Second, while the introduction of the ISS recognition program significantly increased giving for alumni who would be eligible for ISS recognition in one year, the significant positive effect on giving for alumni currently eligible for ISS is greater in magnitude. Third, after controlling for ISS eligibility, we find that the ISS program has no effect for giving to other University priorities.

Focusing next on non-young alumni, Table 26 looks at the effect of the ISS program on donation behavior, controlling for ISS eligibility in future years. Once again, the omitted group is the set of alumni who did not give in the previous year and thus are eligible for ISS recognition in two years. Because the PG program did not publicly recognize non-

[13] Alumni do not receive information regarding the introduction or conclusion of recognition programs until the year in which these changes take place.

young alumni, its effects are not estimated here. Consistent with previous results ignoring dynamic effects, we find evidence that the introduction of the ISS recognition program has a positive and significant effect on both the probability of giving and the dollar amount donated after controlling for potential ISS eligibility in future years and that this effect is greater in magnitude than the effect for alumni eligible for ISS recognition in one year. Public recognition for three or more years of consecutive giving significantly increased giving directly to the Penn Fund and indirectly to other University priorities.

2.5. Do donors exhibit strategic behavior in response to recognition programs?

Having established that recognition for consecutive giving through the ISS and the PG programs increases giving, we next examine whether donors exhibit strategic behavior in response to these programs. In particular, we look at whether bunching behavior increases in response to the enactment of the ISS and PG programs, and if so, whether individuals bunch at a higher or lower support category as their previous donation. To do so, we estimate specifications of the form:

$$\begin{aligned} Y_{it} = \alpha_i + \beta_{1,j}(\text{ELIGIBLE})_{ijt} + \beta_{2,j}(\text{ELIGIBLE})_{ijt} * (\text{INEFFECT})_{jt} \\ + \beta_3 \text{BUNCHED}_{i,t-1} + \beta_4 X_{it} + \eta_t + \epsilon_{it} \end{aligned} \tag{2.2}$$

where i indexes an alumnus, j is the recognition program, and t denotes the fiscal year. The covariates are defined as before in (2.1) with the addition of the indicator $\text{BUNCHED}_{i,t-1}$ which equals one if a donation was made at the lower bound of a support category the previous year. We estimate (2.2) using only the sample of donations made directly to the Penn Fund. Once again, our coefficient of interest is the difference-in-difference estimator $\beta_{2,j}$ of the effect of recognition program $j \in \{\text{ISS,PG}\}$.

Once again, we proceed with our analysis in two steps. First, we restrict our sample for each program to include years through only the first year in which the program was in effect. This restriction is done so as to avoid potential feedback effects and thus provide the cleanest first look for our analysis. We then return to the full sample and re-estimate the effects of the two programs on bunching behavior accounting for dynamics.

2.5.1. Bunching with No Dynamics

We are interested in estimating the effect of the PG and ISS recognition programs on bunching behavior, conditional on having made a positive donation to the Penn Fund in the previous year. While both donations made to the Penn Fund and donations made to other university priorities are listed in the honor roll using categories of support, we focus our analysis on donations made directly to the Penn Fund. Table 27 reports results from estimating (2.2) for the PG program and considers four main outcomes of interest. Column 1 looks at the conditional effect of the PG program on the probability of bunching. The last three columns decomposes the response by whether the dollar amount donated strictly increased, strictly decreased, or did not change from the previous year. As before, the sample is restricted to include years through only the first year in which the PG program was in effect and spans 2002 through 2005. We include only young alumni in our sample since the PG program does not provide recognition for non-young alumni. We find that the enactment of the perfect giving program did not have an effect on bunching behavior, though we note that the sample size is necessarily quite small and likely has resulted in concerns over power.

The effects of the ISS program stand in sharp contrast to the estimated effects for the PG program on bunching behavior. As before, the sample spans 2002 through 2007 and includes only alumni seven or more years out since graduation. Table 28 shows that the introduction of the ISS program significantly increased bunching behavior by 3.57 percentage points. This effect is quite large given a baseline probability of 9.47 percent for ISS-eligible alumni. Columns 2 through 4 show that this large response is relatively evenly distributed between alumni who increase, decrease, and keep their donation amounts the same.

2.5.2. Bunching Accounting for Dynamics

Next, we re-estimate (2.2), this time taking into consideration the potential dynamics of the effect of the ISS and PG programs on donation behavior. We proceed as before and include additional indicators for eligibility in future years. Because we estimate the effects of the two recognition programs conditional on having made a positive donation in the previous year, we can include at most one indicator for ISS eligibility in the future: $\text{ISSEligibleInOneYear}_{it}$ which equals to one for anyone who gave in the previous year but not for the past two years prior and therefore must give again in the next year to become ISS eligible. Again, individuals who are not currently eligible for recognition under the PG program cannot become eligible in future years, so we do not include indicators for PG eligibility in the future.

The analysis is preformed separately for young alumni and non-young alumni and uses the full sample period from 2002 through 2011 during which both programs were in effect. Table 29 presents estimates for the sample of young alumni. Alumni who are eligible for ISS recognition in one year (i.e. who, conditional on having given the previous year, are ineligible for ISS recognition) are the omitted group. We find that the enactment of the PG recognition program significantly increased the probability of bunching by 7.1 percentage points off of a baseline probability of 8.9 percent. Decomposing this response, we see that much of this effect is driven by individuals bunching at a lower support category or bunching in the same support category as their previous donation, though the effect is not statistically significant for bunching in the same support category.

We next focus on non-young alumni and look at the effect of the ISS program on donation behavior, controlling for ISS eligibility in future years. Once again, the omitted group is the set of alumni who are eligible for ISS recognition in one year. In contrast to the young alumni, Table 30 shows that the introduction of the ISS recognition program has a positive and significant effect on the probability of bunching and that this response is driven by individuals bunching either at the same support category as their previous donation or bunching at a *higher* support category.

2.6. Recognition versus Information

Fundraising organizations often provide social recognition to donors as an incentive for giving. The value of such recognition and the motivations behind why it may or may not encourage giving are complex. Potential donors may value social recognition simply as reciprocal recognition or a "thank you" from the recipient organization. Alternatively, donors may value social recognition as a public signal of information such as wealth, generosity, loyalty, or other individual characteristics (citations: Frank 1985, Glazer and Konrad 1996, Harbaugh 1998a and Millet and Dewitte 2007).

In this section, we consider the relative importance of recognition and information in donation behavior. To do so, we exploit the fact that the 5YCG program recognizing donors for five or more consecutive years of giving to the Penn Fund was in effect throughout our sample period. For individuals eligible for recognition under the 5YCG program, the introduction of the ISS program, which recognized three or more consecutive years of giving, would result in additional recognition on the honor roll but would not convey any additional information regarding the alumnus's past donation history. We begin by estimating the following variant of our main equation:

$$\begin{aligned} Y_{it} = \alpha_i &+ \beta_1(\textsc{Eligible})_{it} \\ &+ \beta_{2,1}(\textsc{Eligible})_{it} * (\textsc{InEffect})_t * (\textsc{LessThanFiveYrsConsecGiving})_{it} \\ &+ \beta_{2,2}(\textsc{Eligible})_{it} * (\textsc{InEffect})_t * (\textsc{FivePlusYrsConsecGiving})_{it} \\ &+ \beta_3 X_{it} + \eta_t + \epsilon_{it} \end{aligned} \tag{2.3}$$

where i indexes an alumnus and t denotes the fiscal year. We focus on giving to the Penn Fund and to other University priorities as our main outcomes, Y_{it}, of interest. There are four key indicators in the estimating equation: $\textsc{Eligible}_{it}$ and $\textsc{InEffect}_t$ are defined as before and refer respectively to alumni eligibility for ISS recognition and whether the ISS program was in effect; $\textsc{LessThanFiveYrsConsecGiving}_{it}$ equals one if alumnus i has made consecutive donations for less than five years as of the beginning of fiscal year t; and $\textsc{FivePlusYrsConsecGiving}_{it}$ equals one if alumnus i has made consecutive donations for five or more years as of the beginning of fiscal year t. The remaining covariates are defined as before. Our coefficients of interest are the difference-in-difference estimators $\beta_{2,1}$ and $\beta_{2,2}$ which tell us the differential effect of ISS recognition for alumni who were and were not already recognized under the 5YCG program. Because $\textsc{LessThanFiveYrsConsecGiving}_{it}$ and $\textsc{FivePlusYrsConsecGiving}_{it}$ are separate indicators, the estimates should be interpreted as the total effect of the ISS recognition program for each of the two groups of alumni.

Table 31 reports results from estimating this specification. As before with the estimated aggregate effect of ISS recognition in Section 2.4.1, our estimation sample spans 2002 to 2007 and includes only alumni seven or more years out since graduation. The first column reports the difference-in-difference estimates of the direct effect of the ISS program on the probability of giving to the Penn Fund. The second column reports analogous estimates of the indirect effect of the program on giving to other University priorities. Row 2 reports estimates of the effect of the ISS program on the probability to give when the recognition from the ISS program not only provides recognition but also signals additional information regarding the alumnus's past donation history. Row 3 reports the estimated effect when the ISS program does *not* signal additional information since these alumni already received recognition under the 5YCG program.

The results show that the estimated effect of the ISS program on the probability of giving is significantly higher when recognition under the ISS program signals information in addition to providing recognition (22.82 percentage points or 31.3 percent of baseline) than when the ISS program only provides recognition (8.37 percentage points or 9.3 percent of baseline). In contrast, there is no significant difference in the effect of the ISS program on giving to other University priorities for the two groups of alumni. The absence of any difference for giving to other University priorities may be due to the fact that the ISS program provides recognition only for giving directly to the Penn Fund and thus may induce differential giving only for Penn Fund donations.

Figure 9 further illustrates the role of recognition versus information by showing the difference-in-difference estimates of the effect of the ISS program by the number of years since graduation. Consistent with the results in Table 31, estimates of the effect for giving to the Penn Fund are consistently significantly higher for alumni who were not already recognized under the 5YCG program (estimates to the left of the dashed red line in the figure) than for those who were; there is no significant difference in estimates by years out since graduation for donations made to other University priorities. Taken together, these results suggest that not only do alumni value social recognition but also that they value it even more when such recognition signals additional information about the alumnus.

2.7. Discussion

This paper uses panel data on alumni giving to a major university to address whether social recognition for consecutive donations encourages charitable giving behavior. Using a difference-in-differences identification strategy, we find that publicly recognizing individuals for making consecutive donations has a significant positive effect on both the probability of give and the (log) dollar amount donated. Further examination finds that the recognition programs we study in fact crowd *in* donations made to other University priorities. The magnitudes of the effects we find are quite large. Public recognition for consecutive giving increases the probability of giving by between 15.8 to 19.3 percent of the baseline probability of giving for eligible alumni. By comparison, being solicited by a peer increases the probability of giving by 8.5 percent of the baseline (Meer 2011). This suggests that recognition for consecutive giving has meaningful implications for charitable giving solicitation and rewarding.

There are a number of open questions that remain, which we leave to future research. Though we demonstrate that public recognition encourages charitable giving more so when the recognition conveys some otherwise unobservable information, we do not address the specific motivations driving individuals desire for public recognition for consecutive giving. This desire may stem from a wish to signal generosity, loyalty, wealth, or a combination of such traits. We remain agnostic as to what traits individuals aim to signal by making a publicly recognized donation.

Another question of interest relates to our finding that public recognition for consecutive giving directly the University's main unrestricted giving program leads to crowd in of donations made to other University priorities. One possible explanation for this finding is that the recognition programs increased feelings of loyalty or pride, leading alumni to increase their giving to the University as a whole. Alternatively, there may be potential time and effort savings from making donations to the University's main giving program and to other University priorities simultaneously. To the extent that the two recognition programs increased contributions directly to the main giving program, alumni may then trade off making a donation to other University priorities in the future for making a donation to other University priorities in the present. Understanding what drives the crowd in effect has important policy implications for the provision of public goods more generally.

Finally, there is the important question of how well the results of our study extend to other settings. Higher education institutions by definition have a prior relationship with the alumni they solicit, which may ease their efforts in appealing to potential donors' feelings of loyalty, pride, or reciprocity. Additionally, the fundraising we study in this paper differs from capital campaigns, which may be more one-shot affairs rather than a continuing campaign for contributions. Further research exploring these questions would prove valuable.

Fig. 4.—Honor Roll Sample

Class of 1998 Honor Roll

Gifts made between July 1, 2009 – June 30, 2010

Gifts of $25,000 or more to The Penn Fund

Anonymous (2)
Allison J. Brody ⲷ
Daniel J. Goldring *^
Libby P. Goldring *^
Eric T. Lee *^

Gifts of $25,000 or more to other University Priorities

Vikram Govindan
Jason H. Karp
David C. Wang

Gifts of $10,000 to $24,999 to The Penn Fund

Anonymous *^
Wendy Blank Chaikin ^•
Lance A. Milken *^
Michael A. Pollack
Benjamin J. Whitfield

Gifts of $10,000 to $24,999 to other University Priorities

Arif T. Joshi
Faisel H. Khan
Jennifer Law
Kourtney Matter
Rishi K. Patel
Ashish Y. Shah
Adam Gordon Silfen

Gifts of $5,000 to $9,999 to The Penn Fund

Kenneth S. Henderson *^
Jason H. Karp •
Colin D. Lang ^
Mary Woland Lang ^
Michael G. Linn *^
Maria Magliacano *^ Charles
Lawrence Myers *^
Jessica Resnick Myers *^
Nicole T. Totah ^
Daisy Y. Tung
Michael K. Weaver

Gifts of $5,000 to $9,999 to other University Priorities

Frank D. Catrickes
Eric T. Lee
Michael A. Pollack
Paula Jean Poskon

Gifts of $2,500 to $4,999 to The Penn Fund

Gordon L. Austin
Melissa Friedman Birns *^
Alex Campbell ^
David N. Corleto
Kevin C. Davis *^
Dana B. Klein *^
Amy Kroll Cardillo *^•
Melissa J. Lau
Darren R. Levy *^
Joseph P. Melchiors *^
Youlia K. Miteva *^
Alan K. Miyasaki ^•
Soon H. Pho

Gifts of $2,500 to $4,999 to other University Priorities

Alyse Dann Bodine
Alex Campbell
Lisa K. Forman
Lauren B. Hertzog
Ashley Damron Mohan
Vijai P. Mohan
Arihant Patni
Soon H. Pho
Justin Warren Slatky
Joseph S. Weintraub

Gifts of $1,000 to $2,499 to The Penn Fund

John S. Adractas *^
Michael E. Bogdan *^
Frank D. Catrickes *^
Rebecca D. Flax *^
Lauren Schlenoff Kowal *^
Selma Kwok Lawrence ^
James J. Lees, Jr. *^
Hamilton S. Miller *^
Lisa Ostad *^
Joseph B. Pollack *^
Whitney Namm Pollack *^•
Melissa F. Rice ^
Bryan S. Romano *^
Christopher K. Scarborough ^
Manisha Ahuja Sethi *^
Amanda B. Siegel ^
Daniel Braun Silvers *^•
Monica Scott Tavares ^
Eunice M. Trieu ^

Gifts of $1,000 to $2,499 to other University Priorities

Dafna M. Aaronson
Amir Becher
Thaddeus J. Berdzik
Paul J. Burg
Michael S. Burka
Beth Fluke
Marion Ming-Jung Hsieh
Tony Y. Hsu
Daniel A. Lee
Maria Magliacano
Joseph P. Melchiors
Sheila B. Ong

** Donor to The Penn Fund for five or more consecutive years*
^ Ivy Stone Society member
ⲷ Harrison Society member
• Proud Penn Fund volunteer

1

Note.— This figure provides an example of the Honor Roll for Fiscal Year 2010. Donors are recognized in the appropriate support category and can appear at most twice – once for any donation directly to the Penn Fund and once for any donation indirectly to other University priorities. Recognition for consecutive giving is denoted by a symbol next to the donor's name.

Fig. 5.—Timeline of Recognition Programs

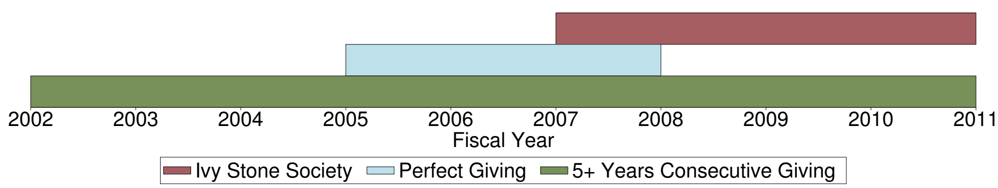

Note.— This figure illustrates the relative timing of three programs recognizing consecutive giving to the Penn Fund during our sample period (2002-2011). The program recognizing five or more years of consecutive giving was in effect during the entirety of our sample period, while the Ivy Stone Society and Perfect Giving programs were introduced midway through the sample period.

Fig. 6.—Donation Frequency for Amounts between \$200 and \$5000

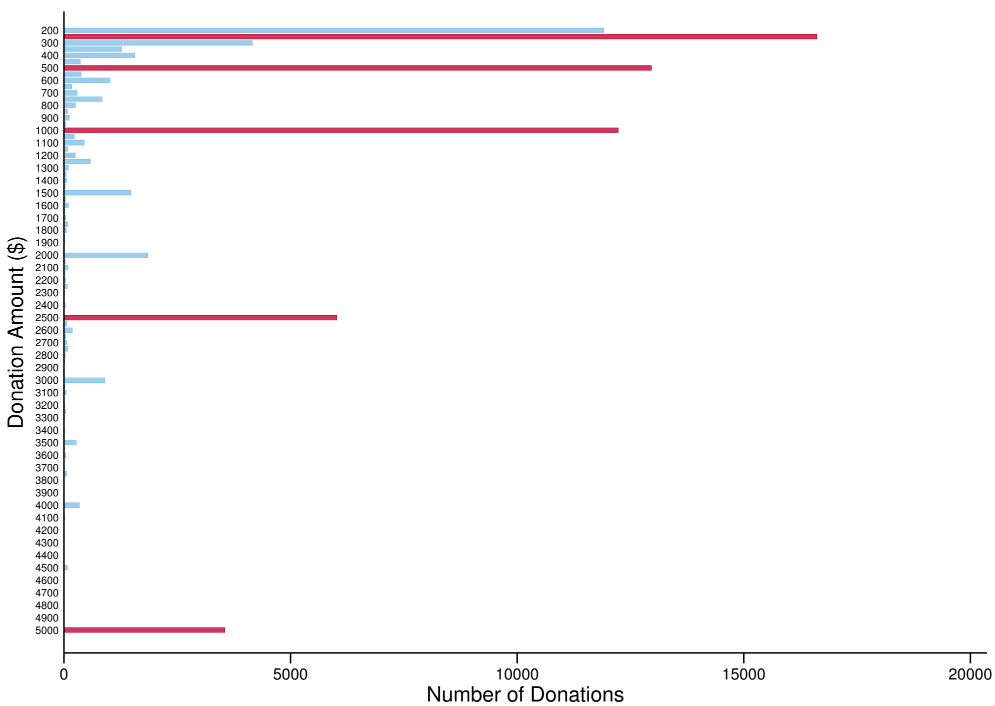

Note.— This figure shows the frequency of non-zero donation amounts (in nominal dollars) between \$200 and \$5000 in increments of 50 dollars in the data. Together, the donations represented in the figure account for over 92.7 percent of non-zero donations made between \$200 and \$5000. Donations under consideration include those made to the Penn Fund and those made to other University priorities. Donation amounts are assigned different colors depending on whether they correspond to the lower bound of a support category (red) or not (blue).

FIG. 7.—Eligibility of Alumni by Recognition Programs

Years of Consecutive Giving at t-1 \ Years Since Graduation at t	1	2	3	4	5	6	7…
0							
1	Perfect Giving ISS						
2		Perfect Giving ISS	ISS	ISS	ISS	ISS	ISS
3			Perfect Giving ISS	ISS	ISS	ISS	ISS
4				Perfect Giving ISS 5 Consec. Years	ISS 5 Consec. Years	ISS 5 Consec. Years	ISS 5 Consec. Years
5					ISS 5 Consec. Years	ISS 5 Consec. Years	ISS 5 Consec. Years
6						ISS 5 Consec. Years	ISS 5 Consec. Years
7…							ISS 5 Consec. Years

NOTE.— This figure illustrates which alumni were eligible for each of the three programs recognizing consecutive giving: PG, ISS, and 5YCG. The rows give the number years of consecutive giving up until but not including the current year. The columns give the number of years since graduation.

FIG. 8.—Mean Residuals for Giving over Time for Eligible and Ineligible

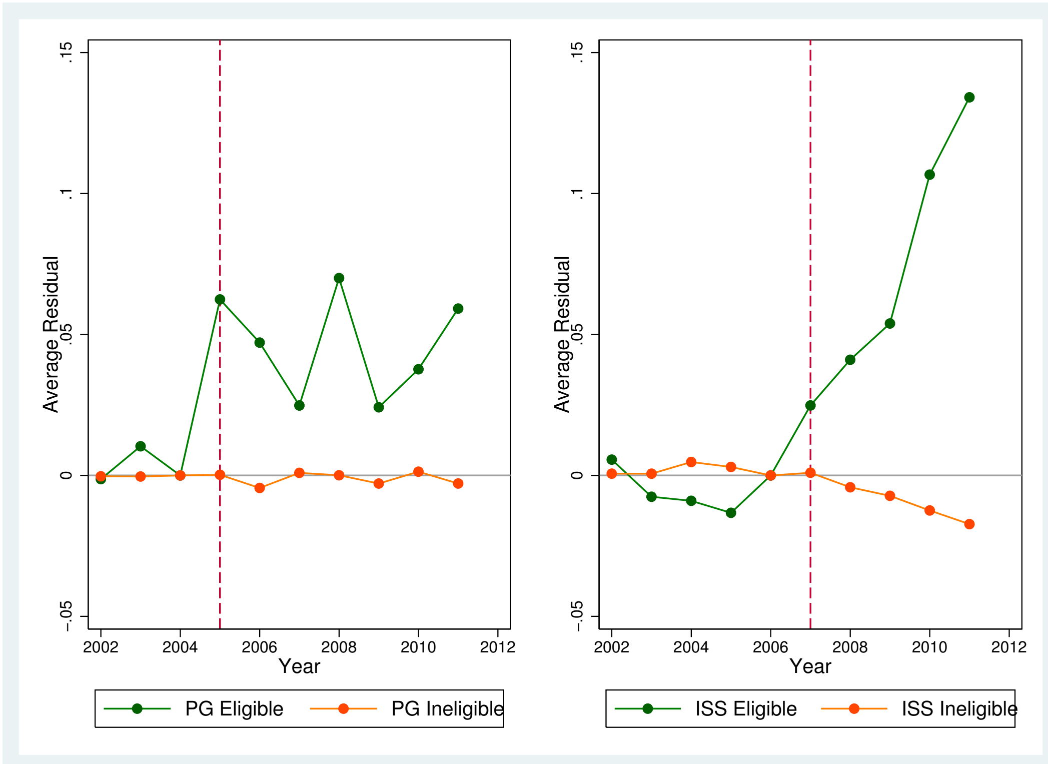

NOTE.—This figure plots the mean residuals over time from estimating giving directly to the Penn Fund for both the PG and ISS programs by eligibility. For each recognition program, residuals are centered at zero in the year prior to the enactment of the program. The PG sample includes only young alumni; the ISS sample includes only alumni 7+ years out since graduation.

FIG. 9.—Effect of Introduction of Ivy Stone Society by Years of Consecutive Giving

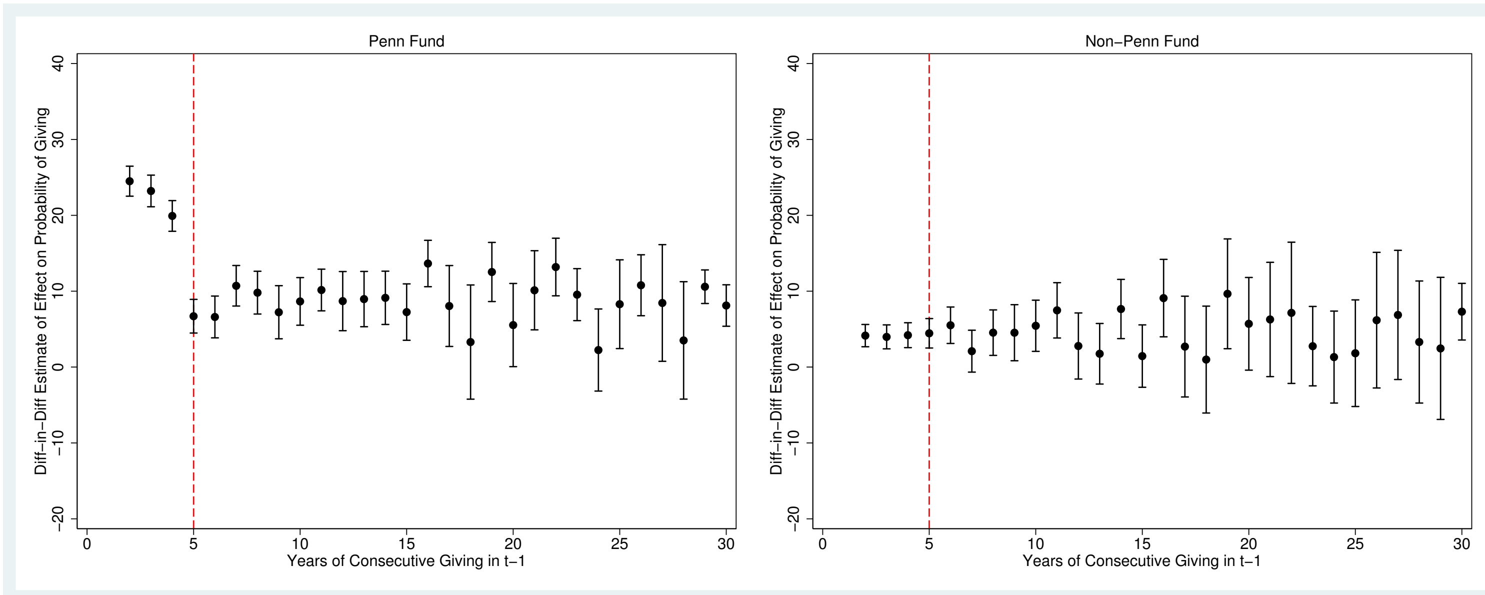

NOTE.—This figure reports estimates of the marginal effect of the ISS program on the probability of giving by the number of years of consecutive giving. The estimation sample used in this figure is the set of donations made by alumni 7 or more years out since graduation between FY02 and FY07. Alumni with 5 or more years of consecutive giving already receive recognition under the 5YCG program.

TABLE 16
Summary Statistics

	Mean	SD
Male	0.57	(0.49)
Married	0.49	(0.50)
Years Since Graduation	27.54	(17.54)
Donations Made to the PennFund†		
Donation ($)	294.22	(8,176.75)
Donation Conditional on Giving ($)	1,723.64	(19,728.65)
Log(Donation) ($)	0.86	(2.01)
Log(Donation) Conditional on Giving ($)	5.04	(1.64)
Giving	0.17	(0.38)
Bunching Conditional on Giving	0.11	(0.31)
Donations Made to Other University Priorities†		
Donation ($)	1,002.42	(53,794.81)
Donation Conditional on Giving ($)	11,759.74	(183,910.39)
Log(Donation) ($)	0.45	(1.60)
Log(Donation) Conditional on Giving ($)	5.31	(2.09)
Giving	0.09	(0.28)
Bunching Conditional on Giving	0.07	(0.26)
N	2,149,234	
Number of Alumni	117937	

Note.—Table entries are means and standard deviations. Observations are on the individual-year-designation level where the designation distinguishes between donations made to the PennFund and donations made to other University priorities. The estimation sample is the set of donations made by undergraduate alumni (one or more years since graduation) during the period 2002-2011, excluding those alumni who are deceased and excluding any donations made by alumni prior to graduating. If no donations were made, the alumni is included as having made a zero dollar donation. All monetary values are in 2010 U.S. dollars. Summary statistics for gender and marital status are calculated at the alumni level. All other statistics are calculated at the observation level. Summary statistics for gender exclude 46 individuals who do not identify as either male or female. Marital status in our sample is time-invariant and represents the marital status of that alumnus at the time the data was extracted by the University. †For each year, alumni appear in the data twice – once for their donation to the Penn Fund and once for their (Non-Penn Fund) donation to other University priorities. Accordingly, donations made to the Penn Fund are half of total observations ($N = 1,074,617$); likewise, donations made to other University priorities are half of total observations ($N = 1,074,617$).

TABLE 17
CATEGORIES OF SUPPORT ($) BY YEARS SINCE GRADUATION

Penn Fund Donations			Non-Penn Fund Donations
$1-4$	$5-9$	$10+$	All Years
$1-249$	$1-499$		
$250-999$	$500-999$	$1-999$	$1-999$
$1,000-2,499$	$1,000-2,499$	$1,000-2,499$	$1,000-2,499$
$2,500-4,999$	$2,500-4,999$	$2,500-4,999$	$2,500-4,999$
$5,000-9,999$	$5,000-9,999$	$5,000-9,999$	$5,000-9,999$
$10,000-24,999$	$10,000-24,999$	$10,000-24,999$	$10,000-24,999$
$\geq 25,000$	$\geq 25,000$	$\geq 25,000^{\dagger}$	$\geq 25,000^{\dagger}$

NOTE.—Table provides the stratification levels of donation (categories of support) by designation type for recognition in the Honor Roll during 2002-2011. Columns 1-3 lists categories for donations made to the Penn Fund. Column 1 lists categories for alumni who graduated one to four years prior (young alumni); Column 2 lists categories for alumni who graduated five to nine years prior; and Column 3 lists categories for alumni who graduated ten or more years ago. Column 4 lists categories for (Non-Penn Fund) donations made to other University priorities for all alumni regardless of years since graduation. The dollar amounts bracketing each category of support are inclusive. †For 25th and 50th reunion years, the highest category of support is instead stratified into two categories: $\$25,000-\$99,999$ and $\$100,000+$.

TABLE 18
Donations at Lower Bound (Bunching)

Categories of Support ($)	Percentage Bunching	Observations
$1-999^{\dagger}$	0.05	204,002
$1-249$	0.23	8,802
$250-999$	68.05	1,105
$1-499$	0.08	13,632
$500-999$	85.36	1,059
$1,000-2,499$	62.21	19,657
$2,500-4,999$	63.45	9,485
$5,000-9,999$	60.31	5,876
$10,000-24,999$	27.70	5,459
$\geq 25,000^{\ddagger}$	13.03	5,950
N		275,027

Note.—Table reports two statistics: (i) the percentage of donations (both Penn Fund and Non-Penn Fund) within each category of support that were made exactly at the category's lower bound (percentage bunching) and (ii) the total number of non-zero donations within each category. The total number of non-zero donations across all categories ($N = 275,027$) excludes 10 non-zero donations which were less than $1. Statistics were calculated using nominal donation amounts. †The categories of support for donations below $1000 depend on the number of years since graduation and the designation type of the donation (see Table 17). Statistics for Row 1 are for Penn Fund donations made by alumni ten or more years since graduation or for all Non-Penn Fund donations within the category bounds. Statistics for Rows 2 and 3 are for Penn Fund donations within the category bounds made by alumni between one and four years since graduation. Statistics for Rows 4 and 5 are for Penn Fund donations within the category bounds made by alumni between five and nine years since graduation. ‡For 25th and 50th reunion years, the highest category of support is instead stratified into two categories. In this table, however, we only show bunching and the number of non-zero donations for the category $25,000+$.

TABLE 19

Summary Statistics by Ivy Stone Society (ISS) Eligibility

	Eligible		Not Eligible	
	Mean	SD	Mean	SD
Male	0.59	(0.49)	0.57	(0.49)
Married	0.72	(0.45)	0.46	(0.50)
Years Since Graduation	29.54	(16.03)	27.28	(17.72)
Donations Made to the PennFund				
Donation ($)	1,701.07	(20,378.09)	107.46	(4,506.79)
Donation Conditional on Giving ($)	2,105.97	(22,655.22)	1,247.66	(15,309.99)
Log(Donation) ($)	4.26	(2.58)	0.41	(1.40)
Log(Donation) Conditional on Giving ($)	5.28	(1.70)	4.74	(1.52)
Giving	0.81	(0.39)	0.09	(0.28)
Bunching Conditional on Giving	0.12	(0.33)	0.08	(0.27)
Donations Made to Other University Priorities				
Donation ($)	2,144.94	(55,286.02)	850.75	(53,591.93)
Donation Conditional on Giving ($)	7,796.57	(105,196.37)	14,170.63	(218,292.34)
Log(Donation) ($)	1.42	(2.53)	0.32	(1.38)
Log(Donation) Conditional on Giving ($)	5.18	(1.97)	5.38	(2.15)
Giving	0.28	(0.45)	0.06	(0.24)
Bunching Conditional on Giving	0.05	(0.23)	0.08	(0.27)
N	251,874		1,897,360	

Note.—Table entries are means and standard deviations. Observations are on the individual-year-designation level. The estimation sample is the set of donations made by undergraduate alumni (one or more years since graduation) during the period 2002-2011. If no donations were made, the alumni is included as having made a zero dollar donation. All monetary values are in 2010 U.S. dollars. Columns 1 and 2 provide statistics for giving opportunities where an individual was considered "ISS eligible" ($N = 251,874$); and Columns 3 and 4 provide statistics for giving opportunities where an individual was not considered ISS eligible ($N = 1,897,360$). Summary statistics for gender and marital status are calculated at the alumni level. All other statistics are calculated at the observation level.

TABLE 20

SUMMARY STATISTICS BY PERFECT GIVING (PG) ELIGIBILITY

	Eligible		Not Eligible (YPA)†		Not Eligible	
	Mean	SD	Mean	SD	Mean	SD
Male	0.47	(0.50)	0.50	(0.50)	0.57	(0.49)
Married	0.22	(0.41)	0.15	(0.36)	0.49	(0.50)
Years Since Graduation	2.73	(0.79)	3.02	(0.82)	27.65	(17.50)
Donations Made to the PennFund						
Donation ($)	79.00	(170.81)	10.78	(224.16)	295.16	(8,194.47)
Donation Conditional on Giving ($)	102.54	(188.30)	109.34	(706.49)	1,755.90	(19,922.68)
Log(Donation) ($)	3.04	(1.93)	0.38	(1.19)	0.85	(2.01)
Log(Donation) Conditional on Giving ($)	3.94	(1.13)	3.81	(1.07)	5.06	(1.64)
Giving	0.77	(0.42)	0.10	(0.30)	0.17	(0.37)
Bunching Conditional on Giving	0.12	(0.32)	0.07	(0.25)	0.11	(0.31)
Donations Made to Other University Priorities						
Donation ($)	19.08	(234.80)	37.13	(1,886.31)	1,006.69	(53,911.46)
Donation Conditional on Giving ($)	191.10	(721.28)	1,262.64	(10,932.71)	11,818.64	(184,376.10)
Log(Donation) ($)	0.42	(1.31)	0.13	(0.80)	0.45	(1.60)
Log(Donation) Conditional on Giving ($)	4.17	(1.20)	4.49	(1.50)	5.31	(2.09)
Giving	0.10	(0.30)	0.03	(0.17)	0.09	(0.28)
Bunching Conditional on Giving	0.02	(0.14)	0.04	(0.19)	0.07	(0.26)
N	9,294		130,324		2,139,940	

NOTE.—Table entries are means and standard deviations. Observations are on the individual-year-designation level. The estimation sample is the set of donations made by undergraduate alumni during the period 2002-2011. All monetary values are in 2010 U.S. dollars. Columns 1 and 2 provide statistics for giving opportunities where an individual was considered "PG eligible" ($N = 9,294$); Columns 3 and 4 provide statistics for giving opportunities where an individual was a young penn alumnus and was not considered PG eligible ($N = 130,324$); and Columns 5 and 6 provide statistics for giving opportunities where an individual (all years) was not considered PG eligible ($N = 2,139,940$). Summary statistics for gender and marital status are calculated at the alumni level. All other statistics are calculated at the observation level. ‡We provide statistics for both the set of giving opportunities for which all alumni are not PG eligible and the subset of giving opportunities for which only young Penn alumni are not eligible.

TABLE 21

EFFECTS OF PG RECOGNITION PROGRAM ON GIVING WITH NO DYNAMICS - 2002 TO 2005

	Direct Effect		Crowd In Effect	
	Giving	Log(Donation)	Giving	Log(Donation)
PGELIGIBLE	-12.768**	-0.509**	-1.460	-0.084
	(5.413)	(0.234)	(2.323)	(0.096)
(PGINEFFECT)*(PGELIGIBLE)	11.730***	0.412***	8.172***	0.409***
	(3.289)	(0.146)	(2.506)	(0.122)
R-squared	0.145	0.111	0.184	0.158
N	27,875	27,875	27,875	27,875
Percent Years Since Grad Gave	Y	Y	Y	Y
Years Since Graduation FEs	Y	Y	Y	Y
Gift Fiscal Year FEs	Y	Y	Y	Y
Individual FEs	Y	Y	Y	Y

NOTE.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made by young alumni between FY02 and FY05. Column 1 reports the marginal effects of the PG recognition program on the probability of giving to the Penn Fund. Column 2 reports the same marginal effects on the log donation amount to the Penn Fund. Columns 3 and 4 reports estimates analogous to Columns 1 and 2 using instead donations to other University priorities. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. "Percent years since grad gave" is the fraction of years since graduation for which an individual has given.

TABLE 22
EFFECTS OF PG RECOGNITION PROGRAM ON GIVING WITH NO DYNAMICS - 2004 TO 2005

	Direct Effect		Crowd In Effect	
	Giving	Log(Donation)	Giving	Log(Donation)
PGELIGIBLE	-5.638	-0.190	4.025*	0.141
	(7.189)	(0.279)	(2.350)	(0.087)
(PGINEFFECT)*(PGELIGIBLE)	21.736***	0.709***	7.741***	0.346***
	(3.895)	(0.166)	(2.403)	(0.116)
R-squared	0.292	0.213	0.285	0.214
N	13,944	13,944	13,944	13,944
Percent Years Since Grad Gave	Y	Y	Y	Y
Years Since Graduation FEs	Y	Y	Y	Y
Gift Fiscal Year FEs	Y	Y	Y	Y
Individual FEs	Y	Y	Y	Y

NOTE.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made by young alumni between FY04 and FY05. Column 1 reports the marginal effects of the PG recognition program on the probability of giving to the Penn Fund. Column 2 reports the same marginal effects on the log donation amount to the Penn Fund. Columns 3 and 4 reports estimates analogous to Columns 1 and 2 using instead donations to other University priorities. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. "Percent years since grad gave" is the fraction of years since graduation for which an individual has given.

TABLE 23

EFFECTS OF ISS RECOGNITION PROGRAM ON GIVING WITH NO DYNAMICS - 2002 TO 2007

	Direct Effect		Crowd In Effect	
	Giving	Log(Donation)	Giving	Log(Donation)
ISSELIGIBLE	1.909***	0.122***	1.825***	0.096***
	(0.357)	(0.019)	(0.238)	(0.013)
(ISSINEFFECT)*(ISSELIGIBLE)	15.533***	0.730***	4.410***	0.243***
	(0.459)	(0.025)	(0.324)	(0.018)
R-squared	0.054	0.048	0.046	0.035
N	546,873	546,873	546,873	546,873
Percent Years Since Grad Gave	Y	Y	Y	Y
Years Since Graduation FEs	Y	Y	Y	Y
Gift Fiscal Year FEs	Y	Y	Y	Y
Individual FEs	Y	Y	Y	Y

NOTE.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made by alumni 7 or more years out since graduation between FY02 and FY07. Column 1 reports the marginal effects of the ISS recognition program on the probability of giving to the Penn Fund. Column 2 reports the same marginal effects on the log donation amount to the Penn Fund. Columns 3 and 4 reports estimates analogous to Columns 1 and 2 using instead donations to other University priorities. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. "Percent years since grad gave" is the fraction of years since graduation for which an individual has given.

TABLE 24

EFFECTS OF ISS RECOGNITION PROGRAM ON GIVING WITH NO DYNAMICS - 2006 TO 2007

	Direct Effect		Crowd In Effect	
	Giving	Log(Donation)	Giving	Log(Donation)
ISSELIGIBLE	-24.642***	-1.219***	-2.925***	-0.124***
	(0.755)	(0.041)	(0.508)	(0.028)
(ISSINEFFECT)*(ISSELIGIBLE)	30.616***	1.363***	7.379***	0.378***
	(0.835)	(0.045)	(0.352)	(0.019)
R-squared	0.268	0.240	0.230	0.192
N	191,619	191,619	191,619	191,619
Percent Years Since Grad Gave	Y	Y	Y	Y
Years Since Graduation FEs	Y	Y	Y	Y
Gift Fiscal Year FEs	Y	Y	Y	Y
Individual FEs	Y	Y	Y	Y

NOTE.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made by alumni 7 or more years out since graduation between FY06 and FY07. Column 1 reports the marginal effects of the ISS recognition program on the probability of giving to the Penn Fund. Column 2 reports the same marginal effects on the log donation amount to the Penn Fund. Columns 3 and 4 reports estimates analogous to Columns 1 and 2 using instead donations to other University priorities. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. "Percent years since grad gave" is the fraction of years since graduation for which an individual has given.

TABLE 25

Effects of ISS and PG Recognition Programs on Giving Accounting for Dynamics - Young Alumni

	Direct Effect		Crowd In Effect	
	Giving	Log(Donation)	Giving	Log(Donation)
ISSEligibleInOneYr	0.376	0.036	1.034*	0.049*
	(1.330)	(0.056)	(0.626)	(0.030)
PGEligible	-9.491***	-0.419***	-5.364***	-0.259***
	(3.084)	(0.130)	(1.791)	(0.077)
ISSNoPGEligible	0.091	0.075	4.430***	0.174***
	(2.548)	(0.111)	(1.334)	(0.059)
(PGOrISSInEffect)*(PGEligible)	11.723***	0.513***	4.997***	0.222***
	(2.703)	(0.116)	(1.806)	(0.080)
(ISSInEffect)*(ISSEligibleInOneYR)	5.759***	0.265***	0.280	0.007
	(1.730)	(0.071)	(0.846)	(0.040)
(ISSInEffect)*(ISSNoPGEligible)	7.558**	0.274**	0.564	0.043
	(2.977)	(0.131)	(1.827)	(0.088)
R-squared	0.135	0.100	0.188	0.154
N	69,809	69,809	69,809	69,809
Percent Years Since Grad Gave	Y	Y	Y	Y
Years Since Graduation FEs	Y	Y	Y	Y
Gift Fiscal Year FEs	Y	Y	Y	Y
Individual FEs	Y	Y	Y	Y

Note.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made by young alumni between FY02 and FY11. Column 1 reports the marginal effects of both the PG and ISS recognition programs on the probability of giving to the Penn Fund. Estimates are relative to individuals who have not given for either of the past two years. Column 2 reports the same marginal effects on the log donation amount to the Penn Fund. Columns 3 and 4 reports estimates analogous to Columns 1 and 2 using instead donations to other University priorities. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. "Percent years since grad gave" is the fraction of years since graduation for which an individual has given.

TABLE 26

EFFECTS OF ISS RECOGNITION PROGRAM ON GIVING ACCOUNTING FOR DYNAMICS - NON YOUNG ALUMNI

	Direct Effect		Crowd In Effect	
	Giving	Log(Donation)	Giving	Log(Donation)
ISSEligibleInOneYr	9.793***	0.488***	0.371*	-0.004
	(0.298)	(0.015)	(0.194)	(0.010)
ISSEligible	12.742***	0.701***	1.653***	0.058***
	(0.315)	(0.017)	(0.228)	(0.013)
(ISSInEffect)*(ISSEligibleInOneYR)	10.910***	0.542***	3.059***	0.176***
	(0.410)	(0.021)	(0.264)	(0.014)
(ISSInEffect)*(ISSNoPGEligible)	22.483***	1.096***	5.171***	0.281***
	(0.353)	(0.019)	(0.248)	(0.014)
R-squared	0.067	0.068	0.015	0.011
N	1,004,808	1,004,808	1,004,808	1,004,808
Percent Years Since Grad Gave	Y	Y	Y	Y
Years Since Graduation FEs	Y	Y	Y	Y
Gift Fiscal Year FEs	Y	Y	Y	Y
Individual FEs	Y	Y	Y	Y

NOTE.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made by alumni five or more years out since graduation between FY02 and FY11. Column 1 reports the marginal effects of the ISS recognition program on the probability of giving to the Penn Fund. Estimates are relative to individuals who have not given for either of the past two years. Column 2 reports the same marginal effects on the log donation amount to the Penn Fund. Columns 3 and 4 reports estimates analogous to Columns 1 and 2 using instead donations to other University priorities. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. "Percent years since grad gave" is the fraction of years since graduation for which an individual has given.

TABLE 27
Effects of PG Recognition Program on Bunching with No Dynamics

	Bunching	Bunch Same	Bunch Up	Bunch Down
PGEligible	-2.201	1.391	-1.253	-2.339
	(8.025)	(6.089)	(5.501)	(1.603)
(PGInEffect)*(PGEligible)	6.107	2.139	2.002	1.966
	(4.696)	(2.936)	(3.696)	(1.818)
Additional Controls:				
BunchedLastYear	-29.938***	35.226***	-57.573***	-7.591*
	(9.576)	(7.853)	(8.439)	(4.164)
R-squared	0.070	0.188	0.288	0.045
N	2,688	2,688	2,688	2,688
Percent Years Since Grad Gave	Y	Y	Y	Y
Years Since Graduation FEs	Y	Y	Y	Y
Gift Fiscal Year FEs	Y	Y	Y	Y
Individual FEs	Y	Y	Y	Y

Note.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made directly to the Penn Fund by young alumni between FY02 and FY05. All estimates are conditional on giving in the prior year. Column 1 reports the conditional marginal effects of the PG recognition program on the probability of bunching. Column 2 reports the conditional marginal effects of the PG program on the probability of bunching and giving the same donation as the prior year. Columns 3 reports conditional effects on the probability of bunching and giving a strictly larger donation than the prior year. Column 4 reports conditional effects on the probability of bunching and giving a strictly smaller donation than the prior year. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. “Percent years since grad gave” is the fraction of years since graduation for which an individual has given.

TABLE 28

Effects of ISS Recognition Program on Bunching with No Dynamics

	Bunching	Bunch Same	Bunch Up	Bunch Down
ISSEligible	-0.581**	-0.018	-0.328*	-0.235*
	(0.235)	(0.161)	(0.177)	(0.123)
(ISSInEffect)*(ISSEligible)	3.571***	1.386***	1.010**	1.175***
	(0.616)	(0.423)	(0.439)	(0.280)
Additional Controls:				
BunchedLastYear	3.442***	32.714***	-17.440***	-11.831***
	(0.885)	(0.678)	(0.610)	(0.585)
R-squared	0.011	0.170	0.064	0.057
N	94,035	94,035	94,035	94,035
Percent Years Since Grad Gave	Y	Y	Y	Y
Years Since Graduation FEs	Y	Y	Y	Y
Gift Fiscal Year FEs	Y	Y	Y	Y
Individual FEs	Y	Y	Y	Y

Note.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made directly to the Penn Fund by alumni 7 or more years out since graduation between FY02 and FY07. All estimates are conditional on giving in the prior year. Column 1 reports the conditional marginal effects of the ISS recognition program on the probability of bunching. Column 2 reports the conditional marginal effects of the ISS program on the probability of bunching and giving the same donation as the prior year. Columns 3 reports conditional effects on the probability of bunching and giving a strictly larger donation than the prior year. Column 4 reports conditional effects on the probability of bunching and giving a strictly smaller donation than the prior year. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. "Percent years since grad gave" is the fraction of years since graduation for which an individual has given.

TABLE 29

Effects of ISS and PG Recognition Programs on Bunching Accounting for Dynamics - Young Alumni

	Bunching	Bunch Same	Bunch Up	Bunch Down
PGEligible	0.702	-2.143	3.955	-1.110
	(6.381)	(3.971)	(4.354)	(2.599)
ISSNoPGEligible	2.203	2.055	1.303	-1.155
	(2.109)	(1.755)	(1.464)	(0.889)
(PGOrISSInEffect)*(PGEligible)	7.037*	3.961	0.399	2.677**
	(3.847)	(2.688)	(2.753)	(1.355)
(ISSInEffect)*(ISSNoPGEligible)	-1.877	0.332	-2.700	0.491
	(2.188)	(1.559)	(1.683)	(0.930)
Additional Controls:				
BunchedLastYear	-19.973***	34.262***	-34.534***	-19.701***
	(3.296)	(2.778)	(2.870)	(2.513)
R-squared	0.046	0.175	0.202	0.134
N	9,132	9,132	9,132	9,132
Percent Years Since Grad Gave	Y	Y	Y	Y
Years Since Graduation FEs	Y	Y	Y	Y
Gift Fiscal Year FEs	Y	Y	Y	Y
Individual FEs	Y	Y	Y	Y

Note.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made by young alumni directly to the Penn Fund between FY02 and FY11. All estimates are conditional on giving in the prior year. Column 1 reports the marginal effects of both the PG and ISS recognition programs on the probability of bunching. Estimates are relative to individuals who gave in the previous year but not two years previously. Column 2 reports the conditional marginal effects of the PG and ISS programs on the probability of bunching and giving the same donation as the prior year. Columns 3 reports conditional effects on the probability of bunching and giving a strictly larger donation than the prior year. Column 4 reports conditional effects on the probability of bunching and giving a strictly smaller donation than the prior year. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. “Percent years since grad gave” is the fraction of years since graduation for which an individual has given.

TABLE 30

EFFECTS OF ISS RECOGNITION PROGRAM ON BUNCHING ACCOUNTING FOR DYNAMICS - NON YOUNG ALUMNI

	Bunching	Bunch Same	Bunch Up	Bunch Down
ISSELIGIBLE	-0.771***	-0.159	-0.286*	-0.326***
	(0.202)	(0.141)	(0.149)	(0.103)
(ISSINEFFECT)*(ISSNOPGELIGIBLE)	1.070***	0.492**	0.405**	0.173
	(0.281)	(0.203)	(0.205)	(0.150)
Additional Controls:				
BUNCHEDLASTYEAR	14.163***	34.282***	-12.213***	-7.905***
	(0.613)	(0.488)	(0.368)	(0.363)
R-squared	0.028	0.196	0.044	0.032
N	171,742	171,742	171,742	171,742
Percent Years Since Grad Gave	Y	Y	Y	Y
Years Since Graduation FEs	Y	Y	Y	Y
Gift Fiscal Year FEs	Y	Y	Y	Y
Individual FEs	Y	Y	Y	Y

NOTE.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made by alumni five or more years out since graduation directly to the Penn Fund between FY02 and FY11. All estimates are conditional on giving in the prior year. Column 1 reports the marginal effect of the ISS recognition program on the probability of bunching. Estimates are relative to individuals who gave in the previous year but not two years previously. Column 2 reports the conditional marginal effects of the ISS program on the probability of bunching and giving the same donation as the prior year. Columns 3 reports conditional effect on the probability of bunching and giving a strictly larger donation than the prior year. Column 4 reports conditional effect on the probability of bunching and giving a strictly smaller donation than the prior year. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. “Percent years since grad gave” is the fraction of years since graduation for which an individual has given.

TABLE 31

Effect of Introduction of Ivy Stone Society by Years of Consecutive Giving

	Direct Effect	Crowd In Effect
	Giving	Giving
ISSEligible	0.498	1.878***
	(0.369)	(0.246)
(ISSInEffect)*(ISSEligible)*(<5YearsConsecGiving)	22.823***	4.111***
	(0.650)	(0.478)
(ISSInEffect)*(ISSEligible)*(5+YearsConsecGiving)	8.374***	4.657***
	(0.515)	(0.429)
R-squared	0.056	0.046
N	546,873	546,873
Percent Years Since Grad Gave	Y	Y
Years Since Graduation FEs	Y	Y
Gift Fiscal Year FEs	Y	Y
Individual FEs	Y	Y

Note.—Each column reports estimates from an OLS regression run at the individual-year-designation level. Standard errors are reported in parentheses and are clustered by individual. The dependent variable in each specification is listed at the head of each column. The estimation sample used in this table is the set of donations made by alumni 7 or more years out since graduation between FY02 and FY07. Column 1 reports the marginal effects of the ISS recognition program on the probability of giving directly to the Penn Fund. Column 2 reports the same marginal effect on the probability of giving indirectly to other University priorities. We allow the marginal effects of the ISS recognition program to vary by whether an ISS eligible individual has given consecutively for the past five years. All specifications include individual fixed effects as well as years since graduation and fiscal year fixed effects. "Percent years since grad gave" is the fraction of years since graduation for which an individual has given.

CHAPTER 3 : Play for Performance

3.1. Introduction

Economists have long recognized that the provision of monetary incentives can lead to increased effort. Pecuniary rewards have been used to motivate student achievement (Angrist and Lavy 2009; Angrist et al. 2009; Fryer 2011; Bettinger 2012) and engagement in healthy behaviors (Charness and Gneezy 2009; Volpp et al. 2009; Acland and Levy 2013). Far less studied is the role of non-pecuniary incentives, or intrinsic motivation, in motivating effort. Deci (1971) defines a person as intrinsically motivated if he is motivated to "...perform an activity when he receives no apparent rewards except the activity itself."

The economics literature on intrinsic motivation has primarily focused on the interplay between intrinsic and extrinsic motivation, with a large body of research dedicated to assessing whether extrinsic rewards or benefits (i.e. extrinsic motivation) may undermine (crowd out) intrinsic motivation (Ryan 1982; Frey 1992, 1997; Bnabou and Tirole 2003; Ryan and Deci 2000; Deci et al. 2001; Gneezy et al. 2011). Past research focusing on intrinsic motivation is limited in part because of the challenges researchers face in identifying relevant settings where possible correlation between pecuniary and non-pecuniary incentives does not exist.

In this paper, we attempt to isolate the role of past performance as a source of intrinsic motivation using a unique setting: an online word-hunting game. Players exert effort in the game by identifying words from grids of letters (game boards) in order to attain points. The objective of the game is to score as many points as possible within a fixed time frame. At the end of each game, a player must decide whether or not to continue playing. The online word-hunting game offers three important advantages for studying the role of past performance as a source of intrinsic motivation. First, performance is directly measurable using the number of points scored by players in each game and is easily observable to players in the game. Second, we are able to use an unusually rich data set of information on both features of each game played over our sample period and player performance data, both individually and as part of a team. In addition, players must register a unique identifier in order to play, allowing us to observe a players pattern of play over time. Third, players are randomly assigned game boards that vary in difficulty, allowing us to exploit quasi-random variation in performance. This random assignment proves key in our identification strategy.

We are interested in estimating the effect of past performance on the persistence of play. However, there are concerns over endogeneity and omitted variables. For example, suppose that individuals with large vocabularies also enjoy playing online games. Then any estimated effect on the persistence of play may be mistakenly attributed to past performance when in reality, it is a result of unobserved preferences. To address this issue, we take advantage of individual variation in the ability to exploit the presence of prefixes and suffixes in game boards. Specifically, we instrument for measures of past performance using number of possible board points that can be scored for each prefix and suffix. Because a random game board is assigned each round, the instruments are necessarily exogenous.

Using the identification strategy outline above, we find that an increase in past performance measures, for both individual players and teams, significantly increases the number of games individuals play consecutively (spell length) and the probability that a spell will end. The one exception is that an increase in the absolute number of points scored significantly decreases spell length and increases the probability that a spell will end. We test to see whether these effects might be explained by players engaging in satisficing behavior there exists some satisficing level of success (e.g. a top ten ranking), attainment of which would induce the player to quit playing. We find no evidence of players engaging in satisficing behavior. Finally, we look at player performance over time and find suggestive evidence that players are learning as they play additional games.

The remainder of the paper proceeds as follows. In 3.2, we provide additional background on the online game and introduce the data. 3.3 describes the relationship between performance and game board features. 3.4 describes our estimation strategy and estimates a relationship between past performance and the persistence of play, using game board features as an instrument for past performance. 3.5 discusses the implications of the results we find and concludes.

3.2. Game Background and Data

3.2.1. Game Play

We study intrinsic motivation using data collected from Wordsplay.com, a free online implementation of the word-hunting game Boggle. Each round of play (a "game") features a randomly generated grid of letters (see 10) and lasts for three minutes. Players score points by finding words in sequences of adjoining letters where no letter is used more than once in a given word and consecutive letters in the word must be adjacent on the game board. The number of points awarded for each word is a non-linear increasing function of the length of the word.

Players have the option to join a team, which will be scored and ranked alongside the individual players. At the end of a round, each player and team receives a final score for the game equal to the cumulative total of points scored for each word that was found. The same word found by two or more players on a team cannot be counted more than once towards the final score for the team, though the word does contribute for the individual players' scores. Players and teams are then ranked together so that for each game, an individual is able to determine the number of points scored, the fraction of total possible points scored, his or her individual ranking, and the ranking of his or her team (conditional on being part of a team). The winner for a game is the player or team with the highest final score (number one ranking). Once a round is complete, players must wait an additional 30 seconds before the next round begins with a new random game board.

3.2.2. Data

Our data provides player-game level information for all games played on Wordsplay.com in 2009. For each observation, the data includes game-specific information on all possible words and possible points for the associated game board as well as player-specific information on the actual points scored by the player and the points scored by the players team (conditional on being on a team). Games are indexed by a timestamp for when the round was played. The data allow us to link games over time for a given individual using a unique identifier that must be created by individuals in order to play. Similarly, a unique team identifier allows us to link players in the same team.

Table 32 presents summary statistics for the 24,433 players in our sample. Players play on average 454 games over the course of our sample period; however, as Figure 11 shows, this distribution is highly skewed with the median number of games played being 16. The distribution of the number of games played also differs by whether the player is participating on a team during the given round with the mean and median number of games played as part of a team far lower than when played individually. This difference is driven primarily by the large fraction of players who never play on a team (i.e. play zero games as part of a team). The degree to which players participate on a team is in general quite stark with over 82 percent of players either never playing on a team or always playing on a team during our sample period.

Because we are able to link observations across time for players, we can observe not only features of a particular game but also a players pattern of play over time. We define a spell as a sequence of consecutive games where a score appears for that game. As Figure 12 shows, the distribution of the number of spells played over the course of our sample period is highly skewed. While the mean number of spells is quite large at 89, the median is 4 spells. The mean length of a spell for players is five games with on average 1.2 days between spells. The distribution of spell lengths and distribution of days between spells are shown in Figures 13 and 14.

Players score on average 30 points in a game or roughly eight percent of possible points that can be scored. Games vary widely in the number of possible points available so that the variance in the share of possible points scored is far greater than the variance in the absolute number of points scored. Figures 15 and 16 show the distribution of points in the first game every played by a player in our sample period and the distribution of points scored in all games, respectively. Comparing across the two distributions, the number of points players score changes as they play additional rounds. Specifically, the distribution of points scored shifts to the right suggesting that player performance increases on average after the first round of play. Section 4.3 discusses the extent to which we observe evidence of learning across rounds of play.

Table 33 presents team-level summary statistics on the number of games played and points scored. Because not all individuals play on teams or on the same team during the entire sample period, the number of games played by teams is on average smaller than the number of games played by individual players. The mean number of points and share of possible points scored in a game are higher for teams than for individual players, however, reflecting the fact that teams scores are always weakly better than the scores of the teams individual players.

3.3. Relationship between Performance and Game Board Features

We are interested in studying the role of past performance as a determinant of intrinsic motivation and its effect on the persistence of play. To do so, we consider four measures of individual player and team performance absolute points scored, points scored as a fraction of total possible points, rank, and a normalized rank. The first two measures depend only on the points scored by the player and are independent of the performance of other players and teams. The third measure, rank, refers to the displayed rank that players observe and is assigned based on the scores for all players and teams in the game. Ties are randomly broken so that players who score the same number of points are randomly assigned a ranking within the range covered by the tie. For example, two players who both score the highest cumulative score for the game would be randomly assigned a rank of one and two respectively. We do not observe the actual displayed rankings in the data and therefore construct player and team rankings based on the observed points scored by each player and team with ties randomly broken. This introduces some measurement error in rankings since our random assignment of players who tie may not exactly correspond to the random assignment generated by Wordsplay.com. However, since any error in rankings is necessarily random, this should only bias the parameter estimates down by introducing noise. Finally, we construct the fourth measure, a normalized rank, by scaling all rankings between zero and one.

The objective of Wordsplay.com is to score points by constructing as many words possible from the letters of a given game board within the allotted time. One way in which players can search for words is by identifying and exploiting the presence of prefixes and suffixes in the game boards. While the total number of prefixes and suffixes is quite numerous, we focus our analysis on the presence of 35 different common prefixes and suffixes that may appear on game boards. Appendix Table 1 lists the prefixes and suffixes used in our analysis along with game-level descriptive statistics on the possible board points from each prefix and suffix. The frequency with which these prefixes and suffixes appear in a game board (column 4) is a function of both the length of prefix or suffix and the frequency with which these prefixes and suffixes occur in the English language. For instance, the fact that s is one of the most common suffixes and is also short in length is reflected in the relatively high frequency (83% of all games) with which it appears in our sample. In contrast, the prefix hyper-, which is both uncommon and long in length, appears in less than one percent of games in our sample.

Our analysis takes advantage of individual level variation in the ability to exploit the presence of prefixes and suffixes that appear on the game board to identify the effect of past performance on the persistence of play. Specifically, we use the number of possible board points that can be scored for each prefix and suffix as instruments for past performance. Appendix Table 5 shows that there is a strong first-stage relationship between the number of possible board points for each prefix and suffix and our measures of past performance for both players and teams. The relationship is strongest for the player points as a fraction of possible board points and team points as a fraction of possible board points. The instruments are somewhat weak for the ranking variables, particularly the normalized rankings. This is unsurprising given that changes in difficulty due to the presence of certain prefixes or suffixes affect all players in a given round and so are more likely to affect points scored rather than relative ranking measures.

3.4. Estimation Framework and Results

In the preceding section, we demonstrated a relationship between the presence of prefixes and suffixes on game boards and performance. We now focus on the effect of past performance on persistence of play, using both spell length ($\textsc{SpellLength}_{is}$) and the probability of a spell ending ($\textsc{SpellEnd}_{ig}$) as measures of persistence. Individual variation in the ability to exploit the presence of prefixes and suffixes, as captured by the possible board points for a given game from the use of the various prefixes and suffixes ($\textsc{AffixPoints}_{is}$), is used to instrument for each of the four past performance measures ($\textsc{Performance}_{is}$) identified in the previous section. We look at the impact of past performance on the persistence of play by estimating variants of the following two equations using two-stage least squares

(2SLS):

$$\textsc{SpellLength}_{is} = \alpha_i + \beta_1(\textsc{Performance})_{is} + \beta_2 X_{is} + \lambda_t + \epsilon_{is} \tag{3.1}$$

and

$$\textsc{SpellEnd}_{ig} = \alpha_i + \beta_1(\textsc{Performance})_{ig} + \beta_2 X_{ig} + \lambda_t + \epsilon_{ig} \tag{3.2}$$

where the subscript i indexes players, s corresponds to spells, and g corresponds to games. The outcome is either spell length or an indicator for whether a spell ends; α_i are player fixed effects; and λ_t are month fixed effects to capture aggregate time trends over the year. Because we are interested in how past performance affects the persistence of play, the four performance measures as denoted by $\textsc{Performance}_{is}$ are based on performance in the first game of a given spell s for player i when estimating the effect of past performance on spell length. In some cases, an indicator, X_{is} or X_{ig}, is included for whether the individual player's team won the round.

The remainder of this section presents our main empirical results from estimating variants of equations (3.1) and (3.2). We first demonstrate the effect of past performance by individual players on the persistence of play, as measured by spell length and the probability of a spell ending. Next, we consider the effect of past performance on the persistence of play for teams rather than individual players. Finally, we investigate whether players demonstrate evidence of learning as they play additional rounds.

3.4.1. The Effect of Player Performance on Persistence of Play

We begin by examining the effect of past performance by individual players on spell length. As a point of departure for the empirical analysis, Figure 17 presents the simple linear relationship between the absolute number of points scored by players during the first game of a spell and the length of the spell. In the figure, spell length increases with the absolute number of points scored, indicating that past performance may motivate continued play during the spell. We next formalize this relationship by estimating (3.1), first using OLS and then instrumenting for $\textsc{Performance}_{is}$ with $\textsc{AffixPoints}_{is}$. While Figure 17 focused on the absolute number of points scored by players, we extend our analysis to consider all four past performance measures of interest. Table 34 reports estimates and standard errors of β_1, with each column representing a separate regression. Each pair of columns presents OLS and IV estimates using a different measure of individual past performance. Column 1 presents the OLS results from estimating (3.1) using the absolute number of points scored by players as the measure of past performance. We see that the inclusion of covariates, and in particular individual fixed effects, reverses the relationship observed in Figure 17 between the absolute number of points scored by players and spell length. Spell length decreases significantly by 0.024 games on average for every ten-point increase in a players score. Column 2 presents estimates after instrumenting for the absolute number of points scored using the possible board points for a given game from exploiting various prefixes and suffixes. The IV estimate is nearly identical to the OLS estimates from Column 1.

Columns 3 and 4 repeat the same exercise using player points scored as a fraction of total possible points as the measure of past performance. While the OLS estimate indicates no effect of past performance on spell length, the IV estimate reported in Column 4 indicates a strong positive relationship between player points as a fraction of total possible points and spell length. Specifically, moving from the 5th to the 95th percentile of the distribution of player points as a fraction of total possible points scored increases spell length by 0.327 games on average. Columns 5 through 9 using player ranking and normalized player rankings present a similar picture to that of player points as a fraction of total possible points. Higher past performance by an individual player significantly increases the length of the spell that is played.

We next consider the effect of past performance by individual players on the probability of ending a spell. A spell ends when a player does not score any points in the following game (i.e. stops playing). For transparency, we focus on linear probability models. Figure 18 plots the simple linear relationship between the absolute number of points scored by players in the first game and the probability of a spell ending. The figure shows a strong negative relationship, suggesting that increasing past performance decreases the probability of a player ending a spell.

Table 35 presents both OLS and IV regressions estimates for each of the four measures of past performance where once again we instrument for $\textsc{Performance}_{is}$ using $\textsc{AffixPoints}_{is}$. Once again, each pair of columns represents a set of OLS and IV estimates using a separate measure of past performance. We find evidence consistent with the results from Table 34 looking at the effect of past performance on spell length. When using the absolute number of points scored by a player, we find a positive relationship between points scored and the probability of a spell ending. Using player points scored as a fraction of possible points and the two ranking measures, however, we find the opposite effect higher past performance decreases the probability of a spell ending.

3.4.1.1 Do Players Exhibit Satisficing Behavior?

We've assumed so far a linear relationship between past performance and persistence of play. Of course, this need not be true. Rather than process game-by-game whether it is optimal to quit, players may instead adopt a simple heuristic decision strategy by which they quit once they've achieved some threshold of success. Individuals who adopt such a strategy can be thought of as exhibiting "satisficing" behavior where there exists some satisficing level of success, attainment of which would induce the player to quit playing (Simon 1955; Caplin et al. 2011). We consider a particular decision strategy where players may decide to quit once they or their team has achieved some satisficing rank level. To test whether players exhibit this form of satisficing behavior, we estimate the following equations:

$$\begin{aligned}\textsc{SpellLength}_{is} = \alpha_i &+ \beta_{11}(\textsc{PlayerRank})_{is} + \beta_{12}(\textsc{PlayerRankAboveThresh})_{is} \\ &+ \beta_2 X_{is} + \beta_{31}(\textsc{TeamRank})_{is} + \beta_{32}(\textsc{TeamRankAboveThresh})_{is} \\ &+ \lambda_t + \epsilon_{is}\end{aligned} \tag{3.3}$$

and

$$\begin{aligned}\textsc{SpellEnd}_{ig} = \alpha_i &+ \beta_{11}(\textsc{PlayerRank})_{ig} + \beta_{12}(\textsc{PlayerRankAboveThresh})_{ig} + \beta_2 X_{ig} \\ &+ \beta_{31}(\textsc{TeamRank})_{ig} + \beta_{32}(\textsc{TeamRankAboveThresh})_{ig} \\ &+ \lambda_t + \epsilon_{ig}\end{aligned} \tag{3.4}$$

where we use both player rankings and team rankings as measures of past performance and we include two indicators, $\textsc{PlayerRankAboveThresh}_{is}$ and $\textsc{TeamRankAboveThresh}_{is}$, which equal to one when the player or his team achieves a ranking above the specified threshold. The remaining covariates are as in (3.1) and (3.2).

With equations (3.3) and (3.4), we next estimate the effect of past performance on persistence of play, accounting for potential "goal cutoffs." A player (or team) is considered to have a ranking above the threshold if his (or his team's) ranking is greater than or equal to twenty. Table 36 reports the results of this estimation. The first two columns look at the effect of past performance, as measure by player and team rank, on spell length while the last two columns look at its effect on the probability of a spell ending. For each pair of columns, the first estimates a straightforward OLS while the second column instruments for past performance with AFFIXPOINTS$_{is}$, the possible board points for a given game from the use of the various prefixes and suffixes. As Table 36 shows, we find no evidence that players exhibit satisficing behavior. Using higher rank thresholds (rank above ten and rank above five) does not significantly change the results.

3.4.2. The Effect of Team Performance on Persistence of Play

Players have the option to join a team for which they will receive a ranking in addition to the individual player ranking. Notably, players must specify the name of team they wish to join and thus cannot choose to simply join any team at random. Alternatively, they may create their own team name and invite others to join. Appendix Tables 3 and 4 presents individual level summary statistics by whether the player was or was not on a team during the first game played. While individual players and team players perform equally as well, players on a team play far more games on average than individual players.

We next consider how a team's past performance affects an individual players persistence of play. As before, we look at two different measures of persistence of play spell length and the probability of a spell ending. We re-estimate (3.1) and (3.2), using the same four measures of performance. Table 37 reports estimates for each performance measure on its effect on spell length. The results are qualitatively similar to those of Table 34. When using the absolute number of points scored by the team as a measure of performance, we find that performance is negatively correlated with spell length. A ten-point increase in a team's total score decreases the number of spells played by 0.078 games on average. In contrast, team points as a fraction of total possible points and team ranking measures exhibit the opposite relationship with spell length increasing with performance. Table 38 considers next the effect of team performance on the probability of a spell ending. Once again, we find that the absolute number of points scored by a team has a positive relationship with the probability of a spell ending. Estimates for the other three performance measures, however, suggest that higher past performance decreases the probability that an individual will end the spell.

3.4.3. Learning Over Time

Because we are able to observe both an individual's pattern of play over time as well as his performance, we can use the data to assess whether players are learning over time. To do so, we look at how player points evolve over time to see whether performance is increasing over time. Because game boards appear randomly and are not a function of past performance, any trend in player points over time cannot be attributed to trends in the difficulty of the game boards themselves. Figure 19 plots the average player points by game order for players who have played at least 100 games during our sample period. For clarity, we restrict the plot to exclude games beyond a player's hundredth game. We also drop players' first game to avoid potentially confounding improvement in the ability to find words and score points with any initial learning curve from a player familiarizing himself with the game. We restrict the sample to include only players who have played at least 100 games (roughly 29 percent of our sample of players) due to potential selection bias in who chooses to only play a few games. Figure 19 shows a clear upward trend as players progress in the number of games played, suggesting that players are in fact improving over time as they play additional games.

3.5. Conclusion

Using data from an online word-hunting game, we examine the role of past performance as a source of intrinsic motivation. We exploiting quasi-random variation in performance to identify the casual effect of past performance, as measured by the absolute number of points scored, points scored as a fraction of possible board points, and both rank and normalized rank, on persistence of play. We find that an increase in the share of board points scored or an increase in rank significantly increases the length of a spell and decreases the probability that a spell will end. This is true of both player performance and team performance measures. We also find suggestive evidence that players are learning over time as they play additional games.

FIG. 10.—4x4 Game Board

L	P	I	M
K	T	I	D
O	H	F	U
T	D	A	M

NOTE.— This figure shows a sample 4x4 game board from Wordsplay.com, an online implementation of the word-hunting game Boggle.

Fig. 11.—Games Per Player

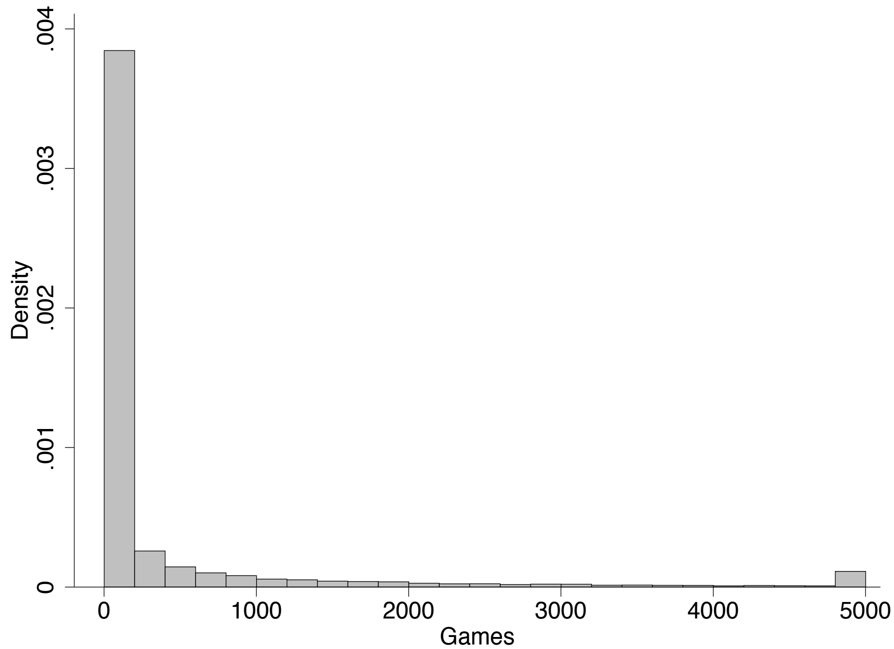

Note.— This figure displays the distribution of the number of games player per player. Games greater than 5000 are pooled in the last bin.

Fig. 12.—Spells Per Player

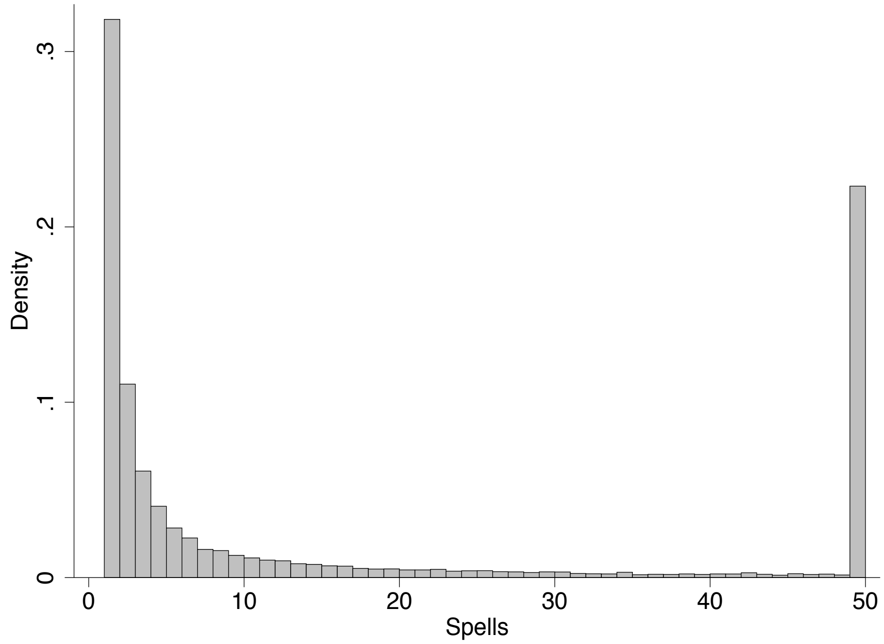

Note.— This figure displays the distribution of the number of spells per player using the first player game. A spell is defined as a sequence of consecutive games where a score appears for that individual, with no interruptions longer than one game. Spells greater than 50 are pooled in the last bin.

Fig. 13.—Length of Spell

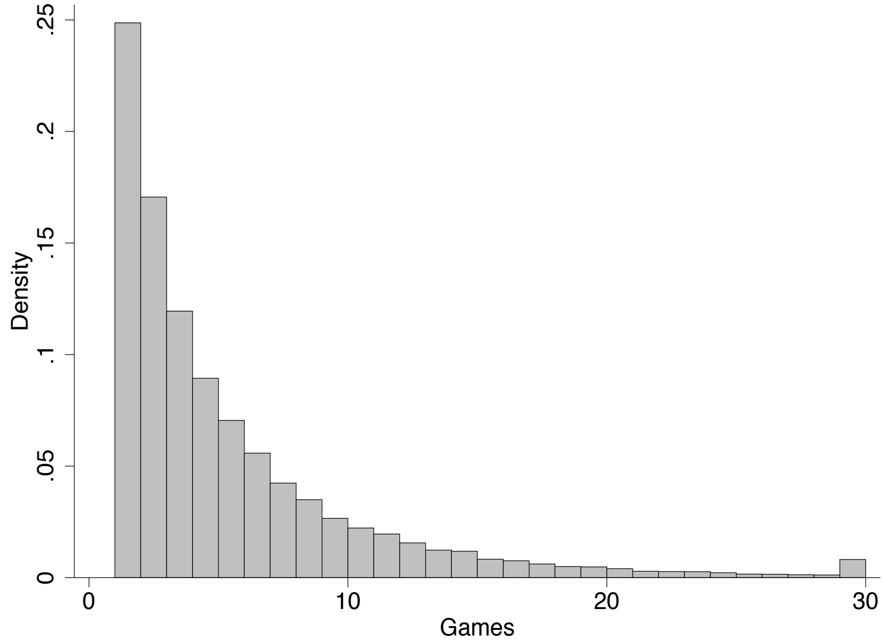

Note.— This figure displays the distribution of the length of a spell in terms of games using the first player game. Spell length is calculated as the number of games between the start and end of a spell. Consecutive games longer than 30 are pooled in the last bin.

Fig. 14.—Days Between Spells

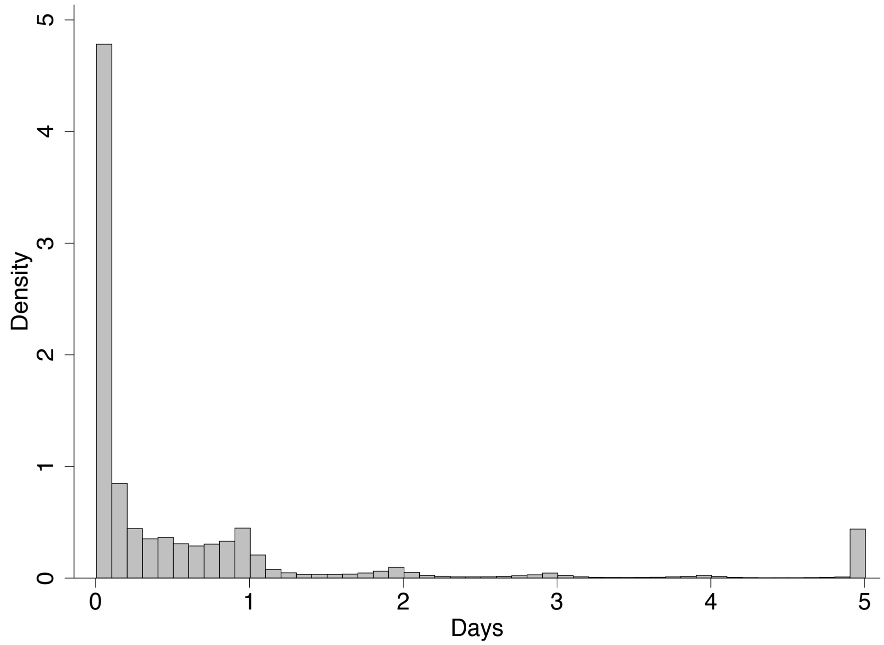

Note.— This figure displays the distribution of time elapsed between spells, in terms of days, using the first game of a spell. Days are calculated as games played in between spells as a function of time. Days greater than 5 are pooled in the last bin.

Fig. 15.—Points Per Player in First Game

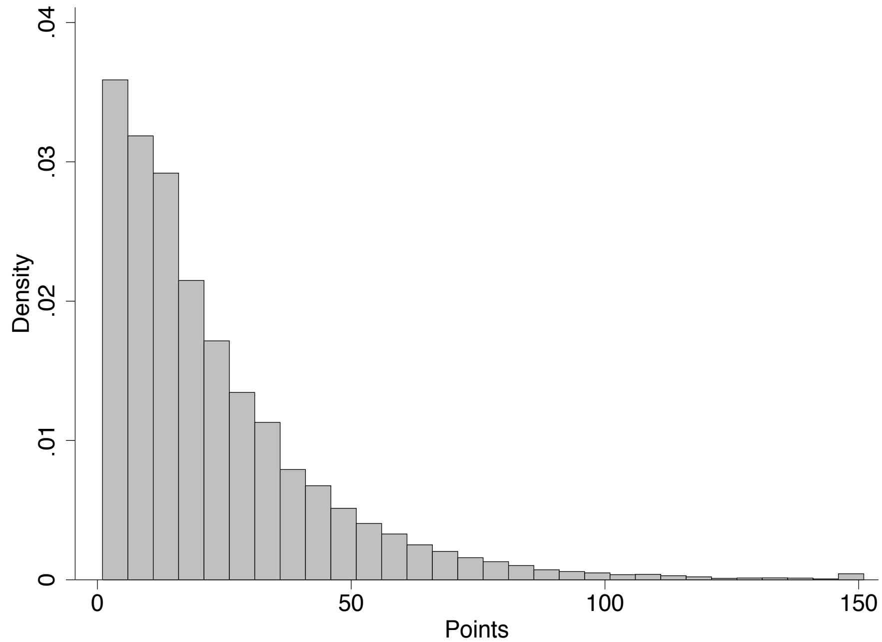

Note.— This figure displays the distribution of points per player in the first game played. Points accumulated greater than 150 are pooled in the last bin.

Fig. 16.—Points Per Player in All Games

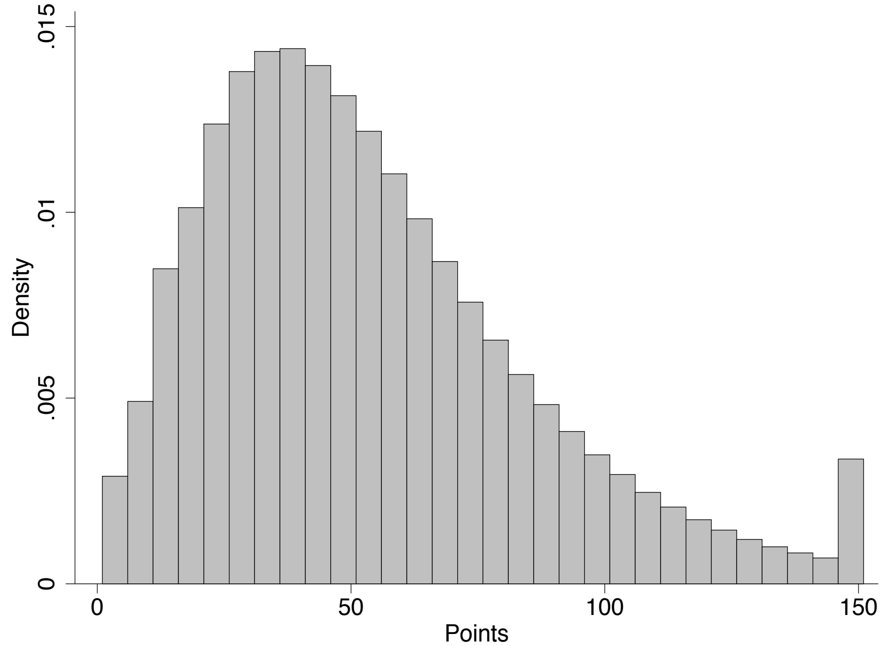

Note.— This figure displays the distribution of points per player in all games played. Points accumulated greater than 150 are pooled in the last bin.

FIG. 17.—Linear Prediction of Spell Length on Player Points in First Game of Spell

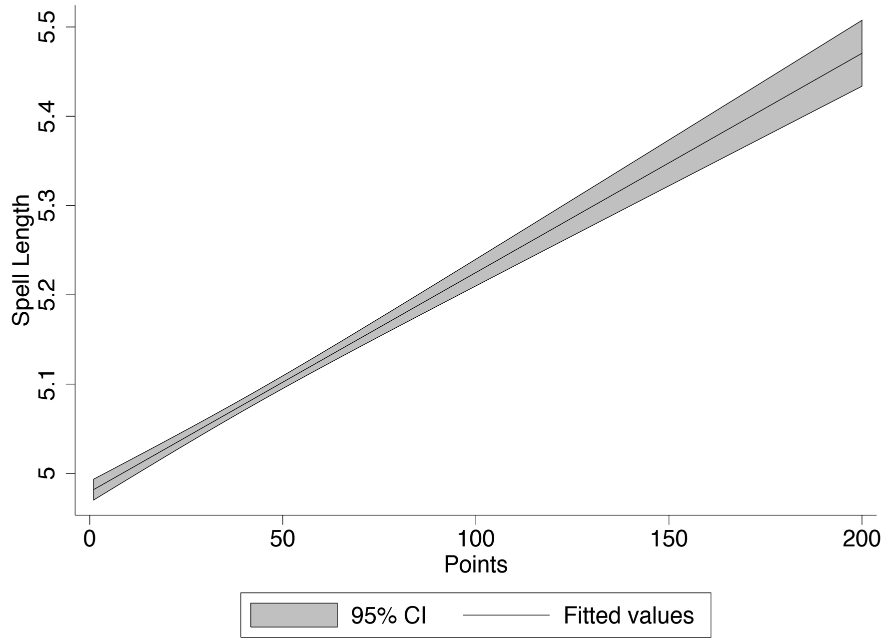

NOTE.— This figure plots the linear prediction of length of spell on player points in the first game of the spell.

Fig. 18.—Linear Prediction of Spell End on Player Points in First Game of Spell

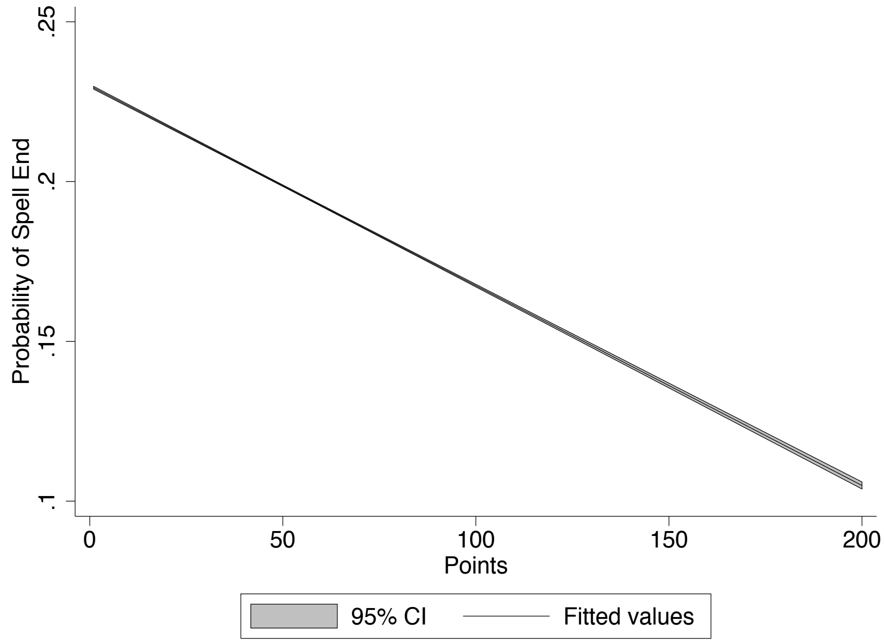

Note.— This figure plots the linear prediction of the probability of spell end on player points.

FIG. 19.—Learning Based on Average Player Points

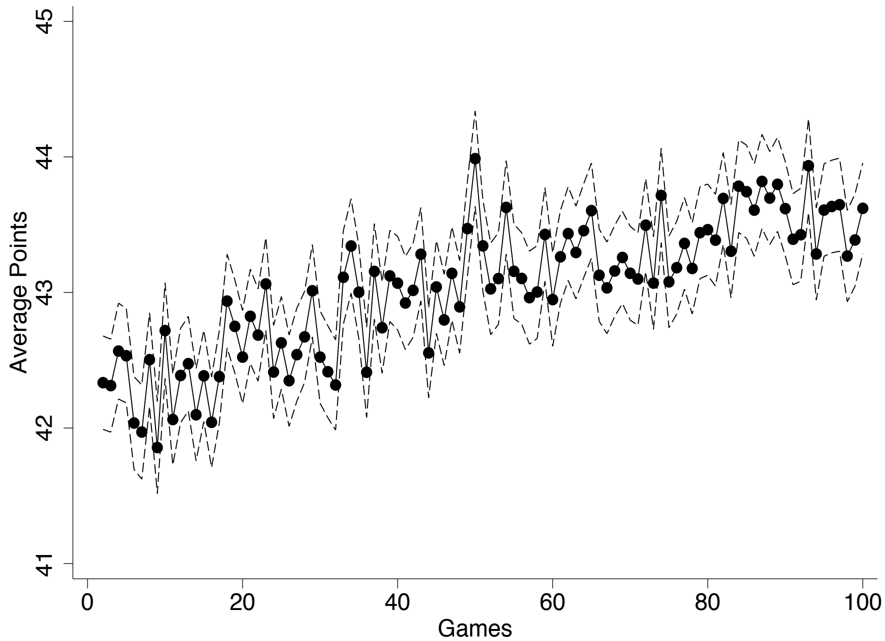

NOTE.— This figure displays mean player points for games 2 to 100 for players who have played at least 100 games.

TABLE 32

INDIVIDUAL LEVEL SUMMARY STATISTICS

	Mean	SD	Median	N
Number of games per player	453.79	1371.27	16	24,433
Number of points scored per player	30.45	19.92	27	24,433
Share of possible points scored	0.08	0.05	0.08	24,433
Number of games played per player on team	118.65	656.56	0	24,433
Number of spells per player	89.28	253.55	4	24,433
Length of spell per player	5.06	5.58	3	24,433
Number of days between spells	1.18	7.72	0	2,156,958

NOTE.—The above summary statistics are calculated using the first player game on Wordsplay.net during 2009. The data in the table reflect the mean, standard deviation and median of the indicated statistics calculated at the individual level. A spell of play for an individual is defined as a sequence of consecutive games where a score appears for that individual, with no interruptions longer than one game.

TABLE 33

Team Level Summary Statistics

	Mean	SD	Median	N
Number of games per team	295.74	4513.52	7	9803
Number of points scored per team	45.72	33.17	40.1	9803
Share of possible points scored	0.13	0.08	0.1	9803

Note.—The above summary statistics are calculated using the first player game on Wordsplay.net during 2009. The data in the table reflect the mean, standard deviation and median of the indicated statistics calculated at the team level.

TABLE 34

Impact of Player Performance on Spell Length

	OLS (1)	IV (2)	OLS (3)	IV (4)	OLS (5)	IV (6)	OLS (7)	IV (8)
Points per Player	-0.0024***	-0.0029 ***						
	(0.0002)	(0.0002)						
Team Win	0.3184***	0.3224***	0.2975***	0.2744***	0.3052***	0.1916***	0.3076***	-0.5270***
	(0.0605)	(0.0608)	(0.0607)	(0.0613)	(0.0609)	(0.0637)	(0.0608)	(0.1465)
Share of Available Points on Board per Player			-0.0108	1.2882***				
			(0.0780)	(0.1177)				
Player Rank per Game					-0.0024***	0.0326***		
					(0.0003)	(0.0031)		
Normalized Player Rank							-0.1376***	11.0512***
							(0.0283)	(1.6809)
Month Indicators	Y	Y	Y	Y	Y	Y	Y	Y
Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
N	2,173,340	2,173,340	2,173,340	2,173,340	2,173,340	2,173,340	2,173,340	2,173,340

Note.—This table presents the relationship between four performance measures and spell length using the first game of a spell. Team win indicates whether the individual player's team won the game. Errors are clustered at the player level.

TABLE 35

IMPACT OF PLAYER PERFORMANCE ON SPELL END

	OLS (1)	IV (2)	OLS (3)	IV (4)	OLS (5)	IV (6)	OLS (7)	IV (8)
Points per Player	-0.0004***	0.0003 ***						
	(0.00002)	(0.00002)						
Team Win	-0.0213***	-0.0263***	-0.0198***	-0.0222***	-0.021**	0.0085***	-0.0156**	-0.3323***
	(0.0018)	(0.0018)	(0.0018)	(0.0018)	((.0018)	(0.0038)	(0.0017)	(0.1307)
Share of Available Points on Board per Player			-0.3025***	-0.1373***				
			(0.007)	(0.0088)				
Player Rank per Game					-0.0007***	-0.0073***		
					(0.00002)	(0.0006)		
Normalized Player Rank							-0.1366**	4.8859***
							(0.0034)	(2.0793)
Month Indicators	Y	Y	Y	Y	Y	Y	Y	Y
Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
N	11,085,064	11,085,064	11,085,064	11,085,064	11,085,064	11,085,064	11,085,064	11,085,064

NOTE.—This table presents the relationship between four performance measures, accounting for team win, and spell end (discontinuing play). Team win indicates whether the individual player's team won the game. Errors are clustered at the player level.

TABLE 36

Impact of Player Performance on Spell Length and Spell End Accounting for "Goal" Cutoffs

	Spell Length		Spell End	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
Player Rank	-0.0046***	0.0161	-0.0008***	0.0289***
	(0.0004)	(0.0299)	(0.00003)	(0.0066)
Team Rank	0.0021***	0.0123	0.0003***	-0.0229***
	(0.0005)	(0.0264)	(0.00002)	(0.0049)
Player Rank Top 20	0.1175***	1.4776	-0.0086***	0.4456
	(0.028)	(3.2572)	(0.0015)	(0.4163)
Team Rank Top 20	0.5039***	4.9209	-0.054***	0.7204
	(0.053)	(4.15)	(0.0019)	(0.6516)
Team Win			-0.0066***	0.315
			(0.0017)	(0.3175)
Month Indicators	Y	Y	Y	Y
Fixed Effects	Y	Y	Y	Y
N	2,173,340	2,173,340	11,085,064	11,085,064

Note.—This table presents the impact of overall player and team rank, top 20 rank for player and team, and whether the team won or not. Team win was only facotred in for spell end. For these purposes, the rank displayed to players was used in calculations.

TABLE 37

Impact of Team Performance on Spell Length

	OLS (1)	IV (2)	OLS (3)	IV (4)	OLS (5)	IV (6)	OLS (7)	IV (8)
Points per Team	0.0046***	-0.0078***						
	(0.0004)	(0.0007)						
Team Win	-0.1572***	1.0663***	-0.2998***	-0.801***	0.2526***	-1.1329***	0.2976***	0.2983***
	(0.0517)	(0.0982)	(0.0512)	(0.1252)	(0.0601)	(0.139)	(0.0614)	(0.0615)
Share of Available Board Points per Team			2.5665***	4.7199***				
			(0.1725)	(0.4666)				
Team Rank					0.001***	0.0319***		
					(0.0004)	(0.0028)		
Normalized Team Rank							0.0135**	0.729
							(0.0062912)	(0.814)
Month Indicators	Y	Y	Y	Y	Y	Y	Y	Y
Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
N	2,173,340	2,173,340	2,173,340	2,173,340	2,173,340	2,173,340	2,173,340	2,173,340

NOTE.—This table presents the relationship between four performance measures and spell length using the first game of a spell. Team win indicates whether the individual player's team won the game. Errors are clustered at the player level.

TABLE 38
Impact of Team Performance on Spell End

	OLS (1)	IV (2)	OLS (3)	IV (4)	OLS (5)	IV (6)	OLS (7)	IV (8)
Points per Team	-0.0003***	0.0006***						
	(0.00001)	(0.00004)						
Team Win	0.0002	-0.0733***	0.0121***	0.0603***	-0.0162***	0.1367***	-0.0241***	-0.0255***
	(0.0017)	(0.0045)	(0.0019)	(0.007)	(0.0018)	(0.0115)	(0.0018)	(0.0018)
Share of Available Board Points per Team			-0.1767***	-0.4114***				
			(0.006)	(0.0303)				
Team Rank					-0.0002***	-0.0045***		
					(0.00001)	(0.0003)		
Normalized Team Rank							-0.0007***	-0.2811***
							(0.0002)	(0.0387)
Month Indicators	Y	Y	Y	Y	Y	Y	Y	Y
Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
N	11,085,064	11,085,064	11,085,064	11,085,064	11,085,064	11,085,064	11,085,064	11,085,064

Note.—This table presents the relationship between four performance measures, accounting for team win, and spell end (discontinuing play). Team win indicates whether the individual player's team won the game. Errors are clustered at the player level.

APPENDIX

A.1. Sample Restrictions

The empirical analysis in this paper uses data between 1990 and 2010 from the quarterly interview portion of the Consumer Expenditure Survey (CEX). This appendix discusses in detail the imposed sample restrictions and follows largely from the sample restrictions described in much of the literature (Lusardi 1996; Souleles 1999, 2002; Hsieh 2003; Parker 1999; and Stephens, Jr. 2008). These restrictions are meant to reduce the potential influence of measurement error or changes in circumstances unrelated to the presence of third paychecks on my estimates. The final sample including weekly, bi-weekly, and monthly heads of households is composed of 24,822 household-month observations for 7776 households, of which 4316 have heads who report working bi-weekly. Table 39 provides information on the number of observations and households remaining after each sample restriction. Additional details on sample restrictions are described below.

I refer to the reference individual in the CEX in this appendix and throughout the paper as the head of household. The reference person is the first member mentioned by the respondent when asked the name of the person or one of the persons who owns or rents the home. The relationships of all other members of the household are defined in relation to this reference person. More information on how the CEX defines the reference person can be found at the CEX website that is maintained by the Bureau of Labor Statistics at http://www.bls.gov/cex.

I first exclude from the sample any consumer units that are composed of multiple households. With this exclusion, consumer units and households can be thought of interchangeably. I drop any households where the head is either 1) employed in farming, forestry, or fishing, 2) self-employed, or 3) working without pay. I further drop any households that report living in student housing or whose head is less than 24 or greater than 64 years of age.

The next set of restrictions relates to data issues in income and expenditure reports. Households are dropped if they are flagged for having an incomplete or top-coded income report or invalid checking or savings report or if their reported income falls below the 1st percentile of the income distribution. Households with heads whose reported gross pay is either missing, is flagged as inconsistent or top-coded, or falls below the 1st percentile of the distribution of reported gross pay are also excluded. I also drop households who do not report the period of time covered by their last gross pay. I drop any households if any component of the aggregate expenditure groups does not have an associated month and year of expenditure. I also exclude households who report ever receiving meals as pay and household-quarters if they report zero food expenditures in any month of the quarter. Finally, I exclude households who report large changes in non-durable expenditure (log change greater than 2) between any two consecutive months.

The last set of restrictions excludes households if the age of any household member changes by more than one year from quarter to quarter or if the household is missing information on family size, age of the head of household, or number of children (defined as members of the households younger than 18).

In addition to the preceding exclusions, I include two additional restrictions as explained in Section 1.3: (i) households where other members work at the same pay frequency as the head of household are not included and (ii) the sample is restricted to individuals who report working full time over the past year (at least 50 weeks) and whose reported gross pay has not changed between the second and fifth interview.

TABLE 39: Sample Restriction Details

Restriction	Obs Level	Obs Count	Hhld Count	Literature
Initial Sample	Hhld-Quarter	541,286	184,893	-
Multiple Hhlds in Same Consumer Unit	Hhld-Quarter	517,924	173,472	Souleles 1999; Souleles 2002
Employed in Farming, Forestry, or Fishing	Hhld-Quarter	510,813	171,179	Lusardi 1996; Souleles 1999; Souleles 2002
Self-Employed or Not Working	Hhld-Quarter	304,037	105,668	Lusardi 1996
Live in Student Housing	Hhld-Quarter	302,632	104,734	Souleles 1999; Souleles 2002; Hsieh 2003
Missing Income Information	Hhld-Quarter	232,308	82,169	Lusardi 1996; Parker 1999; Stephens, Jr. 2008
Top-coded Income Flag	Hhld-Quarter	219,277	78,109	Parker 1999; Stephens, Jr. 2008
Invalid Checking or Savings Report	Hhld-Quarter	176,248	64,742	Lusardi 1996
No Associated Date for Expenditure	Hhld-Quarter	176,248	64,742	Souleles 1999; Souleles 2002
Zero Reported Food Expenditure	Hhld-Quarter	175,961	64,675	Souleles 1999; Souleles 2002; Hsieh 2003
Received Meals as Pay	Hhld-Quarter	164,503	60,840	Souleles 1999; Souleles 2002
Inconsistent/Missing Gross Pay Or Pay Period	Hhld-Quarter	135,369	51,979	-
Gross Pay or Income Less than 1st Pct.	Hhld-Quarter	131,487	50,426	-
Top-coded Gross Pay Flag	Hhld-Quarter	128,281	49,295	-
Gross Pay Changed Across Interviews	Hhld-Quarter	33,599	22,330	-
Pay Frequency Changed Across Interviews	Hhld-Quarter	33,305	22,233	-
Work $<$ 50 Weeks	Hhld-Quarter	25,196	16,598	-
Age Change $>$ 1 Year Across Quarters	Hhld-Quarter	25,060	16,545	Souleles 2002
Missing Information on Family Size	Hhld-Quarter	25,060	16,545	Parker 1999; Hsieh 2003
Not Paid Weekly, Bi-weekly, or Monthly	Hhld-Quarter	23,158	15,304	-
Other Household Members Paid Same Frequency	Hhld-Quarter	16,199	10,931	-
Age of Head Less Than 24 or Greater Than 64	Hhld-Quarter	14,639	9,784	Souleles 1999; Souleles 2002
Convert Observation Level†	Hhld-Month	43,917	9,784	-
Log Change in Non-Durables $>$ 2	Hhld-Month	32,658	7,776	Lusardi 1996; Parker 1999
Final Estimation Sample (Non-missing ΔC_t)	Hhld-Month	24,822	7,776	-

NOTE.—This table provides details on the sample restrictions taken in this paper. Table entries represent the number of observations and households remaining after dropping observations with the indicated characteristic. The last column lists examples of prior literature which impose similar sample restrictions. †Data is reshaped from the household-quarter level to the household-month level using information on the timing of expenditure.

A.2. Classification Correction

The CEX asks individuals to report the period of time covered by their last gross pay. This allows me to identify the pay frequency of the heads of household and, in particular, whether the head is paid bi-weekly. In order to study spending responses following three paycheck months, I create a variable that indicates whether the previous month was a three paycheck month for a given bi-weekly worker's schedule. However, one drawback of using the CEX is that I do not observe the actual date on which the last pay occurred. Because of this, I am unable to observe which of the two possible alternate schedules a given bi-weekly worker is paid by. Table A.2.1 lists the three paycheck months for the two alternate schedules from 1989 to 2010.[1] Each month of the year serves as a three paycheck month on one of the schedules at least once during the sample period.

Because I am unable to observe by which schedule a given bi-weekly worker is paid, I allow the worker to be on either schedule in the Section 1.4 regression specification. The indicator, $\mathbf{1}_{\{t-1 \in S\}}$, from Equation 1.6 is set equal to one if the previous month was a three paycheck month on *either* schedule. The indicator can be thought of as a noisy measure of which months are three paycheck months based on the bi-weekly worker's true schedule. For ease of exposition, let $x = \mathbf{1}_{\{t-1 \in S\}}$ denote this indicator, and let $X = \mathbf{1}_{\{t-1 \in S_j\}}$ be an indicator for whether the previous month is a three paycheck month based on the worker's true schedule $j \in \{1, 2\}$. Then we can write the following relation

$$x = X + u \tag{A.1}$$

where u is an error term taking the value $u = 0$ when X is measured without error and $u = 1$ when X is mis-measured. Again for ease of exposition, let the change in consumption ΔC be denoted by y. Equation 1.6 can then be re-expressed as the following

$$\begin{aligned} y &= X\beta + Z\gamma + \epsilon \\ &= x\beta + Z\gamma + (\epsilon - u\beta) \end{aligned} \tag{A.2}$$

where γ is a $[(k-1) \times 1]$ vector of parameters and Z is an $[n \times (k-1)]$ matrix of the taste shifters and time dummies. The measurement error introduces a bias in the estimation of my parameter of interest, β. Given the binary nature of the indicator variable, the measurement error can be thought of as classification error in the bi-weekly worker schedules. Moreover, this classification error is non-classical in nature because the true value of the indicator variable is necessarily negatively correlated with the error. To see this, first note that

[1] The table includes three paycheck months in 1989 to account for the fact that households who were interviewed in the first quarter of 1990 may report expenditures in 1989.

whenever $x = 0$, the measurement error $u = 0$ since these are months that do not follow three paycheck months on either schedule. On the other hand, whenever $x = 1$, then either $u = 0$ if the month follows a three paycheck month on the worker's true schedule (i.e. $X = 1$) or $u = 1$ if the month follows a three paycheck month on the other schedule (i.e. $X = 0$). Thus we have that $\text{Cov}(X, u) < 0$.

A.2.1. Direction of bias

Classification error of this sort biases downwards the estimates of β in naive OLS regressions. To show this, let $\tilde{X} = [x \quad Z]$ so that

$$y = \tilde{X}b + e. \tag{A.3}$$

I make the following three assumptions

$$\begin{aligned}&\text{(A1) } E(X'e) = 0 \text{ and } E(Z'e) = 0\\ &\text{(A2) } E(Z'u) = 0\\ &\text{(A3) } \text{Cov}(X, u) < 0.\end{aligned}$$

The first assumption (A1) is a standard assumption and states that the regressors of the true population regression are orthogonal to the error terms. The second assumption (A2) states that the classification error from mismeasurement of X is orthogonal to the other regressors Z. The final assumption (A3) is that the classification error is negatively correlated with the true indicator, X. Following Aigner (1973) and Black et al. (2000), applying least squares to Equation A.3 gives the following estimators

$$\hat{b}^{OLS} = \begin{bmatrix} \hat{\beta} \\ \hat{\gamma} \end{bmatrix}^{OLS} = (\tilde{X}'\tilde{X})^{-1}\tilde{X}'y$$

where

$$\tilde{X}'\tilde{X} = \begin{bmatrix} x'x & x'Z \\ Z'x & Z'Z \end{bmatrix} \text{ and } \tilde{X}'y = \begin{bmatrix} x'y \\ Z'y \end{bmatrix}.$$

The sampling error e is then given by

$$e = \begin{bmatrix} \hat{\beta} \\ \hat{\gamma} \end{bmatrix}^{OLS} - \begin{bmatrix} \beta \\ \gamma \end{bmatrix} = (\tilde{X}'\tilde{X})^{-1}\tilde{X}'\epsilon - (\tilde{X}'\tilde{X})^{-1}\tilde{X}'\mu\beta.$$

Given my first assumptions, A1, I can write

$$\text{plim}\hat{b}^{OLS} = \begin{bmatrix} \beta \\ \gamma \end{bmatrix} - \beta\Sigma_{\tilde{X}}^{-1}\text{Cov}(\tilde{X}, u) \tag{A.4}$$

where $\Sigma_{\tilde{X}} = \text{plim}(\frac{1}{n}\tilde{X}'\tilde{X})$.

To determine the direction of the bias, it is necessary to estimate both $\Sigma_{\tilde{X}}^{-1}$ and $\text{Cov}(\tilde{X}, u)$. From assumption A2, it holds that $\text{Cov}(Z, u) = 0$. Thus, for my key parameter of interest, β, Equation A.4 simplifies to

$$\text{plim}\hat{\beta}^{OLS} = \beta - \beta s_{11}\text{Cov}(x, u) \tag{A.5}$$

where s_{11} is the first element in $\Sigma_{\tilde{X}}^{-1}$. The covariance in this expression depends on the joint distribution of (x, u). To determine this, we need to know the probabilities of misclassification and the probability that the previous month is a three paycheck month. Let $\tilde{P} \equiv \text{Prob}(x = 1)$ and $\tilde{Q} = 1 - \tilde{P}$ denote the probabilities of the previous month being and not being a three paycheck month, respectively. Further, recall that months that are not three paycheck months on either schedule ($t-1 \notin S$) are correctly measured so that we can define $\eta \equiv \text{Prob}(u = 0|x = 0) = 1$. For months that are three paycheck months on either schedule ($t - 1 \in S$), the probability of misclassification depends on both the proportion of individuals on each schedule and the probability that a given observation follows a three paycheck month under Schedule 1 versus Schedule 2. Let $\lambda \in [0, 1]$ be the probability an individual is a Schedule 1 individual and $p \in [0, 1]$ be the probability a given observation follows a three paycheck month under Schedule 1 ($t-1 \in S_1$) rather than Schedule 2. Then the probability of misclassification for months that are three paycheck months under either schedule can be defined as $\nu \equiv \text{Prob}(u = 1|x = 1) = \lambda \cdot (1 - p) + (1 - \lambda) \cdot p$. The marginal distribution of x is thus Bernoulli with parameter $\tilde{P}$. Similarly, the marginal distribution of u is Bernoulli with parameter $\nu\tilde{P}$. Thus the covariance between x and u is given by

$$\begin{aligned} \text{Cov}(x, u) &= \nu\tilde{P} - \tilde{P}(\nu\tilde{P}) \\ &= \nu\tilde{P}\tilde{Q} \end{aligned}$$

Note that because $0 \leq \nu, \tilde{P}, \tilde{Q} \leq 1$, it must be that $0 \leq Cov(x, u) \leq 1$.

All that remains is to show that the first element, s_{11}, of $\Sigma_{\tilde{X}}^{-1}$ is positive. Because $\tilde{X}$ is full column rank, it follows that $\Sigma_{\tilde{X}}$ is positive-definite as is its inverse $\Sigma_{\tilde{X}}^{-1}$. The upper left determinants of positive definite matrices are positive, and so $s_{11} > 0$. Therefore, $\text{plim}\beta^{OLS} = \beta - \beta s_{11}\text{Cov}(x, u) \leq \beta$ and the estimate is inconsistent and downward biased.

A.2.2. Relationship between bias and proportion of individuals on either schedule

Because I am unable to observe the proportion of individuals on either schedule, it is important to understand how the bias due to this measurement error varies with that proportion. Let $\beta^R = \frac{\beta}{\beta^{OLS}} = (1 - s_{11}\text{Cov}(x,u))^{-1}$ be the ratio of the true value to the estimated value of the parameter of interest. I establish two facts regarding this bias ratio and its relationship with the proportion of individuals, λ, on Schedule 1.

Proposition A.2.1.

i. If $p = (1-p)$, then the bias ratio is independent of λ.

ii. If $p \neq (1-p)$, then the bias ratio is increasing in λ for $p < \frac{1}{2}$ and decreasing in λ for $p > \frac{1}{2}$.

Proof. To see why these facts hold, recall that λ only enters into the bias ratio through $\text{Cov}(x,u) = \nu\tilde{P}\tilde{Q} = [\lambda(1-p) + (1-\lambda)p]\tilde{P}\tilde{Q}$. When $p = (1-p)$, this expression simplifies to $\text{Cov}(x,u) = \frac{1}{2}\tilde{P}\tilde{Q}$ which does not depend on λ. Hence, the bias ratio does not depend on λ. When $p \neq (1-p)$, then

$$\begin{aligned} \frac{\partial\beta^R}{\partial\lambda} &= (1 - s_{11}\nu\tilde{P}\tilde{Q})^{-2}[s_{11}\tilde{P}\tilde{Q}(1-2p)] \\ &= (\beta^R)^2[s_{11}\tilde{P}\tilde{Q}(1-2p)]. \end{aligned} \tag{A.6}$$

Equation A.6 is positive if $p < \frac{1}{2}$ and negative if $p > \frac{1}{2}$. □

The bias ratio is thus increasing in λ for $p < \frac{1}{2}$ and decreasing in λ for $p > \frac{1}{2}$. Note, however, that because the set of three paycheck months under Schedule 1 is the same size as the set of three paycheck months under Schedule 2 (i.e $|S_1| = |S_2|$), the value for p converges in probability to $\frac{1}{2}$. Thus for an arbitrarily large sample size, the bias ratio is independent of the proportion of individuals on either schedule.

A.2.3. Simulation

I next run a series of simulations to gauge the magnitude of the bias and the extent to which it depends on the proportion of the sample that is on Schedule 1 as opposed to Schedule 2. To do this, I randomly assign a fraction, $\lambda \in [0,1]$, of the sample observations to be Schedule 1 individuals and $1-\lambda$ to be Schedule 2 individuals. I next create an indicator $\mathbf{1}_{\{t-1 \in S_j\}}$ for whether the previous month is a three paycheck month according to the worker's assigned schedule $j \in \{1,2\}$. I further assume that I know the true data generating process for changes in consumption (i.e. the coefficients in Equation A.2 are known) which allows me to generate a "true" consumption path for each household. I then regress these true consumption changes on observed taste shifters, time dummies, and the mis-measured indicator $\mathbf{1}_{\{t-1 \in S\}}$ for whether a given month follows a worker's three paycheck month and compare the estimated coefficients with the values used to generate the consumption variable. To gauge the extent to which the bias introduced by classification error depends on λ, I run this simulation for values of λ ranging from zero to one in increments of 0.01. Figure A.2.1 shows the ratio of the true value of β to the estimated value $\hat{\beta}$ using the mismeasured indicator for different values of λ. The true value of β is on average twice the size of the true value.

TABLE A.2.1: The Timing of Three Paycheck Months from 1989-2010

	Schedule 1			Schedule 2		
Year	Month 1	Month 2	Month 3	Month 1	Month 2	Month 3
1989	Mar	Sept	.	Jun	Dec	.
1990	Mar	Aug	.	Jun	Nov	.
1991	Mar	Aug	.	May	Nov	.
1992	Jan	Jul	.	May	Oct	.
1993	Jan	Jul	Dec	Apr	Oct	.
1994	Jul	Dec	.	Apr	Sept	.
1995	Jun	Dec	.	Mar	Sept	.
1996	May	Nov	.	Mar	Aug	.
1997	May	Oct	.	Jan	Aug	.
1998	May	Oct	.	Jan	Jul	.
1999	Apr	Oct	.	Jan	Jul	Dec
2000	Mar	Sept	.	Jun	Dec	.
2001	Mar	Aug	.	Jun	Nov	.
2002	Mar	Aug	.	May	Nov	.
2003	Jan	Aug	.	May	Oct	.
2004	Jan	Jul	Dec	Apr	Oct	.
2005	Jul	Dec	.	Apr	Sept	.
2006	Jun	Dec	.	Mar	Sept	.
2007	Jun	Nov	.	Mar	Aug	.
2008	May	Oct	.	Feb	Aug	.
2009	May	Oct	.	Jan	Jul	.
2010	Apr	Oct	.	Jan	Jul	Dec

NOTE.—This table lists all three paycheck months in a given year depending on which of the two possible pay schedules a bi-weekly worker may be on. Note that in some years, the calendar is such that bi-weekly workers receive two paychecks each month with the exception of *three* not two months, during which they receive three. The calendar includes three paychecks months for 1989 since some households interviewed in the first quarter of 1990 report expenditures from 1989. All months of the calendar year appear at least once in this table.

Fig. A.2.1: Simulated Bias

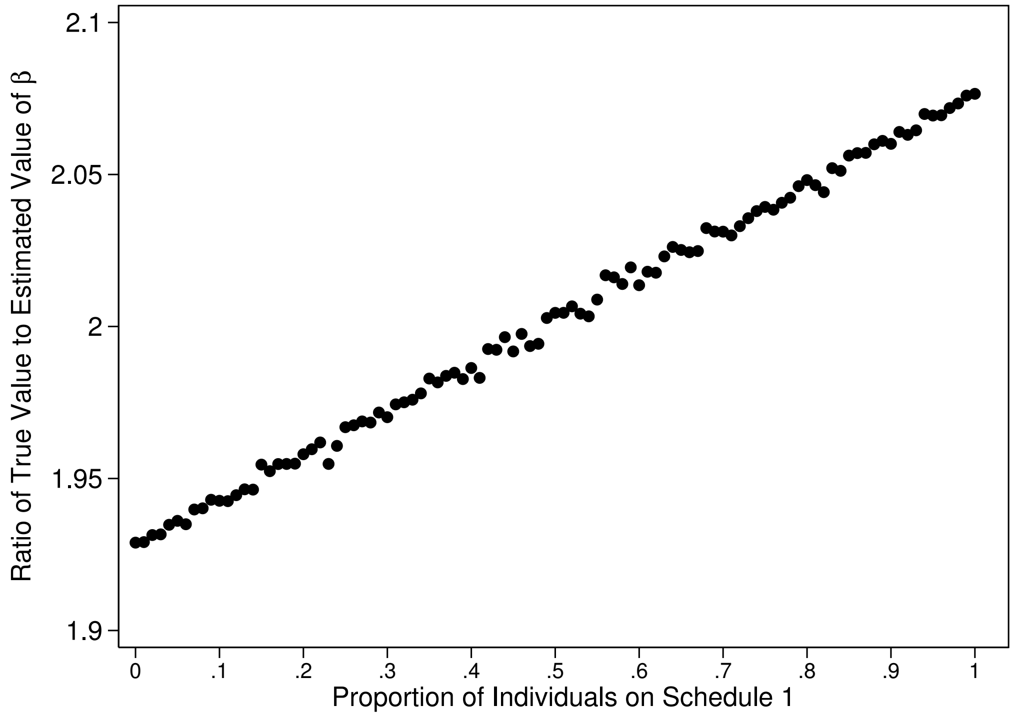

Note.— This figure plots estimates of the bias ratio, β^R, by the proportion, λ, of individuals on Schedule 1 in increasing increments of 0.01. Using simulations, I estimate the bias ratio for each λ by dividing the true value of β by the estimated value of $\hat{\beta}$ using the mismeasured indicator from the main specification in Equation 1.6. Additional details regarding the simulation exercise are in Appendix A.2.3.

A.3. Additional Robustness Checks

The empirical strategy outlined in Section 1.4 takes several steps to reduce the possibility that the estimates I find are not spurious. In this section, I present estimates from a series of variants to the specification in Equation 1.6. For the first alternative specification, I estimate the following equation

$$\Delta log C_{it} = \beta * \mathbf{1}_{\{t-1 \in S\}} + \theta_{it}'\alpha + \gamma_t + \epsilon_{it} \tag{A.7}$$

Table A.3.1 shows the results from this specification for the different aggregate expenditure groups. Here, the parameter of interest β measures the percentage change in expenditure growth following a month with three paychecks. Column 1 of Table A.3.1 shows that there is still a significant percentage increase in total spending following three paycheck months. The results are slightly less significant than the main specification.

In the main specification, the key independent variable was an indicator, $\mathbf{1}_{\{t-1 \in S\}}$, for whether a month follows a three paycheck month. While this specification lends itself to easier interpretation, especially in light of the fact that the response is driven largely by durables, it does not take full advantage of the variation available in the data set. For the next specification, I leverage variation in both the timing and the *size* of the the third paychecks to estimate spending responses. Specifically, I estimate the following equation

$$\Delta C_{it} = \beta * \Delta Y_{it} + \theta_{it}'\alpha + \gamma_t + \epsilon_{it} \tag{A.8}$$

where ΔY_{it} is the monthly change in income. Note that when looking at months with only two paychecks, the monthly change in income is simply zero. However, for months following three paycheck months, there is an additional paycheck available for spending, so the monthly change in income is equal to the size of the paycheck. As stated before, the third paycheck arrives at the end of the last week of the month and so I consider the income from that third paycheck available for spending in the *following* month. Table A.3.2 presents the results from this specification. Column 1 shows that for every \$1 from the third paycheck, 17.4 cents are spent. Given that the average paycheck size for a bi-weekly households is around \$1669, this translates to roughly \$290 in increased spending following three paycheck months. This estimate is similar to the one from our original specification in Table 2. These estimates are slightly more precise because they can exploit the additional variation in the size of paychecks across households. Column 3 shows a statistically significant but small response in non-durable spending that translates roughly to a \$22 increase in non-durable spending following three paycheck months.

TABLE A.3.1: Response to Extra Paychecks by Consumption Categories Using Log Changes in Consumption

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: $\Delta log\ C_t$					
$\mathbf{1}_{\{t-1 \in S\}}$	0.024**	0.033	0.005	-0.003	-0.002
	(0.012)	(0.022)	(0.006)	(0.004)	(0.004)
Age	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Δ Children	0.050	0.117	-0.007	0.038	0.062
	(0.100)	(0.163)	(0.038)	(0.027)	(0.040)
Δ Adults	0.011	-0.038	0.011	0.017	0.107**
	(0.105)	(0.175)	(0.045)	(0.047)	(0.051)
R-squared	0.005	0.002	0.038	0.007	0.007
N	13,707	13,622	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the log month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The indicator $\mathbb{1}_{\{t-1 \in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE A.3.2: Response to Extra Paychecks by Consumption Categories Using Level Changes in Income

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
ΔY_t	0.160***	0.157***	0.003	-0.001	-0.001
	(0.058)	(0.057)	(0.006)	(0.003)	(0.002)
Age	1.736	1.770	-0.034	-0.070	0.063
	(1.568)	(1.543)	(0.270)	(0.151)	(0.077)
Δ Children	318.921	329.226	-10.310	38.728	30.745
	(649.942)	(612.272)	(70.312)	(26.734)	(19.552)
Δ Adults	370.394	386.580	-16.193	17.552	51.503
	(657.372)	(619.442)	(78.217)	(57.838)	(42.091)
R-squared	0.003	0.003	0.034	0.006	0.005
N	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate 2SLS regression run at the household-month level, instrumenting for the change in income ΔY_t using the indicator for the previous month being a three paycheck month. The variable ΔY_t equals the 2010 dollar amount of the head of household's last gross pay if the previous month was a three paycheck month (if three paychecks of income are available in the present month t) and zero otherwise. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The main dependent variable in all specifications is the month-to-month dollar change in wage income. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

A.4. Heterogeneity in Spending Responses

TABLE A.4.1: Response to Extra Paychecks by Consumption Categories by Marital Status

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
$\mathbf{1}_{\{t-1 \in S\}}$	124.704	113.958	10.744	-6.207	-0.268
	(90.540)	(89.399)	(10.236)	(5.902)	(2.774)
$\mathbf{1}_{\{t-1 \in S\}} * \mathbf{1}_{\{Unconstrained\}}$	351.854*	367.430*	-15.595	9.698	-1.761
	(200.538)	(199.146)	(20.355)	(11.600)	(5.548)
Age	2.201	2.240	-0.038	-0.077	0.058
	(1.564)	(1.536)	(0.270)	(0.151)	(0.077)
Δ Children	320.608	334.519	-13.918	37.709	29.969
	(649.363)	(611.908)	(70.704)	(26.693)	(19.537)
Δ Adults	371.466	386.338	-14.878	18.329	51.867
	(661.909)	(624.505)	(78.205)	(57.803)	(42.059)
R-squared	0.003	0.002	0.035	0.006	0.006
N	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The indicator $\mathbb{1}_{\{t-1 \in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). The estimate for this indicator represents the response to extra paychecks for households who are constrained (the omitted category). Constrained households are defined as those whose heads are married; likewise, unconstrained households are those whose heads are unmarried. The indicator $\mathbb{1}_{\{t-1 \in S\}} \mathbb{1}_{\{Unconstrained\}}$ thus give the response to extra paychecks for households that are unconstrained relative to the response for households who are constrained. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE A.4.2: Response to Extra Paychecks by Consumption Categories by Race

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
$\mathbf{1}_{\{t-1 \in S\}}$	232.546	231.541	1.004	-18.156*	-0.577
	(155.319)	(154.264)	(18.382)	(10.625)	(4.819)
$\mathbf{1}_{\{t-1 \in S\}} * \mathbf{1}_{\{Unconstrained\}}$	1.859	-0.904	2.753	17.787	-0.983
	(186.930)	(185.782)	(21.029)	(12.184)	(5.730)
Age	2.063	2.116	-0.052	-0.146	0.045
	(1.621)	(1.594)	(0.276)	(0.151)	(0.078)
Δ Children	999.758	932.869	66.889	29.049	21.719
	(706.893)	(710.349)	(50.138)	(31.638)	(21.806)
Δ Adults	676.847	674.351	2.492	7.338	44.035
	(779.964)	(745.447)	(82.971)	(64.110)	(47.217)
R-squared	0.003	0.002	0.035	0.006	0.005
N	13,013	13,013	13,013	13,013	13,013
Month and Year FEs	Y	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The indicator $\mathbb{1}_{\{t-1 \in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). The estimate for this indicator represents the response to extra paychecks for households who are constrained (the omitted category). Constrained households are defined as those for whom the age of the household is less than the median age for households heads (age 39); likewise, unconstrained households are those for whom the age of the household is greater than or equal to the median age for households heads. The indicator $\mathbb{1}_{\{t-1 \in S\}} \mathbb{1}_{\{Unconstrained\}}$ thus give the response to extra paychecks for households that are unconstrained relative to the response for households who are constrained. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE A.4.3: Response to Extra Paychecks by Consumption Categories by Gender

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
$\mathbf{1}_{\{t-1\in S\}}$	87.660	87.669	-0.010	-6.904	-2.273
	(115.533)	(113.971)	(12.363)	(6.869)	(3.263)
$\mathbf{1}_{\{t-1\in S\}} * \mathbf{1}_{\{Unconstrained\}}$	364.112**	354.306**	9.788	9.353	2.758
	(178.548)	(177.315)	(18.416)	(10.593)	(5.083)
Age	2.311	2.320	-0.008	-0.067	0.064
	(1.562)	(1.533)	(0.272)	(0.152)	(0.077)
Δ Children	320.664	332.205	-11.548	38.494	30.629
	(649.244)	(611.975)	(70.390)	(26.745)	(19.580)
Δ Adults	366.682	382.968	-16.292	17.631	51.529
	(659.887)	(622.104)	(78.215)	(57.892)	(42.157)
R-squared	0.003	0.002	0.034	0.006	0.005
N	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The indicator $\mathbb{1}_{\{t-1\in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). The estimate for this indicator represents the response to extra paychecks for households who are constrained (the omitted category). Constrained households are defined as those whose head of household is female; likewise, unconstrained households are those for whom the head of household is male. The indicator $\mathbb{1}_{\{t-1\in S\}}\mathbb{1}_{\{Unconstrained\}}$ thus give the response to extra paychecks for households that are unconstrained relative to the response for households who are constrained. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

TABLE A.4.4: Response to Extra Paychecks by Consumption Categories by Housing Ownership

	(1) Total	(2) Durable	(3) Non-durable	(4) Strictly ND	(5) Food
Dependent variable: ΔC_t					
$\mathbf{1}_{\{t-1\in S\}}$	166.524*	152.753	13.769	-10.746*	-0.408
	(93.785)	(92.870)	(10.442)	(6.171)	(3.063)
$\mathbf{1}_{\{t-1\in S\}} * \mathbf{1}_{\{Unconstrained\}}$	204.629	224.064	-19.451	17.907*	-1.138
	(181.492)	(179.973)	(18.818)	(10.636)	(5.096)
Age	1.790	1.985	-0.195	-0.150	0.001
	(1.591)	(1.562)	(0.282)	(0.162)	(0.080)
Δ Children	319.892	331.610	-11.724	38.724	30.389
	(649.488)	(611.858)	(70.606)	(26.713)	(19.596)
Δ Adults	363.040	378.751	-15.717	17.500	51.647
	(659.908)	(622.064)	(78.242)	(57.985)	(42.188)
R-squared	0.003	0.002	0.034	0.006	0.006
N	13,707	13,707	13,707	13,707	13,707
Month and Year FEs	Y	Y	Y	Y	Y

NOTE.—Each column reports estimates from a separate OLS regression run at the household-month level. Standard errors are clustered by household and are reported in parentheses. All specifications in this table are estimated using the bi-weekly sample. The dependent variable in all specifications is the (2010) month-to-month dollar change in consumption. The category of consumption for each specification is listed at the head of each column and includes: total, durable, non-durable, strictly non-durable, and food. The indicator $\mathbb{1}_{\{t-1\in S\}}$ equals one if the previous month was a three paycheck month (if three paychecks of income are available in the present month t). The estimate for this indicator represents the response to extra paychecks for households who are constrained (the omitted category). Constrained households are defined as those who are homeowners; likewise, unconstrained households are those who are renters. The indicator $\mathbb{1}_{\{t-1\in S\}}\mathbb{1}_{\{Unconstrained\}}$ thus give the response to extra paychecks for households that are unconstrained relative to the response for households who are constrained. In addition to month and year fixed effects, all four specifications include age of the head of household, changes in the number of children, and changes in the number of adults as controls. Significance is denoted by: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

A.5. Proofs

Lemma 1. *Assume that individuals follow a t-budgeting heuristic and there is misalignment between the timing of consumption and the timing of pay. Then,*

(i) Individuals have overly optimistic beliefs in periods with atypical income.

(ii) Individuals have overly pessimistic beliefs in periods with typical income.

Proof. From Equation 1.3, true lifetime wealth is given by the expression

$$W_t = X_t + \sum_{i=1}^{\infty} \frac{y_{t+i}}{R^i}$$

At time t, the individual expects her remaining lifetime wealth to be given by the expression

$$\begin{aligned} W_t^E &= X_t + \sum_{i=1}^{\infty} \frac{y_{t,t+i}^E}{R^i} \\ &= X_t + \sum_{i=1}^{\infty} \frac{\alpha y_t + (1-\alpha) y_{t+i}}{R^i} \\ &= X_t + \frac{\alpha y_t}{R-1} + (1-\alpha) \sum_{i=1}^{\infty} \frac{y_{t+i}}{R^i}. \end{aligned} \tag{A.9}$$

Note that for rational individuals ($\alpha = 0$), expected wealth simply equals true wealth. For all other individuals ($\alpha > 0$), subtracting the first expression from the second then implies that

$$W_t^E - W_t \gtreqless 0 \quad \text{if} \quad y_t \gtreqless (R-1) \sum_{i=1}^{\infty} \frac{y_{t+i}}{R^i} \tag{A.10}$$

where $y_{t+i} = Y + bY \cdot \mathbb{1}_{\{(t-1+i) \bmod n = 0\}}$. Equation A.10 can further be rearranged in the form

$$W_t^E - W_t \gtreqless 0 \quad \text{if} \quad y_t - Y \gtreqless (R-1) \sum_{i=1}^{\infty} \frac{bY}{R^i} \cdot \mathbb{1}_{\{(t-1+i) \bmod n = 0\}} \tag{A.11}$$

Given that $n > 1$ and $b > 0$, we can bound the right-hand side of Equation A.11 as follows

$$bY > (R-1) \sum_{i=1}^{\infty} \frac{bY}{R^i} \cdot \mathbb{1}_{\{(t-1+i) \bmod n = 0\}} > 0 \tag{A.12}$$

Therefore,

$$
\begin{aligned}
W_t^E - W_t > 0 \quad &\text{if} \quad y_t = (1+b)Y, \quad \text{and} \\
W_t^E - W_t < 0 \quad &\text{if} \quad y_t = Y
\end{aligned}
$$

Individuals are overly optimistic in periods with atypical income and overly pessimistic in periods with typical income. □

Lemma 2. *Assume that individuals follow a t-budgeting heuristic and there is misalignment between the timing of consumption and the timing of pay.*

(i) Consumption is weakly greater than the optimal consumption level under correct beliefs in periods with atypical income.

(ii) Consumption is weakly less than the optimal consumption level under correct beliefs in periods with typical income.

Proof. This follows directly from Equation 1.5 and Lemma 1. □

Proposition 1. *In the presence of a t-budgeting heuristic and misalignment between the timing of consumption and the timing of income, the marginal propensity to consume out of additional income bY in atypical income periods is $(1 - (\delta R^{1-\rho})^{1/\rho})\frac{R}{R-1}\alpha$ and is increasing in α.*

Proof. Recall from Equation A.9, the individual's expected remaining lifetime wealth in period t is

$$W_t^E = X_t + \frac{\alpha y_t}{R-1} + (1-\alpha)\sum_{i=1}^{\infty}\frac{y_{t+i}}{R^i}.$$

Looking forward one period, expected wealth in period $t+1$ is

$$W_{t+1}^E = R(X_t - c_t) + y_{t+1} + \frac{\alpha y_{t+1}}{R-1} + (1-\alpha)\sum_{i=1}^{\infty}\frac{y_{t+1+i}}{R^i}.$$

Dividing both sides by R and then rearranging gives

$$
\begin{aligned}
\frac{W_{t+1}^E}{R} &= X_t - c_t + \alpha\frac{y_{t+1}}{R} + \frac{1}{R}\cdot\alpha\frac{y_{t+1}}{R-1} + (1-\alpha)\sum_{i=1}^{\infty}\frac{y_{t+i}}{R^i} \\
&= W_t^E - c_t + \alpha\frac{y_{t+1}-y_t}{R-1}.
\end{aligned}
$$

Multiplying both sides by R and then combining with Equation 1.5, expected remaining lifetime wealth is

$$W_{t+1}^E = (\delta R)^{1/\rho} W_t^E + \frac{R}{R-1}\alpha(y_{t+1} - y_t)$$

and period $t+1$ consumption is

$$c_{t+1} = (\delta R)^{1/\rho} c_t + (1 - (\delta R^{1-\rho})^{1/\rho})\frac{R}{R-1}\alpha(y_{t+1} - y_t)$$

We can thus see that receiving additional income bY leads to $(1 - (\delta R^{1-\rho})^{1/\rho})\frac{R}{R-1}\alpha(bY)$ in additional consumption so the marginal propensity to consume out of the additional income is simply MPC $= (1 - (\delta R^{1-\rho})^{1/\rho})\frac{R}{R-1}\alpha$. Given our assumptions about individual impatience and the gross interest rate, it follows that $\frac{\partial \text{MPC}}{\partial \alpha} > 0$.

□

Proposition 2. *For all α, the presence of borrowing constraints of the form $X_t - c_t \geq 0$ leads to weakly larger increases in consumption in periods with atypical income than in the absence of borrowing constraints.*

Proof. Recall from the proof of Proposition 1 that expected wealth in period $t+1$ can be written as

$$\frac{W_{t+1}^E}{R} = W_t^E - c_t + \alpha\frac{y_{t+1} - y_t}{R-1}.$$

Suppose that the borrowing constraint binds so that

$$c_t = \min\{X_t, (1 - (\delta R^{1-\rho})^{1/\rho})W_t^E\} = X_t.$$

Let c_t^* be the preferred consumption level under no borrowing constraints in period t, and let $s_t = c_t^* - X_t \geq 0$ be the difference between this preferred consumption level and cash-on-hand. Assuming the borrowing constraint binds in period t, expected wealth in period $t+1$ can be rewritten as

$$\frac{W_{t+1}^E}{R} = W_t^E - (c_t^* - s_t) + \alpha\frac{y_{t+1} - y_t}{R-1}.$$

Multiplying both sides by R and then combining with Equation 1.5, expected remaining lifetime wealth and consumption in period $t+1$ are

$$W_{t+1}^E = (\delta R)^{1/\rho} W_t^E + s_t + \frac{R}{R-1}\alpha(y_{t+1} - y_t)$$
$$c_{t+1} = (\delta R)^{1/\rho} c_t + (1 - (\delta R^{1-\rho})^{1/\rho})[\frac{s_t}{y_{t+1} - y_t} + \frac{R}{R-1}\alpha](y_{t+1} - y_t)$$

Note that $s_t = 0$ in the absence of borrowing constraints. Given that s_t is weakly positive and $y_{t+1} - y_t = bY$ in months with atypical income, it follows that the presence of borrowing constraints leads to weakly larger increase in consumption in periods with atypical income than would otherwise be implied by α in the absence of borrowing constraints. □

Proposition 3. *Let wealth entering period 0 be distributed according to the CDF $G(W_0^E)$ and density $g(W_0^E)$, where the density $g(\cdot)$ depends on the joint distribution of (X_0, α). For a given initial level of committed consumption d_0, the probability that an individual discretely adjusts her committed consumption in period 1 ($d_1 \neq d_0$) is increasing in α.*

Proof. First note that within-period preferences are Cobb-Douglas between adjustable and committed consumption. The optimal level of consumption of each type of good is proportional to the individuals expected lifetime wealth in that period. Thus, in period 0, the individual's optimal consumption is given by

$$\{c_0, d_0\} = \left\{(1-\gamma)(1 - (\delta R^{1-\rho})^{1/\rho})W_0^E, \gamma(1 - (\delta R^{1-\rho})^{1/\rho})W_0^E\right\}$$

where expected lifetime wealth in period 0 is

$$W_0^E = X_0 + y_0 + \frac{\alpha y_0}{R-1} + (1-\alpha)\sum_{i=1}^{\infty} \frac{y_{t+i}}{R^i},$$

From Chetty and Szeidl, for a given d_0, there exists $s_1 < S_1$ such that the optimal policy is to move when $W_1^E \notin (s_1, S_1)$. Then for a given α, there exists $m = S_1 - (\alpha bY - kd_0) < \infty$ such that individuals with $W_0^E > m$ have $W_1^E = W_0^E + \alpha bY - kd_0 > S_1$. For individuals with $W_0^E > m$, it is therefore optimal to adjust her level of committed consumption. For a given α, the probability that an individual has wealth entering into period 1 sufficiently high that it is optimal for her to move is then given by $1 - G(m(\alpha))$. To show that this probability is increasing in α, we take the derivative of our expression with respect to α,

$$\frac{\partial(1 - G(m(\alpha))}{\partial \alpha} = -\frac{dG(W)}{dW} \cdot m'(\alpha) = -g(W)(-bY) > 0.$$

Thus the probability of an individual adjusting her level of committed consumption is strictly increasing in α. $\square$

BIBLIOGRAPHY

D. Aaronson, S. Agarwal, and E. French. The spending and debt response to minimum wage hikes. *American Economic Review*, 102(7):3111–39, 2012.

J. Abaluck and J. Gruber. Choice inconsistencies among the elderly: Evidence from plan choice in the medicare part d program. *American Economic Review*, 101(4):1180–1210, 2011.

D. Acland and M. Levy. Naivet, projection bias, and habit formation in gym attendance, 2013. Working Paper.

W. Adams, L. Einav, and J. Levin. Liquidity constraints and imperfect information in subprime lending. *American Economic Review*, 99(1):49–84, 2009.

S. Agarwal, C. Liu, and N. S. Souleles. The reaction of consumer spending and debt to tax rebates: Evidence from consumer credit data. *Journal of Political Economy*, 115(6): 986–1019, 2007.

D. J. Aigner. Regression with a binary independent variable subject to errors of observation. *Journal of Econometrics*, 1:49–60, 1973.

J. Andreoni. Philanthropy. In S. Kolm and J. M. Ythier, editors, *Handbook of the Economics of Giving, Altruism and Reciprocity*, volume 2, pages 1201–69. North-Holland, The Netherlands, 2006.

J. Andreoni and B. D. Bernheim. Social image and the 50 - 50 norm: A theoretical and experimental analysis of audience effects. *Econometrica*, 77(5):1607–36, 2009.

J. Andreoni and R. Petrie. Public goods experiments without confidentiality: A glimpse into fund-raising. *Journal of Public Economics*, 88:1605–23, 2004.

J. Angrist and V. Lavy. The effects of high stakes high school achievement awards: Evidence from a randomized trial. *American Economic Review*, 99(4):1384–1414, 2009.

J. Angrist, D. Lang, and P. Oreopoulos. Incentives and services for college achievement: Evidence from a randomized trial. *American Economic Journal: Applied Economics*, 1 (1):136–63, 2009.

D. Ariely, A. Bracha, and S. Meier. Doing good or doing well? image motivation and monetary incentives in behaving prosocially. *American Economic Review*, 99(1):544–55, 2009.

O. Attanasio, P. K. Goldberg, and E. Kyriazidou. Credit constraints in the market for consumer durables: Evidence from micro data on car loans. *International Economic Review*, 49(2):401–36, 2008.

N. Barberis, A. Shleifer, and R. Vishny. A model of investor sentiment. *Journal of Financial Economics*, 49(3):307–43, 1998.

S. Barbieri and D. A. Malueg. Increasing fundraising success by decreasing donor choice. *Journal of Public Economics*, forthcoming.

G. S. Becker. A theory of social interactions. *Journal of Political Economy*, 82(6):1063–93, 1974.

R. Bénabou and J. Tirole. Incentives and prosocial behavior. *American Economic Review*, 96(5):1652–78, 1996.

S. Bernatzi and R. Thaler. Heuristics and biases in retirement savings behavior. *Journal of Economic Perspectives*, 21(3):81–104, 2007.

E. P. Bettinger. Paying to learn: The effect of financial incentives on elementary school test scores. *Review of Economics and Statistics*, 94(3):686–98, 2012.

D. A. Black, M. C. Berger, and F. A. Scott. Bounding parameter estimates with nonclassical measurement error. *Journal of the American Statistical Association*, 95(451):739–48, 2000.

I. Bohnet and B. S. Frey. The sound of silence in prisoner's dilemma and dictator games. *Journal of Economic Behavior Organization*, 38(1):43–57, 1999.

A. Bracha and L. Vesterlund. How low can you go? charity reporting when donations signal income and generosity, 2013. Working Paper.

M. Browning and D. Collado. The response of expenditures to anticipated income changes: Panel data estimates. *American Economic Review*, 91(3):681–92, 2001.

M. Browning and T. F. Crossley. The life-cycle model of consumption and saving. *Journal of Economic Perspectives*, 15(3):3–22, 2001.

M. Browning and A. Lusardi. Household saving: Micro theories and macro facts. *Journal of Economic Literature*, 34(4):1797–1855, 1996.

R. Bnabou and J. Tirole. Intrinsic and extrinsic motivation. *Review of Economic Studies*, 70(3):489–520, 2003.

J. Y. Campbell and G. N. Mankiw. Consumption, income and interest rates: Reinterpreting the time series evidence. *NBER Macroeconomics Annual: 1989*, pages 185–216, 1989.

J. Y. Campbell and G. N. Mankiw. The response of expenditures to anticipated income changes: Panel data estimates. *American Economic Review*, 91(3):681–92, 2001.

A. Caplin, M. Dean, and D. Martin. Search and satisficing. *American Economic Review*, 101(7):2899–2922, 2011.

E. Cartwright and A. Patel. How category reporting can improve fundraising. *Journal of Economic Behavior Organization*, 87:73–90, 2013.

G. Charness and U. Gneezy. Incentives to exercise. *Econometrica*, 77(3):909–31, 2009.

R. Chetty and A. Szeidl. Consumption commitments and risk preferences. *The Quarterly Journal of Economics*, 122(2):831–77, 2007.

J. J. Choi, D. Laibson, and B. C. Madrian. $100 bills on the sidewalk: Suboptimal investment in 401(k) plans. *The Review of Economics and Statistics*, 93(3):748–63, 2011.

Council for Aid to Education. Voluntary support of education 2013 survey, 2014.

J. Dana, D. M. Cain, and R. M. Dawes. What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes*, 100:193–201, 2006.

A. Deaton. Saving and liquidity constraints. *Econometrica*, 59(5):1221–48, 1991.

E. L. Deci. Effects of externally mediated rewards on intrinsic motivation. *Journal of Personality and Social Psychology*, 18(1):105–15, 1971.

E. L. Deci, R. Koestner, and R. M. Ryan. Extrinsic rewards and intrinsic motivation in education: Reconsidered once again. *Review of Educational Research*, 71(1):1–27, 2001.

S. DellaVigna. Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–72, 2009.

T. Ellingsen and M. Johannesson. Pride and prejudice: The human side of incentive theory. *American Economics Review*, 98(3):990–1008, 2008.

A. J. Elliot and J. M. Harackiewicz. Goal setting, achievement orientation, and intrinsic motivation: A meditational analysis. *Journal of Personality and Social Psychology*, 66 (5):968–80, 1994.

A. Filippin and J. C. van Ours. Run for fun: Intrinsic motivation and physical performance, 2012. Working Paper.

B. S. Frey. Tertium datur: Pricing, regulation and intrinsic motivation. *Kyklos*, 45(2): 161–84, 1992.

B. S. Frey. On the relationship between intrinsic and extrinsic work motivation. *International Journal of Industrial Organization*, 15(4):427–39, 1997.

M. Friedman. *A Theory of the Consumption Function*. Princeton University Press, 1957.

R. Fryer. Financial incentives and student achievement: Evidence from randomized trials. *Quarterly Journal of Economics*, 126(4):1755–98, 2011.

X. Gabaix, D. Laibson, G. Moloche, and S. Weinberg. Costly information acquistion: Experimental analysis of a boundedly rational model. *American Economic Review*, 96 (4):1043–68, 2006.

T. Gilovich, D. Griffin, and D. Kahneman, editors. *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press, 2002.

Giving USA. Giving usa 2013 report, 2013.

A. Glazer and K. A. Konrad. A signaling explanation for charity. *American Economic Review*, 86(4):1019–28, 1996.

U. Gneezy and A. Rustichini. Pay enough or dont pay at all. *Quarterly Journal of Economics*, 115(3):791–810, 2000.

U. Gneezy, S. Meier, and P. Rey-Biel. When and why incentives (dont) work to modify behavior. *Journal of Economic Perspectives*, 25(4):191–210, 2011.

R. Hai, D. Kreuger, and A. Postlewaite. On the welfare cost of consumption fluctuations in the presence of memorable goods. January 2013. NBER Working Paper No. 19386.

R. E. Hall and F. S. Mishkin. The sensitivity of consumption to transitory income: Estimates from panel data on households. *Econometrica*, 50(2):461–81, 1982.

W. T. Harbaugh. The prestige motive for making charitable transfers. *American Economic Review*, 88(2):277–82, 1998a.

W. T. Harbaugh. What do donations buy? a model of philanthropy based on prestige and warm glow. *Journal of Public Economics*, 67(2):269–84, 1998b.

J. S. Hastings and J. M. Shapiro. Fungibility and consumer choice: Evidence from commodity price shocks. *The Quarterly Journal of Economics*, Forthcoming, 2013.

J. Hattie and H. Timperley. The power of feedback. *Review of Educational Research*, 77 (1):81–112, 2007.

O. Heffetz and R. H. Frank. Preferences for status: Evidence and economic implications. In M. O. J. Jess Benhabib and A. Bisin, editors, *Handbook of Social Economics*, volume 1A, page 6991. North-Holland, The Netherlands, 2011.

H. Hollander. A social exchange approach to voluntary cooperation. *American Economic Review*, 80(5):1157–67, 1990.

J. Holmes. Prestige, charitable deductions and other determinants of alumni giving: Evidence from a highly selective liberal arts college. *Economics of Education Review*, 28(1): 18–28, 2009.

C.-T. Hsieh. Do consumers react to anticipated income changes? evidence from the alaska permanent fund. *American Economic Review*, 93(1):397–405, 2003.

D. Huffman and M. Barenstein. A monthly struggle for self-control? hyperbolic discounting, mental accouting, and the fall in consumption between paydays. Working Paper, December 2005.

N. J. Ireland. On limiting the market for status signals. *Journal of Public Economics*, 53 (1):91–110, 1994.

T. Jappelli. Who is credit-constrained in the us economy? *Quarterly Journal of Economics*, 105:219–34, 1990.

T. Jappelli and L. Pistaferri. The consumption response to income changes. *Annual Review of Economics*, 2:479–506, 2010.

T. Jappelli, J.-S. Pischke, and N. S. Souleles. Testing for liquidity constraints in euler equations with complementary data sources. *The Review of Economics and Statistics*, 80:251–62, 1998.

D. S. Johnson, J. A. Parker, and N. S. Souleles. Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610, 2006.

D. S. Johnson, J. A. Parker, and N. S. Souleles. The response of consumer spending to rebates during the expansion: Evidence from the 2003 child tax credit. Working Paper, April 2009.

D. Karlan and M. A. McConnell. Hey look at me: The effect of giving circles on giving, 2013. Working Paper.

B. Kőszegi and A. Szeidl. A model of focusing in economic choice. *The Quarterly Journal of Economics*, 128(1):53–104, 2013.

N. Lacetera, D. G. Pope, and J. R. Sydnor. Heuristic thinking and limited attention in the car market. *American Economic Review*, 102(5):2206–36, 2012.

D. Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–78, 1997.

J. Li and Y. E. Riyanto. Category reporting in charitable giving: An experimental analysis, 2009. Working Paper.

S. Linardi and M. A. McConell. No excuses for good behavior: Volunteering and the social environment. *Journal of Public Economics*, 95(5-6):445–54, 2011.

G. Loewenstein, T. O'Donoghue, and M. Rabin. Projection bias in predicting future utility. *The Quarterly Journal of Economics*, 118:1209–48, 2003.

A. Lusardi. Permanent income, current income, and consumption: Evidence from two panel data sets. *Journal of Business and Economic Statistics*, 14(1):81–90, 1996.

J. Meer. Brother, can you spare a dime? peer pressure in charitable solicitation. *Journal of Public Economics*, 95(7-8):926–41, 2011.

J. Meer and H. S. Rosen. The impact of athletic performance on alumni giving: An analysis of microdata. *Economics of Education Review*, 28(3):287–94, 2008a.

J. Meer and H. S. Rosen. Does genorsity beget genoristy? alumni giving and undergraduate financial aid. *Economics of Education Review*, 28(3):287–94, 2008b.

J. Meer and H. S. Rosen. Altruism and the child cycle of alumni donations. *American Economic Journal: Economic Policy*, 1(1):258–86, 2009.

J. Meer and H. S. Rosen. The abcs of charitable solicitation. *Journal of Public Economics*, 95(5-6):363–71, 2011.

K. L. Milkman and J. Beshears. Mental accounting and small windfalls: Evidence from an online grocer. *Journal of Economic Behavior and Organization*, 71(2):384–94, 2009.

F. Modigliani. The role of intergenerational transfers and life cycle saving in the accumulation of wealth. *Journal of Economic Perspectives*, 2(2):15–40, 1988.

T. O'Donoghue and M. Rabin. Doing it now or later. *American Economic Review*, 89(1): 103–24, 1999.

M. Olson. *The Logic of Collective Action*. Harvard University Press, Harvard, 1965.

J. A. Parker. The reaction of household consumption to predictable changes in social security taxes. *American Economic Review*, 89(4):959–73, 1999.

J. A. Parker, N. S. Souleles, D. S. Johnson, and R. McClelland. Consumer spending and the economic stimulus payments of 2008. January 2011. NBER Working Paper No. 16684.

C. Parsons and E. D. V. Wesep. The timing of pay. *Journal of Financial Economics*, 109 (2):373–97, 2013.

C. H. Paxson. Consumption and income seasonality in thailand. *Journal of Political Economy*, 101(1):39–7, 1993.

E. S. Phelps and R. A. Pollak. On second-best national saving and game-equilibrium growth. *Review of Economic Studies*, 35(2):185–99, 1968.

D. Pope and U. Simonsohn. Round numbers as goals: Evidence from baseball, sat takers, and the lab. *Psychological Science*, 22(1):71–9, 2011.

D. G. Pope and M. E. Schweitzer. Is tiger woods loss averse? persistent bias in the face of experience, competition, and high stakes. *American Economic Review*, 101(1):129–57, 2011.

J. Potters, M. Sefton, and L. Vesterlund. After you - endogenous sequencing in voluntary contribution games. *Journal of Public Economics*, 89(8):1499–1419, 2005.

C. Prendergast. Intrinsic motivation and incentives. *American Economic Review*, 98(2): 201–05, 2008.

M. Rabin and J. L. Schrag. First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics*, 114(1):37–82, 1999.

M. Rege and K. Telle. The impact of social approval and framing on cooperation in public good situations. *Journal of Public Economics*, 88(7-8):1625–44, 2004.

J. A. Roberts, I. Hann, and S. A. Slaughter. Understanding the motivations, participation, and performance of open source software developers: A longitudinal study of the apache projects. *Management Science*, 52(7):984–99, 2006.

D. Rondeau and J. A. List. Matching and challenge gifts to charity: Evidence from laboratory and natural field experiments. *Experimental Economics*, 11(3):253–67, 2008.

E. Rosch. Cognitive reference points. *Cognitive Psychology*, 7(4):532–47, 1975.

D. E. Runkle. Liquidity constraints and the permanent income hypothesis: Evidence from panel data. *Journal of Monetary Economics*, 27(1):73–98, 1991.

R. M. Ryan. Control and information in the intrapersonal sphere: An extension of cognitive evaluation theory. *Journal of Personality and Social Psychology*, 43(3):450–61, 1982.

R. M. Ryan and E. L. Deci. Intrinsic and extrinsic motivations: Classic denitions and new directions. *Contemporary Educational Psychology*, 25(1):54–67, 2000.

A. Samak and R. Sheremeta. Recognizing contributors: An experiment on public goods. *Experimental Economics*, forthcoming.

J. M. Shapiro. Is there a daily discount rate? evidence from the food stamp nutrition cycle. *Journal of Public Economics*, 89(2-3):303–25, 2005.

M. D. Shapiro and J. Slemrod. Consumer response to the timing of income: Evidence from a change in tax witholding. *American Economic Review*, 85(1):274–83, 1995.

J. Shea. Union contracts and the life-cycle/permanent-income hypothesis. *American Economic Review*, 85(1):186–200, 1995.

H. M. Shefrin and R. H. Thaler. The behavioral life-cycle hypothesis. *Economic Inquiry*, 26(4):609–43, 1988.

H. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1): 99–118, 1955.

A. R. Soetevent. Anonymity in giving in a natural context a field experiment in 30 churches. *Journal of Public Economics*, 89(11-12):2301–2323, 2005.

N. S. Souleles. The response of household consumption to income tax refunds. *American Economic Review*, 89(4):947–58, 1999.

N. S. Souleles. Consumer responses to the reagan tax cuts. *Journal of Public Economics*, 85(1):99–120, 2002.

R. Steinberg. Does government spending crowd out donations? interpreting the evidence. *Annals of Public and Cooperative Economics*, 62(4):591–612, 1991.

M. Stephens, Jr. 3rd of that month: Do social security recipients smooth consumption between checks? *American Economic Review*, 93(1):406–22, 2003.

M. Stephens, Jr. Paycheque receipt and the timing of consumption. *The Economic Journal*, 116(513):680–701, 2006.

M. Stephens, Jr. The consumption response to predictable changes in discretionary income: Evidence from the repayment of vehicles loans. *Review of Economic Studies*, 90(2):241–52, 2008.

M. Stephens, Jr. and T. Unayama. The consumption response to seasonal income: Evidence from japanese public pension benefits. *American Economic Journal: Applied Economics*, 3(4):86–118, 2011.

R. H. Strotz. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3):165–80, 1956.

R. H. Thaler. Saving, fungibility, and mental accounts. *Journal of Economic Perspectives*, 4(1):193–205, 1990.

R. H. Thaler. Mental accounting matters. *Journal of Behavioral Decision Making*, 12(3): 183–206, 1999.

R. H. Thaler. Mental accounting and consumer choice. *Marketing Science*, 27(1):15–25, 2008.

R. H. Thaler and H. M. Shefrin. An economic theory of self-control. *Journal of Political Economy*, 89(2):392–410, 1981.

A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–31, 1974.

L. Vesterlund. The informational value of sequential fundraising. *Journal of Public Economics*, 87(2-4):627–57, 2003.

L. Vesterlund. Why do people give? In R. Steinberg and W. W. Powell, editors, *The Nonprofit Sector*, pages 568–87. Yale Press, 2 edition, 2006.

K. G. Volpp, A. B. Troxel, M. V. Pauly, H. A. Glick, A. Puig, D. A. Asch, R. Galvin, J. Zhu, F. Wan, J. DeGuzman, E. Corbett, J. Weiner, and J. Audrain-McGovern. A randomized trial of financial incentives for smoking cessation. *New England Journal of Medicine*, 350 (7):699–709, 2009.

D. W. Wilcox. Social security benefits, consumption expenditures, and the life cycle hypothesis. *Journal of Political Economy*, 97(2):288–304, 1989.

S. P. Zeldes. Consumption and liquidity constraints: an empirical investigation. *Journal of Political Economy*, 97(2):305–46, 1989.