

NO. 1208
OCTOBER 2026

Consumption, Savings, and Earnings Responses to Financial Windfalls

Rajashri Chakrabarti | Philippe d'Astous | Kory Kroft |
Sheisha Kulkarni | Vyacheslav Mikhed | Matthew J. Notowidigdo |
Sahil Raina | Barry Scholnick

Consumption, Savings, and Earnings Responses to Financial Windfalls

Rajashri Chakrabarti, Philippe d'Astous, Kory Kroft, Sheisha Kulkarni, Vyacheslav Mikhed, Matthew J. Notowidigdo, Sahil Raina, and Barry Scholnick

Federal Reserve Bank of New York Staff Reports, no. 1208

October 2026

<https://doi.org/10.59576/sr.1208>

Abstract

We estimate the causal effects of a financial windfall on consumption, savings, and wage earnings using linked administrative data on a large sample of Canadian lottery winners. Using separate linkages of lottery winners to income tax records and credit bureau data, we estimate how lottery winnings affect consumption measured through credit card spending, savings through financial asset accumulation and debt repayment, and wage earnings from tax records. We then examine how these responses vary across the income distribution. We find that high-income individuals allocated a larger share of lottery winnings to savings and leisure, while low-income individuals spent a larger share on consumption. We illustrate the policy relevance of these heterogeneous responses through three applications: designing fiscal stimulus programs, evaluating a Universal Basic Income (UBI), and determining optimal top tax rates on wage earnings and savings.

JEL classification: D12, D15, E21, E62, H2, J22

Key words: lottery, marginal propensity to consume, marginal propensity to save, marginal propensity to earn, consumption, saving

Chakrabarti: Federal Reserve Bank of New York (email: rajashri.chakrabarti@ny.frb.org). d'Astous: HEC Montreal (email: philippe.dastous@hec.ca). Kroft: Department of Economics, University of Toronto, NBER (email: kory.kroft@utoronto.ca). Kulkarni: Department of Finance, McIntire School of Commerce, University of Virginia (email: sk7nn@virginia.edu). Mikhed: Consumer Finance Institute, Federal Reserve Bank of Philadelphia (email: slava.mikhed@phil.frb.org). Notowidigdo: University of Chicago, Booth School of Business, NBER (email: noto@chicagobooth.edu). Raina: Alberta School of Business, University of Alberta (email: sraina@ualberta.ca). Scholnick: University of Alberta (email: bscholni@ualberta.ca). The authors gratefully acknowledge funding from the Social Sciences and Humanities Research Council of Canada. They are very grateful to Lee Lockwood, Jesse Shapiro, Stephen Zeldes, and conference and seminar participants at the NBER Public Economics, NBER Innovative Data in Household Finance, NBER Household Finance Across the Lifecycle, University of Virginia, University of Wisconsin–Madison, the Junior Household Finance Seminar, the Early Career Women in Finance Conference, American Economic Association Annual Meeting, Chicago Household Finance Conference, the Federal Reserve Bank of New York, the Federal Reserve Bank of Philadelphia, and Queen's University for their helpful comments. They sincerely appreciate the support from Roshani Agrawal, Peter McAlpine, Subhadip Sarkar, Lisa Weir, and ChengDa (Dennis) Zheng, who worked at the credit bureau that provided them with the data. They thank Stephen Claassen, Julian Haas, and Muhammad Ali Sajid for outstanding research assistance.

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the author(s) and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve Bank of Philadelphia, or the Federal Reserve System. Any errors or omissions are the responsibility of the author(s).

To view the authors' disclosure statements, visit
https://www.newyorkfed.org/research/staff_reports/sr1208.html.

1 Introduction

How do individuals' economic decisions change when their financial resources increase? Economists have long conjectured that the answer depends on an individual's initial position in the income distribution. Thomas Malthus observed that “the laboring poor . . . seem always to live from hand to mouth,” while John Maynard Keynes wrote that “the propensity to consume will be weaker in a rich community because it is well off, and stronger in a poor community because it is deprived” (Malthus, 1798; Keynes, 1936–2018). These observations suggest that identical financial windfalls may be allocated very differently across consumption, savings, and labor supply depending on recipients' initial resources.

This heterogeneity is central to the evaluation of policies that change individuals' disposable resources, including fiscal stimulus (Jappelli and Pistaferri, 2014; Auclert, 2019; Patterson, 2023), tax-and-transfer policies (Ferey, Lockwood, and Taubinsky, 2024; Hellwig and Werquin, 2026), and the transmission of monetary policy to the macro economy (Acharya, Challe, and Dogra, 2023; Ampudia et al., 2024; Crawley and Kuchler, 2023). Since these policies often target particular segments of the income distribution, their incidence depends on the behavioral responses of the affected population rather than on the average responses across the population as a whole. Moreover, translating responses to changes in wealth into the income effects required for evaluating these policies often depends critically on how individuals allocate resources over time. Thus, policy evaluation requires understanding both how behavioral responses vary across individuals and how wealth is translated into a flow of resources over time.

A growing literature has studied how households respond to financial windfalls (Jappelli and Pistaferri, 2010), but measuring how a given windfall is allocated across consumption, saving, and labor supply, how this allocation evolves over time, and how this allocation varies across the income distribution, remains challenging. This is because it requires an exogenous wealth shock paired with longitudinal data on all outcomes for a broad set of individuals. Several recent strands of research have made important progress on this question. First, Golosov et al. (2024) examine the

causal effects of lottery winnings on consumption and earnings across the income distribution using U.S. income tax data. However, they do not directly observe consumption and must instead infer it from tax records on labor and capital income, using the accounting identity that total spending equals income plus capital gains minus the change in wealth over the period. Eika, Mogstad, and Vestad (2020) show that consumption measures constructed using this approach can be sensitive to measurement error.¹ Second, Bartik et al. (2025) and Vivalt et al. (2025) study the effects of cash transfers on multiple outcomes using a sample of primarily low-income households. Although they directly measure consumption using surveys, they cannot examine how consumption, saving, and labor supply responses jointly vary across the income distribution.² Finally, several other studies examine heterogeneity in specific responses to financial windfalls, rather than how the windfall is allocated across consumption, saving, and earnings. Fagereng, Holm, and Natvik (2021) study heterogeneity in consumption and saving responses, while Jappelli and Pistaferri (2020) and Lewis, Melcangi, and Pilossoph (2026) study heterogeneity in spending responses.

In this paper, we address these challenges by creating two new administrative data linkages covering a large sample of Canadian lottery winners that together allow us to measure how an exogenous wealth shock is allocated across consumption, savings, and wage earnings across the income distribution. Credit bureau data from Equifax Canada provide novel, high-frequency measures of credit card spending and debt balances, while tax records from Statistics Canada provide annual measures of wage earnings and financial assets. Applying a common event-study design that treats the timing and amount of the lottery prize as random conditional on individual and time fixed effects, we estimate how an additional \$1 of wealth is allocated to consumption, debt repayment, asset accumulation, and wage earnings, and test whether this allocation varies

¹Specifically, Eika, Mogstad, and Vestad (2020) show that measures of household consumption expenditure based on income and wealth tax records can result in severe measurement errors unless complemented with individual-level data on financial and real estate transactions. Golosov et al. (2024) observe income tax records, but not wealth or asset transactions. Fagereng, Holm, and Natvik (2021), who estimate the effect of lottery winnings on consumption and savings (but not earnings), observe Norwegian tax records on income and wealth but similarly lack individual-asset transactions data.

²The broader literature studying consumption responses to income shocks using surveys includes Imbens, Rubin, and Sacerdote (2001), Kuhn et al. (2011), Jappelli and Pistaferri (2014), Fuster, Kaplan, and Zafar (2021), and Lewis, Melcangi, and Pilossoph (2026).

between low- and high-income individuals.³

The event-study figures show sharp, immediate effects at the time of the win on consumption (including both durable and nondurable consumption), savings (including both debt repayment and financial asset accumulation), and wage earnings. Across all outcomes, we find no evidence of differential trends prior to the win, which supports the validity of the research design. Aggregated over the five years after the win, lottery recipients increased consumption by about 38 cents per dollar, reduced after-tax wage earnings by roughly 15 cents per dollar, and saved 19 cents. The savings response consists of 9 cents in debt repayment and 10 cents in investment. Summing across all responses shows that we can account for 71 cents per dollar of lottery winnings. The remainder is largely accounted for by increased spending on durable goods like housing, which is primarily allocated to consumption after the first five years following the win. Consistent with this interpretation, if we instead sum expenditures rather than consumption, we can account for 89 cents per dollar of lottery winnings.

Next, we examine heterogeneity in wealth effects for low- and high-income individuals. Our findings reveal that low-income winners increased consumption by roughly 25 cents of every dollar won over two years, compared with 15 cents among high-income winners. By contrast, high-income winners reduced wage earnings by roughly 14 cents per dollar, compared with only 4 cents among lower-income winners. Finally, low-income individuals increased their savings by 24 cents while high-income individuals increased savings by 32 cents. Thus, our evidence supports the conjectures of Keynes and Malthus: low-income individuals boosted their spending more in response to winning the lottery, while high-income individuals were more likely to invest their winnings and consume more leisure.

A central motivation for studying lottery winners is to learn about the effects of policies that generate persistent changes in disposable resources, such as a Universal Basic Income (UBI). This

³As we describe in more detail below, we translate the effects of lottery winnings on credit card spending into overall consumption using time- and income-group-specific scaling factors derived from the Bank of Canada's Methods-of-Payment surveys and two Statistics Canada sources: the Survey of Household Spending (SHS) and the Distribution of Household Economic Accounts (DHEA). Additionally, we validate our imputed measure of consumption by showing that it closely tracks official consumption aggregates over time.

requires extrapolating from the response to a one-time wealth shock to the response to a recurring flow of unearned income. Marginal propensities to consume and earn out of unearned income provide portable behavioral parameters for making this translation and for comparing responses across settings. However, recovering these parameters from lottery wealth effects requires knowing the flow of resources generated by the lottery prize, which depends on how winners allocate their wealth over time. Two common approaches to measuring unearned income either involve assuming that winnings are perfectly smoothed over time (“annuitization approach”) or imputing changes in wealth from taxable capital income using an assumed rate of return (“capitalization approach”). We follow a different approach. Because we have a high-quality consumption proxy derived from credit card spending measured in the credit report data, we are able to recover unearned income directly from the budget constraint using a two-sample instrumental variables (IV) estimator.

We estimate an MPC of 0.71 and an MPE of -0.42 .⁴ By contrast, if we use the annuitization approach to construct our estimates, we obtain an MPC estimate of 0.40 and an MPE of -0.86 . The difference arises because lottery winners do not perfectly smooth their winnings; instead, our evidence indicates that they draw down their prizes substantially faster than the annuitization benchmark implies. This highlights the value of directly measuring the stream of unearned income rather than inferring it from an assumed annuitization of lottery wealth. Consistent with this interpretation, we calibrate a standard life-cycle model and show that discount rates substantially higher than the interest rate can rationalize the rapid drawdown of lottery wealth. The calibrated discount rate is higher for low-income individuals, implying greater impatience and a stronger incentive to “front-load” resources.

Accurately recovering the MPC and MPE addresses the first translation problem, from lottery wealth to unearned income. Applying these parameters to *persistent policies* creates a second problem, which is determining the flow of resources generated by the policy in each period. We illustrate the importance of both issues, and the role of income heterogeneity, through three policy

⁴As we describe in more detail below, because changes in (pre-tax) earnings change taxes paid, there is an accounting identity $MPC - MPE + MPT = 1$, where MPT is the marginal propensity to pay taxes out of unearned income.

applications. Our first application considers how a policymaker can target fiscal transfers based on income to obtain the greatest “bang for the buck.” We find that, given a fixed fiscal budget, targeting transfers to individuals below the 25th income percentile of the income distribution generates a larger increase in aggregate consumption than distributing the budget equally across all Canadians. This follows directly from our finding of a higher MPC among low-income individuals.

Our second application evaluates the fiscal cost of a Universal Basic Income (UBI), which has recently drawn considerable attention in Canada and other countries.⁵ Applying our estimated marginal propensities to a persistent transfer requires translating each dollar of the UBI into the flow of resources available in each period. Using the discount rate calibrated from the life-cycle model to simulate the front-loading of unearned income in response to a UBI, we find a larger surcharge on labor income relative to a perfect consumption-smoothing benchmark. Thus, even when the underlying MPC and MPE are correctly estimated, assumptions about how households allocate the resources generated by a persistent policy over time can materially affect its calibrated fiscal cost.

Lastly, we consider optimal top tax policy using the model in Hellwig and Werquin (2026), which extends Saez (2001) to a dynamic setting with preference heterogeneity. Using the credit report and tax data, we estimate that consumption is less concentrated than income at the top in Canada. This can arise either from identical, non-homothetic preferences (i.e., savings is a luxury good), or because higher-productivity individuals have a stronger preference for savings relative to consumption, even conditional on resources. Under a separability assumption, we show that our income effect estimates can be used to identify non-homotheticities and, when combined with the Pareto coefficients for consumption and income, heterogeneity in preferences for saving across the population. Our findings establish a role for preference heterogeneity and, using behavioral responses specific to high earners, imply an optimal labor income tax smaller than the Saez (2001) benchmark in part because of a substitution from income taxation to a tax on savings of

⁵For example, a UBI is currently under investigation in the Senate as part of Bill S-206, which would result in every Canadian over age 17 gaining access to a guaranteed livable basic income without work requirements. See <https://www.parl.ca/legisinfo/en/bill/45-1/s-206>.

a similar magnitude. Calibrating the same formulas using our full-sample estimates yields larger optimal labor income and savings taxes, which shows the value of providing policymakers with estimates of behavioral responses specific to the group that is affected by the policy. Like in the UBI application, we find that the optimal tax rates are sensitive to assumptions about intertemporal resource allocation: relaxing the assumption of perfect consumption-smoothing materially changes both the optimal labor income and savings tax rates.

A key novelty of our study is the use of a consumption measure that is constructed from directly observed expenditure data in credit reports. The credit report data contain credit card spending, which captures expenditure on nondurable goods, and debt balances on auto loans and mortgages, which can be used to infer expenditure on durable goods.⁶ This relates to studies that measure consumption spending directly in administrative credit card and bank records (Agarwal, Liu, and Souleles, 2007; Agarwal and Qian, 2014; Baker, 2018; Olafsson and Pagel, 2018; Ganong and Noel, 2019; Agarwal, Qian, and Zou, 2021; Chetty et al., 2024). While these studies typically feature accounts at selected lenders, our data cover the universe of individual-level accounts across all lenders that are included in the credit bureau data.

Since our savings measure combines debt repayments observed in the credit bureau data with financial asset accumulation observed in the tax records, we provide a broader measure of the savings responses to financial windfalls. Existing studies using administrative data often observe only one of these margins: those using credit report data observe debt repayment but not financial assets (e.g., Di Maggio et al., 2017; Brown, 2021; Cookson, Gilje, and Heimer, 2022), while those using tax records provide capital income flows but not debt repayment (e.g., Golosov et al., 2024).⁷ Other studies use administrative balance-sheet data that contain both assets and liabilities, but observe them only as annual stocks and therefore do not capture the account-level repayment

⁶Credit bureau reporting on credit cards outside of Canada is typically limited to end-of-month credit card balances, not actual payments (Gibbs et al., 2024; Guttman-Kenney and Shahidinejad, 2025). Changes in balances over time thus necessarily reflect spending and debt repayment. Canadian credit reports have the advantage that lenders report actual credit card payments to credit bureaus, allowing us to separately observe spending and debt repayment.

⁷There are several papers studying debt repayment that rely on proprietary data from financial institutions (e.g., Agarwal, Liu, and Souleles, 2007; Agarwal and Qian, 2014) or surveys (e.g., Hisnanick and Kern, 2018; Parker et al., 2022; Koşar et al., 2023; Boutros and Mijakovic, 2024).

behavior recorded in credit bureau data (e.g., Fagereng, Holm, and Natvik, 2021; Druedahl and Martinello, 2022). Studies that report both margins using higher-frequency measures of debt repayment and asset accumulation typically rely on survey responses (Sahm, Shapiro, and Slemrod, 2010; Coibion, Gorodnichenko, and Weber, 2020; Bartik et al., 2025), while we observe both margins in administrative credit and tax data.

On the methodological side, our two-sample IV estimator allows us to avoid making the strong behavioral and measurement assumptions associated with both the annuitization and capitalization approaches by working directly from the budget constraint. Graber et al. (2026) also work from the budget constraint, but on the earnings side: they use it to measure unearned income directly from capital income and asset transactions and estimate the marginal propensity to earn from earnings responses to that measure.⁸ Our approach is the consumption-side mirror image approach: by directly observing consumption alongside wage earnings and labor taxes, we recover the effect of lottery winnings on unearned income from the budget constraint. Both approaches highlight the value using the budget constraint along with high-quality administrative data to infer how wealth is allocated over time.

Finally, our earnings results relate to studies examining how lottery winnings and other exogenous changes in wealth or unearned income affect earnings (Imbens, Rubin, and Sacerdote, 2001; Cesarini et al., 2017; Picchio, Suetens, and van Ours, 2018; Bulman et al., 2021; Golosov et al., 2024; Vivalt et al., 2025; Georgarakos et al., 2025; Graber et al., 2026).⁹ Throughout the paper, we compare our estimates primarily with prior lottery studies, and, following Golosov et al. (2024), we distinguish between the earnings response per dollar of lottery wealth and the marginal propensity to earn out of unearned income, and we report both measures. This distinction is

⁸Like us, Graber et al. (2026) also estimate heterogeneity in income effects across the income distribution, and additionally they estimate both the income and substitution effects of income taxation across the income distribution. Relative to this work, our goal is to use exogenous variation in lottery winnings to study income effects on consumption, savings, and labor supply across the income distribution using unusually rich consumption data and more comprehensive savings data which covers both asset accumulation and debt repayment.

⁹Our findings also relate to the literature estimating labor supply responses to income transfer programs (Jacob and Ludwig, 2012; Feinberg and Kuehn, 2018; Jones and Marinescu, 2022). These studies provide valuable quasi-experimental evidence, although comparisons across settings can be complicated by differences in the nature of the income shock, including whether it is temporary or permanent and whether it is anticipated or unanticipated.

important because marginal propensities out of unearned income provide a natural way to transport evidence from one-time wealth shocks to policies that generate different streams of resources over time. Recovering these parameters from lottery data, however, requires either observing the relevant components of the household budget constraint or imposing assumptions about how lottery wealth is allocated over time. Once recovered, applying these parameters to a persistent policy requires a second translation from the policy transfer into the flow of resources available in each period. Our analysis shows that assumptions about intertemporal allocation can therefore matter both for estimating marginal propensities from lottery data and for using those parameters to evaluate persistent policies.

The remainder of our paper is organized as follows. Section 2 describes our lottery, credit report, and tax data. Section 3 describes identification and estimation of the treatment effect. Section 4 reports our empirically estimated wealth effects on consumption, savings, and wage earnings and benchmarks them against estimates from other lottery studies in the literature. Section 5 translates our wealth effects to marginal propensities to consume and earn out of unearned income. Section 6 presents our policy applications. Section 7 concludes.

2 Data and Descriptive Statistics

This section describes the lottery data obtained from an anonymous province in Canada, credit bureau data from Equifax Canada, and administrative tax records from Statistics Canada. We then describe the procedures used to link lottery winners separately to the credit bureau and tax records.

2.1 Lottery Data

Our lottery dataset consists of the universe of lottery winners in excess of \$999.99 from 2004 through 2021 in a Canadian province.¹⁰ The majority of the lottery products in our data are “classic

¹⁰All lottery products in Canada are administered by Crown corporations. Winners of prizes of less than \$1,000 are not required to report their names and addresses to the lottery corporation and are thus excluded from our data. An anonymous corporation provided us with the lottery data; due to a non-disclosure agreement, we do not disclose the name of the province or lottery corporation.

lotteries” in which a winner is decided by a random draw of numbers. For example, the most common lottery product in our dataset is called Lotto 6/49, where players pick six numbers (from one to 49), and a lottery ticket is designated a winner if at least two of the chosen numbers match the randomly drawn numbers. The prize amount is primarily dictated by two factors. First, the prize amount is increasing in how many of a player’s chosen numbers match the drawn numbers, irrespective of order. For example, if the winning numbers are 1-3-17-20-26-32, a winning ticket could contain only 26 and 32 (a smaller prize) or 20, 26, and 32 (a slightly larger prize). Second, the prize amount decreases with the number of winning tickets and increases with the overall size of the lottery pool. The overall size of the lottery pool depends primarily on the number of tickets purchased; however, it can also depend on prior lotteries if those lotteries had no “jackpot” lottery winners (i.e., a winner who guesses all six numbers correctly). In this scenario, the ensuing lottery grows in size.

For all solo lottery winners, we observe first and last names, six-character postal codes,¹¹ lottery payment date, lottery product (e.g., Lotto 6/49), and lottery winnings. For prizes with two winners, we observe complete information on both individuals and their common ticket number; for such prizes, we assign half of the prize amount to each winner. For prizes with more than two winners, we do not observe all winners’ names, only the winner who received the prize “in trust” for the rest. In these cases, since we are unable to assign prizes to all winning individuals, we exclude them from our sample. We restrict our sample to winners who win only one prize over the entire sample period, and we exclude wins larger than \$250,000 (the 98th percentile), examining robustness to this threshold. The single-prize restriction allows us to cleanly identify the effect of a single wealth shock.

The lottery data we use offer a few advantages compared with similar data used in previous studies. First, the data include the exact payment date and the amount won for each prize, which provides precise information about the timing and size of the wealth windfall. This feature is valuable for estimating short-run consumption responses. Second, all prizes in the data are lump-

¹¹Six-character postal codes are similar to ZIP+4 extensions (i.e., 9-digit ZIP codes) in the U.S.

sum prizes, which avoids mixing lump-sum prizes and “annuity-like” prizes that are paid out gradually over time. Finally, the lottery winnings are non-taxable, so the observed prize amount equals the financial windfall received by the winner. As Golosov et al. (2024) show, using the pre-tax lottery value when winnings are taxable can lead to a substantial underestimate of wealth effects.

2.2 Credit Bureau Dataset

We use two anonymized quarterly borrower-level credit bureau datasets from Equifax Canada. The first contains credit records for lottery winners, constructed by matching our sample of winners to Equifax data (Section 2.4). The second is a 0.1 percent random sample of the Canadian population covering 2016 to 2021, which we use to compare the baseline credit characteristics of lottery winners with those of the broader population and thereby assess the external validity of our results. By 2015, nearly all major Canadian credit card lenders were reporting monthly actual payment information to credit bureaus, enabling Equifax to construct our measure of credit card spending, as described in Section 2.2.1. This reporting coverage determines the sample window for our analysis, described in Section 2.4.

The primary variables from the credit bureau data used in our empirical analysis are credit card spending; outstanding debt, measured in total and separately for credit cards, unsecured lines of credit, home equity lines of credit, installment loans, mortgages, auto loans, and student loans; indicators corresponding to credit card and mortgage ownership; Equifax’s risk score; and indicators of financial distress, including delinquency, collections, and bankruptcy.¹² Table 1 reports summary statistics for these variables. To assess whether pre-lottery credit characteristics predict win amounts, we augment them with a comprehensive set of credit bureau characteristics, including product-specific measures of number of accounts, balance, delinquency, credit limit, and utilization, as well as an indicator for a recent address change.

¹²For most continuous credit bureau outcomes, we winsorize observations at the 99.5th percentile within calendar years, following Allen, Boutros, and Guttman-Kenney (2026). The exception is credit card spending, which we winsorize at the 99th percentile within calendar years.

The credit bureau data contain limited demographic information. For each individual, we observe an age band of roughly 5 years and residential location, which we map to Census Dissemination Areas (DAs). To examine heterogeneity by income, we assign each winner the average household income in their DA of residence, measured in the most recent census preceding the lottery win. DAs contain approximately 400 to 700 people and therefore provide a geographically detailed measure of neighborhood income (Government of Canada, 2021).¹³

2.2.1 Constructing and Characterizing Credit Card Spending in Canada

A unique feature of Canadian credit reporting is that lenders report actual credit card payments to credit bureaus; such payment data are generally unavailable in other countries, including the U.S. (Guttman-Kenney and Shahidinejad, 2025). Observing actual payments is essential for constructing a reliable measure of credit card spending because changes in balances alone do not distinguish new spending from debt repayment (e.g., Gibbs et al., 2024). Equifax combines data on monthly credit card payments with changes in credit card balances to construct monthly credit card spending according to the following identity:

$$\text{Credit card spending}_{it} = (\text{Balance}_{it} - \text{Balance}_{i,t-1}) + \text{Payment}_{it},$$

where i indexes individuals and t indexes months, and Balance_{it} and Payment_{it} are the statement balance and the payment towards that balance, both reported by the credit card lender.¹⁴ Equifax then sums this monthly measure within calendar quarters to produce the quarterly credit card spending series that we observe. For individuals who are transactors, paying off their statement balances in full at the end of each month, the identity perfectly captures their credit card spending.

¹³DAs are even smaller in size than Census block groups in the U.S., and as a result neighborhood income should be a reasonably accurate proxy for individual income, especially when we create above-median and below-median income groups in our empirical analysis. Consistent with a limited amount of measurement error, when we use individual income observed in tax data and compare the results with those obtained using DA neighborhood income, we find very similar results for all of the main outcomes observed in the tax data.

¹⁴The identity is written for an individual holding a single credit card. For holders of multiple cards, Equifax calculates spending at the card level and then sums across all cards to the individual level.

For revolvers, who carry positive balances from month to month, the identity overstates credit card spending because some of the payment goes toward paying down accrued interest on balances. Appendix Section B.1 details how we approximate accrued interest on persistent balances and highlights why isolating actual spending from interest charges similarly hinges on observing actual payments.

Credit cards were widely held and accepted in Canada during our sample period. At the sample midpoint, 89 percent of Canadians owned a credit card (Henry, Huynh, and Welte, 2018; Tompkins and Galociova, 2019), and 90 percent of small- and mid-sized merchants (6–50 employees) accepted them (Huynh, Nicholls, and Nicholson, 2019). Credit cards generally cannot be used for major durable purchases, such as housing, or to service other debt, so their use is concentrated in nondurable goods and services. In 2017, for example, they were the dominant payment method for everyday purchases, accounting for about half of grocery and gasoline spending, while durable goods (such as electronics, furniture, and appliances) accounted for less than 10 percent of credit card spending (Henry, Huynh, and Welte, 2018). Because nondurables are consumed at or shortly after purchase, expenditure closely approximates consumption for these goods, so credit card spending directly measures the portion of nondurable consumption purchased by credit card. Given the share of nondurable consumption purchased by credit card, we can scale observed credit card spending to total nondurable consumption. We can then use the nondurable share of total consumption to recover total consumption. We now describe our procedure, which implements these two scaling steps using available data.

2.2.2 From Credit Card Spending to Total Consumption

We construct total consumption in two steps. First, we scale credit card spending to total non-housing consumption using the share of non-housing purchases made by credit card. We use non-housing rather than nondurable consumption because payment-method shares are available only for non-housing purchases as a whole. Second, we scale non-housing consumption to total consumption using the non-housing share of total consumption. Appendix Section B.2 provides

additional details.

Step 1: From Credit Card Spending to Non-Housing Consumption. The first scaling factor accounts for non-housing purchases made using payment methods other than credit cards. For example, if credit cards account for half of non-housing purchases, then \$1 of observed credit card spending corresponds to \$2 of non-housing consumption. The credit card share varies across households and over time. In 2017, credit cards accounted for 56 percent of non-housing purchases by value overall, compared with 44 percent among households earning less than \$45,000, while the overall share increased from 46 percent in 2013 to 62 percent in 2021 (Henry, Huynh, and Welte, 2018; Henry, Shimoda, and Zhu, 2022).¹⁵ We therefore construct scaling factors that vary by income group and year.

Step 2: Adding Housing Consumption. The second scaling factor accounts for housing consumption. We use the non-housing share of total consumption from the Distribution of Household Economic Accounts (Statistics Canada, 2026), which includes imputed rent for owner-occupied housing. This approach treats homeowners as consuming the housing services provided by their homes, following Eika, Mogstad, and Vestad (2020). As in Step 1, we allow the scaling factor to vary by year and income group.

Combining the two steps yields our measure of total consumption. We validate this measure against official consumption data in both the time series and the cross-section. First, Figure 1 compares the growth of our consumption proxy between 2016 and 2021 with the growth of aggregate household consumption measured in the Distribution of Household Economic Accounts (Statistics Canada, 2026) and the Survey of Household Spending (Government of Canada, 2017). Our proxy closely tracks the official measures over this period. Second, Table A.1 compares consumption levels overall and across income groups in 2017. Because the official data are

¹⁵Non-housing spending consists primarily of nondurables, for which expenditure closely approximates consumption, but can also include vehicle purchases, for which we do not directly observe consumption. We treat vehicle expenditure as vehicle consumption. The two coincide if the vehicle stock is stable and vehicles are replaced at a constant rate, so that new purchases offset depreciation of the existing stock.

measured at the household level, we aggregate our individual-level measure to the household level using equivalence scales that adjust for household size (Statistics Canada, 2026). The resulting consumption levels are nearly identical overall and across income groups.

The key assumption underlying our procedure is that lottery winnings do not change the share of total consumption devoted to housing, so that the non-housing share used in Step 2 applies both before and after the win. We assess this assumption in Section 4.1, where we present our estimated wealth effects on consumption.¹⁶

2.2.3 Constructing Savings from Debt Repayments

We observe debt balances quarterly, reflecting the balance on the last monthly statement of each quarter, as reported to Equifax Canada. For each individual, we measure total outstanding debt as the sum of balances across all credit and debt accounts in a given quarter. We infer debt repayment from changes in these balances around the time of a lottery win, capturing the portion of the savings response that operates through paying down existing debt.

2.3 Statistics Canada Tax Data

We use individual-level tax records from Statistics Canada’s T1 Personal Master File (T1PMF), which covers the universe of Canadian individual income tax returns from 2001 to 2018. Our primary outcomes are wage earnings, employment status, total taxes paid, and financial assets. We winsorize continuous outcomes at the 99th percentile. To construct financial assets, we combine taxable investment income flows with information on contributions to and holdings of Canada’s tax-preferred savings accounts. The tax records also contain demographic characteristics, including age, sex, and marital status. Table 1 reports summary statistics for these variables.

Our measure of total wage earnings is employment income reported on the T1 tax return; it is

¹⁶Step 1 requires an analogous assumption: that lottery winnings do not change the share of non-housing spending paid by credit card. Back-of-the-envelope calculations indicate that moving our estimates outside their 95 percent confidence intervals would require the credit card share of non-housing spending to fall by roughly a quarter or to rise by more than half. Steps 1 and 2 also require that an individual’s credit card share of non-housing spending and housing share of total consumption, respectively, be uncorrelated with the size of their consumption response within each income group and year, so that the average shares are the appropriate scaling factors.

gross of tax, measured at the individual level, and excludes self-employment income. We define an individual as employed in a year if total wage earnings are at least \$1,000. The T1 return also reports total income and total taxes paid, measured as net federal tax plus provincial tax, which we use to measure the labor income tax response.

We observe two taxable investment income flows: interest and other investment income (tax return line 121) and taxable capital gains (tax return line 127). Interest and other investment income includes foreign interest or dividend income, interest credited on bank accounts, term deposits and guaranteed investment certificates, Canada Savings Bonds, Treasury bills, and earnings on life insurance policies. Taxable capital-gains income represents the net capital gain on sold or transferred property such as land, buildings, shares, bonds, and fund and trust units.

We observe Canada's two main tax-preferred savings accounts: Registered Retirement Savings Plans (RRSPs), which are similar to traditional IRAs in the United States, and Tax-Free Savings Accounts (TFSA), which are similar to Roth IRAs. For RRSPs, we observe both contributions and withdrawals, and use net contributions to measure savings in the account. For TFSA, we observe the year-end fair market value, available from the program's introduction in 2009. Because investment income earned within these accounts is tax-sheltered, savings held there generate no taxable investment income. Inferring savings solely from capitalized taxable investment income would therefore miss this important saving margin. This distinction is particularly important given our focus on heterogeneity in the ownership of financial assets across income levels.

Our measure of savings in financial assets combines three components: implied asset stocks, obtained by capitalizing taxable capital gains and interest and other investment income at asset-specific rates of return; the fair market value of TFSA, a directly observed year-end stock; and cumulative net RRSP contributions. Capital gains and investment income are observed as annual taxable income flows rather than as asset stocks, so we capitalize them into implied stocks as $\text{stock} = \text{flow}/r$, using capitalization factors of 43.9 for capital gains and 71.8 for interest and other investment income (implied nominal returns of 2.28 and 1.39 percent), based on estimates from Hempel (2025). Because a single year of realized gains or interest receipts is a noisy signal of

the underlying stock, we average annual flow responses over all post-treatment years through year t before capitalizing, so that each point reports the mean annual return flow generated through that year rather than one year's realization. TFSA fair market value enters in levels, with no capitalization or averaging. RRSP net contributions (contributions minus withdrawals) are a saving flow rather than a stock, so we cumulate the estimated net-contribution responses from the win year through year t . Each component thus measures the additional asset stock accumulated by year t per dollar of winnings, so the three can be summed into a single savings response.

2.4 Linking Datasets

Our analyses draw on two matched datasets: lottery winners linked to their credit bureau records, and lottery winners linked to their tax records. Although both draw on the same underlying lottery data from an unnamed Canadian province, privacy law in Canada prevents us from linking an individual winner's tax and credit bureau records to each other.

For the first dataset, Equifax matched the 40,872 winners in our lottery sample to their credit records using their postal codes, names, and payout dates. The payout date indicates when each winner resided at the recorded postal code, allowing Equifax to match against the address on file at that time. A full match requires agreement on postal code and both first and last name; a partial match requires agreement on postal code, last name, and first initial. Because Canadian six-character postal codes are extremely small (15 households on average), these criteria identify winners precisely. In a representative data subset Equifax provided to us, 70 percent of winners are full matches and 29 percent are partial matches, for an overall match rate of 99 percent. After matching, Equifax deleted all personally identifiable information (names and postal codes). We restrict this matched sample to winners who won between 2015Q1 and 2020Q1, the period over which credit card spending is most reliably measured (Section 2.2.1) and which precedes the onset of the COVID-19 pandemic. We further require that each winner be observed over our full event window, which spans two years before and after the lottery win, yielding a final analysis sample of

11,703 winners between 2013Q1 and 2022Q1.¹⁷

For the second dataset, Statistics Canada linked the 40,872 winners in our lottery sample to their administrative tax records through its Social Data Linkage Environment (SDLE). The SDLE links individuals to a Social Insurance Number (SIN) using the winner’s first and last names and postal codes as recorded at the payout date, combining deterministic matches with weighted probabilistic matches validated by Statistics Canada. This process links 84 percent of the winners to a SIN, which then permits linkage to the tax records. Similar to the credit bureau data, Statistics Canada removed all personally identifiable information from the matched data, including first and last names and postal codes (and the SIN as an intermediary step). We restrict this matched sample to winners who won by 2017, so that each winner can be observed for at least one year after winning the lottery in the tax data, which span 2001 to 2018. The tax-matched sample therefore covers wins from 2004, the start of our lottery data, through 2017, yielding a final analysis sample of 30,000 winners, with the count rounded to the nearest 1,000 in accordance with Statistics Canada confidentiality requirements.

In both matched datasets, we classify winners into “cohorts” based on the timing of their win, defined at the year-quarter level in the credit bureau data and at the calendar-year level in the tax data to align with the frequency of each dataset. We refer to these cohorts when describing our empirical strategy in Section 3.

2.5 Descriptive Statistics

Table 1 presents summary statistics measured in the period prior to winning. Panel A draws on the credit bureau data for our sample of lottery winners and a random 0.1 percent sample of the Canadian population. Panel B draws on tax data and compares the tax-matched winners to tax filers from the same province.

Among our credit bureau sample (Panel A), the average prize amount is \$8,363. As is typical

¹⁷Because the event window extends beyond the range of win dates, the estimation sample includes some quarters before 2015, when credit card payment data were not yet fully reported, and some quarters after 2020Q1, during the COVID-19 pandemic. Our estimates are not sensitive to either cutoff.

in lottery studies, the distribution of lottery prizes is right-skewed (Figure A.1). Winners in this sample resemble the general population. They are 48 years old on average, as in the comparison sample. Their balance sheets are also similar: winners carry total debt balances of \$107,485 against \$103,435, most of which is mortgage debt in both groups. Differences in credit outcomes are modest: risk scores are about 28 points lower (727 versus 755), quarterly credit card spending is somewhat lower (\$2,960 versus \$3,615), and fewer winners hold a credit card (70 versus 83 percent), while delinquency (7 versus 4 percent) and bankruptcy (2 versus 1 percent) remain rare in both groups.

Winners in the tax-matched sample (Panel B) likewise broadly resemble the provincial tax-filing population, although they are older (50 years old versus 45), more likely to be male (64 percent versus 50 percent), and somewhat more attached to the labor market, with an employment rate of 75 versus 70 percent and wage earnings of \$43,036 versus \$33,416. Taken together, these comparisons suggest that lottery winners are not highly selected relative to the broader population along the financial characteristics we observe. Lottery participation is also widespread in Canada: roughly two-thirds of Canadians gamble at least once per year, with an overwhelming majority of players participating in lotteries, instant lotteries, or sports betting (Marshall, 2011; Rotermann and Gilmour, 2022).

3 Identification and Estimation of Wealth Effects

3.1 Identification

This section describes how we identify wealth effects on consumption, savings, and earnings. Our setting is a continuous-treatment design in which all individuals receive a positive treatment dose. Winners within a cohort receive lottery prizes of different sizes, and we identify the effect of an additional dollar of lottery wealth by comparing changes in outcomes across winners receiving different prize amounts. This corresponds to the no-untreated-unit continuous-treatment design

considered by Callaway, Goodman-Bacon, and Sant’Anna (2024).¹⁸

Following Callaway, Goodman-Bacon, and Sant’Anna (2024), we assume there are multiple time periods $t = 1, \dots, T$. The period in which individual i is first treated (e.g., wins a lottery) is the unit’s *timing group* and is denoted by $G_i \in \mathcal{G}$. The treatment is defined by a dose $d \in \mathcal{D}$ where $\mathcal{D} = (0, \infty)$. The observed dose (lottery prize amount) for individual i is denoted by D_i . We also define $W_{i,t} = D_i \mathbf{1}\{t \geq G_i\}$ which is the dose amount that unit i experiences in period t . For all units that are not yet treated by t , we have $W_{i,t} = 0$.

Individual i ’s potential outcome at time t is given by $Y_{i,t}(g, d)$. The untreated potential outcome is defined as $Y_{i,t}(0) \equiv Y_{i,t}(\infty, 0)$. For each unit i , we observe their outcome in period t , Y_{it} , which satisfies the following:

$$Y_{i,t} = Y_{i,t}(0) \mathbf{1}\{t < G_i\} + Y_{i,t}(G_i, D_i) \mathbf{1}\{t \geq G_i\}$$

Our main causal effects of interest are the wealth effects. The treatment effect concepts that correspond to these terms are causal response parameters. In order to define these, we first define the average treatment effect on the treated:

$$ATT(g, t, d|g, d) \equiv \mathbb{E}[Y_t(g, d) - Y_t(0)|G = g, D = d]$$

The causal effect $ATT(g, t, d|g, d)$ is the average treatment effect of dose d , timing group g , in period t , among units in group g that experienced dose d . Given this, the average causal response (ACR) parameter is defined as follows:

$$ACRT(g, t, d|g, d) \equiv \frac{\partial ATT(g, t, l|g, d)}{\partial l} \Big|_{l=d} = \frac{\partial \mathbb{E}[Y_t(g, l)|G = g, D = d]}{\partial l} \Big|_{l=d}$$

¹⁸An alternative specification, as implemented in Cesarini et al. (2017), exploits information on ticket-purchase counts, which we do not possess. We note, however, that Cesarini et al. (2017) find very similar results when they exclude information on the number of lottery tickets purchased, suggesting that this limitation need not be a source of bias. Our event-study results are consistent with this and show strong support for the validity of lottery windfalls as a credible idiosyncratic wealth shock.

Our identification relies on comparisons across lottery winners receiving different prize amounts within the same win cohort. We require that, absent differences in lottery wealth, winners receiving larger and smaller prizes would have experienced the same evolution of outcomes. More generally, to interpret differences in outcome changes across prize amounts as causal responses to an additional dollar of lottery wealth, we require that the evolution of potential outcomes under a given prize amount does not depend on the prize amount actually received. This corresponds to the *strong parallel trends condition* for continuous treatments in Callaway, Goodman-Bacon, and Sant’Anna (2024).¹⁹

3.2 Estimation: Event-Study Specification

For all our analyses, we use the following event-study specification:

$$Y_{itg} = \alpha_i + \gamma_{tg} + \sum_{\tau, \tau \neq -1} \beta_{\tau} D_i \times \mathbf{1}\{\tau = t - G_i\} + \varepsilon_{itg}. \quad (1)$$

For credit bureau outcomes (spending, saving via debt repayment), the indicators $\mathbf{1}\{\tau = t - G_i\}$ are dummies for the eight quarters before and after the lottery win (omitting the coefficient for the quarter immediately preceding the lottery win). For tax outcomes (earnings, saving via investment), the indicators $\mathbf{1}\{\tau = t - G_i\}$ are dummies for the five years before and after the lottery win (omitting the coefficient for the year immediately preceding the lottery win). We include individual fixed effects (α_i) and cohort-by-calendar-time fixed effects (γ_{tg}). The latter allow each lottery-win cohort to have an unrestricted time path, so identification of β_{τ} comes from within-cohort variation in prize amounts conditional on individual fixed effects.²⁰ Lastly, we cluster standard errors at the individual level.

¹⁹Callaway, Goodman-Bacon, and Sant’Anna (2024) show that continuous dose variation can identify average causal responses without an untreated comparison group (Theorem 3.3 and Proposition S4). Because we observe the exact (and untaxed) prize amount of each winner, we can include all cohorts of winners, in contrast to a traditional “stacked design,” which must exclude cohorts for which no not-yet-treated controls remain.

²⁰Because our specification includes individual fixed effects, balance in pre-win characteristics is not a direct test of our identifying assumption, which concerns differential outcome trends across prize amounts. We nevertheless report these tests as complementary evidence that prize size is not systematically related to observable winner characteristics. Table A.2 shows that 50 pre-win credit bureau characteristics are generally balanced across prize amounts within win cohorts.

Our estimator aligns closely with the dynamic two-way fixed effects (TWFE) estimator with a continuous treatment. Callaway, Goodman-Bacon, and Sant’Anna (2024) show that under the assumptions considered above, the weights of average causal responses across doses are positive and sum to one, providing a (weakly) causal interpretation of TWFE. A pitfall is that the TWFE weighting scheme between doses may not equal the dose distribution among the treated. However, linearity of realized outcomes resolves such interpretation challenges and fully restores a causal interpretation in terms of ACRs.

To evaluate heterogeneity of responses based on income, we estimate Equation (1) separately for individuals with above- and below-median income. In the credit bureau data, income is not observed directly, so we assign individuals the neighborhood-average household income of their census DA of residence in the closest census year prior to winning. Below-median winners have DA incomes below \$80,502 (group mean \$50,259), while above-median winners have DA income above \$80,502 (group mean \$134,405), all in 2015 Canadian dollars. In the tax data, we observe income directly and split at the median of own pre-win individual income.

4 Wealth Effects

In this section, we describe how an individual allocates their lottery windfall to consumption (durables and nondurables), savings (via debt reduction and asset accumulation), and wage earnings. All our results are based on estimating the dynamic specification in Equation (1). When we examine heterogeneity by income, we estimate this equation separately for winners with above- and below-median income prior to winning. We present all main results visually in the form of event-study plots. For credit report outcomes, we plot quarterly event-study estimates from eight quarters prior to and eight quarters after the lottery win ($t = -8$ to $t = +8$). For tax outcomes, we report annual event-study estimates from five years before to five years after the win ($t = -5$ to $t = +5$). To provide a concise summary of the overall post-treatment effect, we aggregate these results by taking a simple unweighted average of the event-time coefficients in the period following the lottery win. For all outcomes, we also report wealth effect estimates that aggregate the event-

study coefficients over a two-year period and examine heterogeneity in these estimates by income. All standard errors are clustered at the individual level. For reference, the aggregated two-year point estimates are provided in Table 2. Finally, at the end of each subsection, we benchmark our estimates against other lottery studies (Imbens, Rubin, and Sacerdote, 2001; Cesarini et al., 2017; Picchio, Suetens, and van Ours, 2018; Fagereng, Holm, and Natvik, 2021; Golosov et al., 2024) since they similarly consider unanticipated wealth shocks, providing the closest “apples-to-apples” comparison.

4.1 Consumption

Figure 2(a) shows event-study estimates corresponding to the effect of lottery winnings on credit card spending (dashed line) and overall consumption (solid line) as well as 95 percent confidence intervals. The latter measure is constructed using the scaling procedure described in Section 2.2.1 and Appendix Section B. The pre-trends are flat up to the time of the lottery win, and are followed by an immediate increase of 0.013 (s.e. 0.002) in credit card spending and 0.029 (s.e. 0.005) in overall consumption per \$1 of lottery winnings in the first fully treated quarter, with positive effects continuing for 8 quarters post-win. Across all post-treatment quarters, the estimates imply an average quarterly increase of 0.009 (s.e. 0.002) in credit card spending and 0.020 (s.e. 0.004) in overall consumption per \$1 of lottery winnings. In Table A.3 we assess the sensitivity of our consumption estimates to the lottery win size cutoff, set in our baseline at the 98th percentile of the prize distribution (\$250,000). Raising the cutoff yields smaller estimates, whereas retaining all winners and winsorizing prizes at the 99th percentile (\$350,000) yields larger ones.

Durable Spending. Our measure of overall consumption combines durable and nondurable spending. We estimate the effect of lottery winnings on durable spending in two ways. The first is an indirect approach based on our scaling procedure, which allocates a fixed share of the overall *consumption* response to durables under the equiproportionality assumption described in Section 2.2.2: winners increase durable and nondurable consumption in the same proportion. The

second is a direct approach that follows the literature (e.g., Di Maggio et al., 2017; Ganong and Noel, 2020) in inferring durable *expenditure* from origination balances on mortgages, auto loans, and installment loans. Specifically, we define loan originations as increases in outstanding non-revolving loan balances. These two approaches estimate inherently different objects: consumption measures the flow of services from a durable good over time, whereas expenditure measures spending at the time of purchase. Comparing the two estimates therefore requires converting the expenditure effect recovered by the direct approach into a consumption effect (see, e.g., Eika, Mogstad, and Vestad, 2020; Laibson, Maxted, and Moll, 2022).

Appendix Table B.1 reports the results of the direct approach. The two-year effect on total origination balances is 0.075 (s.e. 0.022) per \$1 of lottery winnings and is driven almost entirely by mortgage originations, consistent with housing being the dominant durable category. Because origination balances capture only debt-financed purchases, they miss changes in housing consumption that do not involve mortgage borrowing, and the estimate should therefore be viewed as a lower bound. We convert this response into a consumption response in two steps. First, we translate each dollar of mortgage originations into the stream of mortgage payments it generates, assuming a standard 25-year amortization at the average mortgage rate prevailing over our sample period (4.9 percent), and use these payments as a proxy for the rental value of the home. Second, because this measure captures only mortgage holders, we scale it up by the inverse of the share of winners holding a mortgage. This scaling assumes that households without mortgages, including renters and outright homeowners, adjust their housing consumption by the same amount per dollar won. Further details are provided in Appendix Section B.3.

Figure B.2(a) shows that our two approaches produce similar wealth effects on housing consumption, in the 0.03–0.05 range per \$1 of lottery winnings, consistent with our equiproportionality assumption.²¹ The large difference between the consumption and expenditure responses highlights the importance of distinguishing between them. Figure B.2(b) illustrates this point by comparing our estimated overall consumption response (Figure 2) to our estimated overall expendi-

²¹The conversion can also be inverted, instead translating our implied consumption response into an implied expenditure response. This yields estimates in the 0.28–0.52 range.

ture response, which is constructed by summing our non-housing (nondurable) consumption effect and the housing *expenditure* effect estimated directly from origination balances. The expenditure response of 0.43 per \$1 of winnings is more than twice our estimated consumption response of 0.20.

Heterogeneous Responses. Next, we examine heterogeneity in consumption responses by income and credit utilization. We begin with income, for which we sum the event-study coefficients on overall consumption across all quarters following the lottery win. Figure 2(b) reports the estimates and shows that the two-year wealth effect on consumption is 0.15 (s.e. 0.04) for above-median income winners, compared with 0.25 (s.e. 0.06) for below-median income winners. Thus, lower-income winners spent a larger share of their lottery winnings on consumption.

We next examine heterogeneity by credit utilization, defined as the ratio of total revolving credit balances (credit card debt, home equity lines of credit, and unsecured lines of credit) to credit limits. Because we do not observe the scaling factors needed to translate credit card spending into overall consumption across the utilization distribution, we conduct this analysis using credit card spending and view the results as more suggestive. As shown in Figure A.2, the credit card spending response is larger among winners with above-median credit utilization than among those with below-median utilization. Taken together, these results indicate that consumption responses to lottery wealth are larger among lower-income winners and suggest that they are also larger among winners with higher credit utilization.

Comparison to Literature. We benchmark our consumption estimates against similar estimates from other lottery studies. Two lottery studies estimate a consumption wealth effect: Golosov et al. (2024) and Fagereng, Holm, and Natvik (2021).²² We estimate that winners spent 0.20 per dollar

²²There are two lottery studies that measure the consumption response for specific goods based on surveys. Imbens, Rubin, and Sacerdote (2001) study lottery prizes paid in annual installments over 20 years, with winners surveyed, on average, halfway through the payout period. Relative to the cumulative prize received by the survey date, they estimate that winners spend roughly 0.01 per dollar on cars and 0.04 per dollar on housing. Kuhn et al. (2011) find that winners increase expenditures on non-car durables by 0.03 per euro won and that a €10,000 prize reduces the age of a household's main car by approximately 0.4 years.

won over the first two years. Our estimate falls within the range of estimates in these studies, although differences in time horizons make direct comparisons challenging. Golosov et al. (2024) do not report a wealth effect directly, but we can infer it using their reduced-form effect of lotteries on imputed consumption, along with reported mean post-tax winnings. Their estimates imply a smaller annual response of about 0.024 per dollar of lottery wealth averaged over five years post-win. Fagereng, Holm, and Natvik (2021) estimate a larger consumption response than we do, with winners spending about 0.52 per dollar in the year of the win, 0.20 in the following year, and almost 0.90 per dollar over five years.²³

Finally, two related papers estimate the effects of income shocks on consumption. Lewis, Melcangi, and Pilossoph (2026) study the 2008 Economic Stimulus payment and report an average quarterly consumption response of 0.42, with heterogeneity ranging from 0.04 to 1.33. In contrast to our results, they find that income is positively associated with the consumption response, where higher-income households spend a larger share of the rebate. Bartik et al. (2025) find that recipients spend about 0.33 per month per dollar received, with larger estimates after adjusting for underreporting in their survey data. Their larger estimate likely reflects the sample composition (low-income households for whom we estimate a relatively larger wealth effect), and the nature of the treatment.

4.2 Savings

Next, we examine the effect of lotteries on savings, which we measure through both debt reduction (via credit data) and financial asset accumulation (via tax data). We discuss each of these estimates in turn.

²³While consumption in their study is measured at the household level, their estimates remain considerably larger than ours even after normalizing by the average number of individuals per household in their sample (1.91). Part of this difference may reflect their measure of consumption, which captures expenditures. In our setting, for example, attributing housing expenditure to consumption leads to estimates more than twice as large (see Figure B.2(b) and the discussion in Section 4.1).

4.2.1 Debt Reduction

Figure 3(a) displays the event-study estimates corresponding to the effect of lottery winnings on total debt balances, along with 95 percent confidence intervals. Mirroring the pattern of our consumption estimates, the pre-trends are flat up to the time of the lottery win, and are followed by an on-impact reduction in total debt balances of 0.119 (s.e. 0.024) per \$1 of lottery winnings in the first fully treated quarter, which persists throughout the estimation window (to $t = +8$). On average, across all post-win quarters, the estimates imply that winners reduced their outstanding debt by 0.10 (s.e. 0.03) per \$1 of lottery winnings.

Heterogeneous Responses. We next examine heterogeneity in debt repayment, first by income and then by credit utilization, similar to our analysis for consumption. Because total debt balance is a stock variable, we use the average post-treatment wealth effect for this exercise. Figure 3(b) reports the estimates by income. High-income winners repaid 0.12 (s.e. 0.05) per \$1 of winnings, compared with 0.07 (s.e. 0.02) among low-income winners. Thus, despite having greater resources prior to winning, high-income winners used a larger share of their lottery wealth to repay debt.

The pattern by credit utilization is different. As shown in Figure A.2, winners with above-median credit utilization repaid more debt out of their lottery winnings than those with below-median utilization. This result is consistent with Cookson, Gilje, and Heimer (2022), who find that subprime borrowers repay more of their outstanding debt following a wealth shock than prime borrowers.

Comparison to Literature. Our wealth effect on debt repayment of 0.10 per dollar won is similar to the response in Fagereng, Holm, and Natvik (2021), who found that debt fell by about 0.07 per dollar of lottery winnings in the year of the win. Bartik et al. (2025), in contrast, found that debt increased by about 0.05 per dollar of the monthly transfer. These differences may partly reflect the nature of the shock, since a lottery provides a large amount of liquidity at once, while the treatment in Bartik et al. (2025) arrives gradually over three years.

4.2.2 Financial Asset Accumulation

Figure 4(a) displays event-study estimates of the effect of lottery winnings on financial assets, along with 95 percent confidence intervals (and Figure A.3 reports separate event-study estimates for each of the four components of financial assets). We construct our measure of financial assets by combining four components observed in the tax data. For taxable assets, we capitalize income from taxable capital gains, interest, and other investments using the procedure described in Section 2.3. For tax-preferred assets, no capitalization is required: we observe the fair market value of Tax-Free Savings Accounts (TFSA) directly and construct Registered Retirement Savings Plan (RRSP) balances by cumulating net contributions. Similar to our other plots, pre-trends are flat up to the time of the lottery win, after which financial assets rise sharply in the first fully treated year and remain elevated through the end of the estimation window.

Because our financial asset measure is a stock variable, we interpret the response aggregated over the two fully treated years following the win as the wealth effect on saving via investment. Figure 4(b) displays this effect, overall and by income group. Winners accumulated 0.18 (s.e. 0.04) in financial assets per \$1 of lottery winnings, with high-income winners accumulating 0.20 (s.e. 0.08) and low-income winners 0.16 (s.e. 0.02). Thus, similar to our findings on debt repayment, high-income individuals saved a larger share of their lottery winnings than low-income individuals.

Comparison to Literature. Our two-year financial-asset response of 0.18 per dollar won is smaller than the immediate balance-sheet response in Fagereng, Holm, and Natvik (2021). In the year of winning, they find that about 0.42 per dollar is added to deposits and another 0.06 is allocated to stocks, bonds, and mutual funds, although deposits are subsequently drawn down as winners continue to consume their prizes. Golosov et al. (2024) also find an increase in capital income following lottery wins. Using their assumed rate of return to translate this capital-income response into an implied asset response gives an estimate of about 0.16 per dollar won, although financial-asset stocks are not directly observed in their data. Bartik et al. (2025), in contrast, estimate an increase in financial assets of only about 0.03 per dollar received, although it is

important to keep in mind that their sample consists primarily of low-income individuals.

4.3 Wage Earnings

Figure 5(a) shows event-study estimates corresponding to the effect of lottery winnings on total wage earnings, along with 95 percent confidence intervals. We define wage earnings as employment income—wages, salaries, and commissions before deductions—excluding self-employment income. The pre-trends are flat up to the time of the lottery win, and are followed by an immediate reduction in wage earnings of -0.039 (s.e. 0.007) per \$1 of lottery winnings in the first fully treated year. The estimates imply that lottery winners experienced an average annual wage-earnings effect of -0.038 (s.e. 0.004) per \$1 of lottery winnings, and the effects remain negative throughout the post-treatment period ($t = +5$).

Heterogeneous Responses. To examine income heterogeneity, we sum the event-study coefficient estimates on wage earnings over the first two years following the lottery win. The wealth effect estimates for each income group are presented in Figure 5(b). For above-median income winners, we find a two-year wealth effect on earnings of -0.14 (s.e. 0.02), and for below-median income winners, we find -0.04 (s.e. 0.01). Thus, high-income individuals reduced their labor earnings after winning the lottery by more than low-income individuals.

Extensive vs. Intensive Margin. Since the estimates in Figure 5 combine intensive- and extensive-margin labor supply responses, we decompose the total earnings response into extensive and intensive-margin earnings responses. As shown in Figure A.4, winning the lottery changed the probability of employment by -3.2×10^{-5} percentage points (s.e. 6.7×10^{-6}) per \$1 of winnings after two years. Among winners who remain employed, earnings changed by -0.077 (s.e. 0.013) per \$1 of winnings. After scaling this conditional earnings response by the post-win employment share, the intensive margin accounted for 88.7 percent of the overall two-year earnings decline of -0.09 (s.e. 0.01) per \$1 of winnings. Thus, the earnings response operated primarily through the intensive margin.

Comparison to Literature. Our earnings response is larger than the estimates from previous lottery studies. Averaged over the first five years after winning, we find an annual earnings response of about -0.038 per dollar of lottery wealth. Golosov et al. (2024) estimate an annual response of about -0.023 per dollar of lottery wealth, averaged over the first five years after winning.²⁴ Estimates from European lottery studies are somewhat smaller. Cesarini et al. (2017) find an average annual response of about -0.011 per dollar won, while Picchio, Suetens, and van Ours (2018) estimate a response of about -0.012 in the first year that grows to roughly -0.018 by year 3.²⁵ Graber et al. (2026) study Norwegian lottery winners and estimate an annual response of about -0.052 per NOK of the median prize.²⁶ The reduction grows with pre-win earnings, from about -0.029 to -0.058 per year between the bottom and top income deciles.

In sum, our combined evidence in this section supports the conjectures of Keynes and Malthus: low-income individuals are more likely to be hand-to-mouth consumers who boost their spending in response to winning the lottery, while high-income individuals are more likely to invest their winnings and consume more leisure.

4.4 Fully Accounting for the Lottery Winnings

So far, we have reported wealth effect estimates for consumption, savings, and earnings. Figure 6 summarizes these responses over a five-year horizon for the full sample and the low- and high-income subsamples. In the full sample, we account for roughly 71 cents of every dollar of lottery winnings through consumption, savings, and earnings, with similar shares accounted for in both

²⁴It is important to note that Golosov et al. (2024) restrict their sample to lottery wins of at least \$30,000 per winner. This contrasts with our sample restriction, noted in Section 2.1, which conditions on winnings above \$1,000 and below \$250,000. Further, Golosov et al. (2024) note that their estimates are larger if they include wins smaller than \$30,000, and they estimate an average effect of -0.055 when they condition on winning between \$30,000 and \$300,000, which is broadly similar to our estimate. Thus, the difference in the lottery prize distributions across the studies can largely account for the gap in estimates.

²⁵Although their primary focus is the effect of lottery winnings on children’s college outcomes, Bulman et al. (2021) also estimate earnings responses using U.S. administrative data. For this analysis, they restrict the sample to lottery winners whose children were at least 19 years old at the time of the win. They estimate a household earnings response of approximately -0.012 per dollar of pretax lottery winnings.

²⁶They find that average earnings fall by about NOK 18,000 per year three years after the win, which we compare to the median prize size of NOK 347,000 that they report.

income subsamples.²⁷

Much of the remaining gap can be explained by the distinction between durable expenditure and consumption. Our baseline accounting exercise converts housing expenditure observed in the credit report data into a flow of housing consumption. Figure 7 shows that if we instead treat housing as expenditure in the period of purchase, we can account for roughly 89 cents of every dollar of lottery winnings. The larger expenditure response reflects the fact that winners use part of their lottery wealth to finance durable purchases whose consumption services accrue over time. This distinction is consistent with recent work emphasizing that marginal propensities for expenditure and consumption can differ substantially (Laibson, Maxted, and Moll, 2022).

What accounts for the remainder? We consider two potential explanations. First, our consumption, savings, and earnings estimates are measured at the individual level, while some lottery winnings may be allocated to other household members. For example, Cesarini et al. (2017) find that the spouses of the lottery winners also reduced their labor earnings following a lottery win, although by less than the winners themselves. Labor supply responses by other household members may therefore account for part of the remaining winnings.

Second, some of the remainder may reflect measurement error or unobserved components of consumption and savings. Although measurement error cannot be fully ruled out, our validation exercises indicate that our overall consumption measure closely tracks external benchmarks. Our savings measure is more difficult to validate and does not capture all forms of saving. For example, some winners may hold part of their winnings in low-interest savings or checking accounts, which we do not directly observe and cannot reliably infer from tax records. Consistent with this possibility, Briggs et al. (2023) find that lottery winners retain a modest share of their winnings in bank accounts for several years, potentially reflecting inaction.

²⁷Bartik et al. (2025) conduct a similar accounting exercise and account for 60 cents of every dollar, or 71 cents after rescaling their consumption expenditure estimates using the Personal Consumption Expenditure or Consumer Expenditure Survey to account for the under-reporting of consumption expenditures. Their calculation does not account for changes in taxes generated by the earnings response to the transfer (footnote 33). In our setting, ignoring these tax effects would overstate the share accounted for by roughly 6 cents (see five-year estimate for Total Labor Taxes in panel (b) of Table 2). This reflects the decline in earnings following a lottery win: lower earnings reduce taxes paid, leaving winners with additional resources that can be allocated to other uses. Table A.4 provides further details.

Taken together, our results show that we can account for a large share of lottery winnings. The remaining gap likely comes from a mix of responses by other household members and forms of saving that are not captured in our data.

5 Propensities to Consume and Earn out of Unearned Income

The wealth effect estimates in the previous section capture dynamic marginal propensities out of lottery wealth: the period- t consumption and earnings responses per dollar of lottery winnings. We now turn to estimating static marginal propensities out of unearned income, where unearned income is the flow of resources generated by the lottery win and allocated to period t . To estimate these responses, we introduce a two-sample empirical approach that directly recovers both the change in unearned income and the marginal propensities to consume and earn out of it. We can implement this approach because we observe consumption alongside wage earnings and taxes. We then present our estimates and compare them with those from other lottery studies, particularly studies using the widely employed annuitization approach, as well as those from studies of shocks other than lotteries.

5.1 Two-Sample Instrumental Variables (IV) Estimator

Our empirical approach is a two-sample instrumental variables (IV) estimator, with lottery winnings serving as the instrument for unearned income. The first stage estimates the causal effect of lottery winnings on unearned income, while the second stage uses this variation to estimate the effects of unearned income on consumption (MPC), earnings (MPE), and taxes (MPT).

We begin by defining unearned income using the budget constraint:

$$c_t = y_t - \mathcal{T}(y_t) + \underbrace{\overbrace{(1+r)a_{t-1} - a_t - \mathcal{T}_a(ra_{t-1})}^{n_t} + L \times 1\{t=0\}}_{\text{Unearned income} \equiv N_t}. \quad (2)$$

Here, c_t denotes consumption in period t , y_t denotes wage earnings, a_t denotes assets held at the end of period t , r is the interest rate, and $\mathcal{T}$ and $\mathcal{T}_a$ denote taxes on labor and capital income,

respectively. We define n_t as unearned income in period t excluding lottery winnings and $N_t = n_t + L \times 1\{t = 0\}$ as total unearned income. Thus, in the period of the win ($t = 0$), lottery winnings increase unearned income by L .

The budget constraint allows us to recover the first-stage effect of lottery winnings on unearned income from observable responses. Specifically, we take expectations of Equation (2) over the population of lottery winners (conditional on the win amount L) and differentiate with respect to L :

$$\frac{d\mathbb{E}[N_t]}{dL} = \underbrace{\frac{d\mathbb{E}[c_t]}{dL}}_{\text{credit bureau data}} - \underbrace{\frac{d\mathbb{E}[y_t]}{dL}}_{\text{tax data}} + \underbrace{\frac{d\mathbb{E}[\mathcal{T}(y_t)]}{dL}}_{\text{tax data}}. \quad (3)$$

Equation (3) illustrates our two-sample approach to estimating the causal effect of lottery winnings on unearned income: the credit bureau data provide an estimate of the consumption response, while the tax data provide estimates of the earnings and labor income tax responses.

For an outcome $x \in \{c, y, \mathcal{T}(y)\}$, the IV estimand $\frac{d\mathbb{E}[x_t]/dL}{d\mathbb{E}[N_t]/dL}$ is the ratio of the reduced-form effect of lottery winnings on x to the first-stage effect on unearned income. Under an exclusion restriction that lottery wealth affects outcomes only through N_t , this identifies MPx ; i.e., the marginal propensity to consume, earn, or pay taxes out of unearned income. We estimate this ratio using our wealth effect estimates from Equation (1) in the numerator and the first-stage estimate from Equation (3) in the denominator. Because consumption is measured quarterly in the credit bureau data, while earnings and labor taxes are measured annually in the tax data, we aggregate the responses for each outcome over common two- and five-year horizons. Note that by construction, $MPC - MPE + MPT = 1$.

As described in Section 4.4, we do not fully account for the lottery winnings through the estimated effects on consumption, savings, earnings, and taxes. However, the two-sample IV approach described here remains valid if the unaccounted-for lottery winnings are allocated to forms of savings that we do not measure, rather than unmeasured consumption or earnings.²⁸ Intuitively,

²⁸An alternative assumption that would also deliver valid two-sample IV estimates is that each component is measured with the same proportional error. Specifically, if each term in Equation (3) is biased by the same constant factor α , then α cancels when constructing the ratio $\frac{d\mathbb{E}[x_t]/dL}{d\mathbb{E}[N_t]/dL}$.

Equation (3) recovers the change in unearned income from the observed changes in consumption, earnings, and taxes, without requiring us to observe the corresponding change in savings. We believe this assumption is likely to be a reasonable approximation because our consumption proxy is designed to capture total consumption, while the tax data provide a comprehensive measure of labor earnings. By contrast, our measures of savings constructed from credit report and tax data omit some forms of savings, such as checking account balances.

5.2 Estimates of the MPC, MPE, and MPT

Figure 8 presents the marginal propensities out of unearned income, constructed using wealth effect estimates averaged over the five years following the lottery win. We estimate that \$1 of lottery wealth raised unearned income by 0.10 (s.e. 0.02) per year. Scaling the wealth effects by this first-stage response yields an MPC out of unearned income of 0.71 (s.e. 0.05), an MPE of -0.42 (s.e. 0.06), and an MPT of -0.13 (s.e. 0.04). Note that these estimates satisfy the budget constraint $\text{MPC} - \text{MPE} + \text{MPT} = 1$.

We next examine heterogeneity by income. Consistent with savings out of lottery wealth increasing slightly with income, the first-stage effect on unearned income is slightly larger for below-median income winners: \$1 of lottery wealth raised annual unearned income by 0.11 (s.e. 0.02) for below-median income winners, compared with 0.10 (s.e. 0.02) for above-median income winners. The marginal propensities out of unearned income largely mirror the heterogeneity in the underlying wealth effects. The MPC is 0.88 (s.e. 0.04) for below-median income winners against 0.54 (s.e. 0.09) for above-median income winners, while the MPE is -0.16 (s.e. 0.04) against -0.68 (s.e. 0.12). The MPT is -0.03 (s.e. 0.01) for below-median income winners and -0.22 (s.e. 0.07) for above-median income winners. The larger-magnitude MPT for above-median income winners comes from both higher average tax rates as well as a larger-magnitude MPE. Table 3 summarizes these estimates together with their two-year counterparts, constructed using wealth effect estimates averaged over the first two years following the win.

Annuitization Approach. As discussed, several studies estimate the marginal propensity to earn out of unearned income by assuming that lottery winners annuitize their winnings over their remaining lifetime at a predetermined interest rate. To benchmark our estimates against this approach, we annuitize lottery winnings assuming a 3 percent interest rate, close to the average yield on long-term Government of Canada bonds over our sample period, an average winner age of 50, close to the average age in our sample, and a life expectancy of 80 years, approximately the average in Canada during our sample period (see Appendix Section C.1 for more details). Figure 9 reports the resulting annuitization estimates alongside our IV estimates for comparison. We estimate that over the five years after the win, each lottery dollar generates 4.9 cents of unearned income per year under the annuitization approach. This contrasts with an average of about 10 cents of unearned income per year using our IV approach. Thus, winners draw down their prizes roughly twice as fast as the level annuity benchmark implies. Combined with our wealth effect estimates, this yields an annuity-based MPE estimate of -0.86 (s.e. 0.09) and an MPC of 0.40 (s.e. 0.10). These patterns hold for both income groups.

One possible explanation for the rapid “drawdown” of the prize implied by our estimates is that lottery winners discount the future at a higher rate than the interest rate. To examine this, we assume that the lottery winners have Stone-Geary preferences over consumption and leisure and discount the future at a rate ϑ , which may differ from the interest rate r . Appendix Section C.2 shows that at $r = 3$ percent, matching the five-year average estimate on unearned income requires $\vartheta = 18.4$ percent. Thus, winners draw the prize down about twice as fast as a level annuity, and a discount rate that is substantially higher than the interest rate can rationalize this pattern. Furthermore, when we examine heterogeneity by income levels, we find that low-income individuals have a subjective discount rate of $\vartheta = 22.8$ percent, which is higher than the discount rate of high-income individuals of $\vartheta = 15.1$ percent.

Comparison to Literature. Comparisons across lottery studies are complicated by differences in conventions, including the annuitization rate, the tax treatment of winnings, and whether earn-

ings are measured before or after tax and at the individual or household level. Golosov et al. (2024) harmonize earlier estimates along these dimensions, providing a useful basis for comparison. Under common conventions, they report individual-level MPEs of -0.43 for the U.S. and -0.27 for Sweden.²⁹ Because lottery winnings are tax-exempt in Canada, our setting is most directly comparable to the after-tax prize convention in Golosov et al. (2024). On this basis, our MPE of -0.42 is close to the U.S. estimate of -0.43 and larger in magnitude than the Swedish estimate of -0.27 , consistent with the U.S.–Europe difference documented by Golosov et al. (2024). The denominators differ in construction, however: Golosov et al. (2024) normalize by annuitized unearned income, whereas ours is the unearned-income flow recovered from observed consumption. The two coincide when winners draw down their winnings at the annuity rate, which is approximately the case in their sample but not in ours; with the annuitization convention, our MPE is -0.86 . Graber et al. (2026) define unearned income from the budget constraint and measure it directly from capital income and asset transactions rather than by annuitizing the prize. They estimate an MPE of -0.26 that is fairly flat across the income distribution.

A complementary comparison comes from Vivalt et al. (2025), who study the earnings response to a guaranteed income transfer rather than lottery winnings. Like our approach, they define the MPE relative to the resources allocated to the period, accounting for changes in participants’ net worth. For wage earnings, the outcome most comparable to ours, they estimate MPEs ranging between -0.18 and -0.14 in administrative records (their Table III). These estimates bracket our MPE of -0.16 for below-median income winners, the most relevant comparison given that their sample consists of lower-income households.

Fewer lottery studies report MPCs out of unearned income because consumption is rarely observed in administrative data. Golosov et al. (2024) infer a household-per-adult MPC of 0.58 from the budget constraint; applying their approach at the winner level implies an MPC of approximately

²⁹These adjustments are quantitatively important. For example, converting the Swedish estimate of Cesarini et al. (2017) from after-tax to gross earnings increases the MPE in magnitude from -0.14 to -0.28 , while expressing the winnings in Imbens, Rubin, and Sacerdote (2001) on a pre-tax basis reduces the corresponding U.S. MPE from -0.43 to -0.27 . Golosov et al. (2024) can further reconcile their individual-level MPE with that of Imbens, Rubin, and Sacerdote (2001) by adopting the same 10 percent discount rate.

0.66, broadly comparable to our estimate of 0.71. They also find that the MPC declines with pre-win income, from 0.73 in the bottom quartile to 0.50 in the top quartile, consistent with our estimates of 0.88 and 0.54 for below- and above-median income winners, respectively. Graber et al. (2026) do not observe consumption directly, but assuming a 40 percent marginal tax rate, their MPE estimate implies an MPC out of unearned income of 0.84. Outside the lottery literature, Bartik et al. (2025) estimate a marginal propensity to spend of 0.56–0.74 after accounting for unmeasured consumption or savings.

6 Policy Applications

This section considers three policy applications that highlight the value of our lottery-based estimates. Section 6.1 illustrates the simplest use of our estimates by showing how fiscal stimulus transfers can have the greatest “bang for the buck” if they are targeted toward individuals with the largest MPC. We next use our income effect estimates in Section 6.2 to calibrate the fiscal externality associated with a UBI program. Lastly, we calibrate optimal top income and savings taxes in Section 6.3.

6.1 Fiscal Stimulus Policy

We start with an exercise modeled after Canada’s “middle-class” tax cut in 2025 (Canada, 2025). All below-median income individuals are given a one-time cash transfer of \$1,000. To estimate the aggregate consumption impact, we first use the naïve estimate of the population-average wealth effect on consumption. This is the estimate most readily available to policymakers, since existing studies typically report only population averages. However, this would underestimate the true consumption impact because we find that below-median individuals have larger wealth effects than the population average. Applying our population-average wealth effect to the 17 million individuals below the 50th percentile of income yields a two-year aggregate consumption increase of \$3.42 billion (Table 4, column (4)). Next, we use the consumption wealth effect of 0.25 for low-income households and find that aggregate consumption increases to \$4.32 billion (column

(5)). This represents a 26 percent gain in expected consumption and reduces the fiscal cost by 21 percent (from \$4.97 to \$3.94). These differences grow as targeting increases: for the bottom quartile (row 5), using the income-specific wealth effect raises aggregate consumption by 76 percent relative to the population-average calculation and lowers fiscal cost by 43 percent. This stylized exercise demonstrates the importance of evaluating fiscal policy with context-specific wealth effects. Gross, Notowidigdo, and Wang (2020) make a similar point for the business cycle: in a stylized calibration, applying a typical-year MPC out of liquidity to a credit expansion during a recession understates the consumption response by roughly 40 percent.

6.2 Universal Basic Income (UBI)

In this section and the next, we consider the dynamic model of consumption (C), savings (S), and earnings (Y) based on Hellwig and Werquin (2026). Preferences are given by $U(Y, C, S)$ and we consider a linear tax function $T(Y, S) = -B + \tau_Y Y + \tau_S S$.³⁰ Throughout the remainder of this section, we assume that preferences are separable in consumption, savings, and earnings. Every individual receives the transfer B , which the government finances with a proportional surcharge $\delta(B)$ on labor income. We suppress the index for individual types, though the endogenous variables vary across the population. Government revenue is tax revenue less government outlays on transfers:

$$\mathcal{R}(B) = \mathbb{E}[-B + (\tau_Y + \delta(B))Y(B) + \tau_S S(B)] \quad (4)$$

where $Y(B)$ is labor income, $S(B)$ is savings, and τ_Y and τ_S are the linear tax rates on labor income and savings, respectively. This framework extends the analysis of a UBI considered in Golosov et al. (2021) to allow for a tax on savings. We include the savings tax because the optimal savings tax is nonzero when preferences for savings are heterogeneous, which is what we find in Section 6.3.

We consider budget-neutral reforms, so the surcharge $\delta(B)$ satisfies $\mathcal{R}(B) = \mathcal{R}(0)$. Differ-

³⁰Hellwig and Werquin (2026) denote the lump-sum transfer by $-T_0$. We label it B since it corresponds more directly to the UBI.

entiating this condition with respect to B and solving for the surcharge (Appendix Section D.1) yields:

$$\frac{\partial \delta}{\partial B} = \frac{1 - (\tau_Y + \delta(B)) \mathbb{E} \left[\frac{\partial Y}{\partial B} \right] - \tau_S \mathbb{E} \left[\frac{\partial S}{\partial B} \right]}{\mathbb{E}[Y] - \frac{\tau_Y + \delta(B)}{1 - \tau_Y} \mathbb{E} \left[Y \zeta_{Y, \tau_Y}^U \right] - \frac{\tau_S}{1 - \tau_Y} \mathbb{E} \left[S \zeta_{S, \tau_Y}^U \right]} \quad (5)$$

where $\mathbb{E}[\partial Y / \partial B]$ and $\mathbb{E}[\partial S / \partial B]$ are the average marginal propensities to earn and save out of the transfer B , respectively, and $\mathbb{E} \left[Y \zeta_{Y, \tau_Y}^U \right]$ and $\mathbb{E} \left[S \zeta_{S, \tau_Y}^U \right]$ are weighted-average responses of earnings and savings to a proportional change in the net-of-tax rate $1 - \tau_Y$, in dollars, where ζ_{Y, τ_Y}^U and ζ_{S, τ_Y}^U are the individual uncompensated elasticities of earnings and savings with respect to $1 - \tau_Y$. There are two effects of a savings tax on the required surcharge which can work in opposite directions: first, in the numerator, if savings is a normal good, the income effect reduces savings, which generates less tax revenue and increases the required surcharge; second, in the denominator, a higher effective tax on earnings (due to budget neutrality) raises savings if the substitution effect dominates the income effect, thus increasing the tax base and decreasing the required surcharge. In summary, whether a savings tax increases or decreases the surcharge is theoretically ambiguous.

We evaluate Equation (5) at $B = 0$, where $\delta(0) = 0$. This involves calibrating the following sufficient statistics: mean earnings ($\mathbb{E}[Y]$), the income effects on earnings ($\mathbb{E}[\partial Y / \partial B]$) and savings ($\mathbb{E}[\partial S / \partial B]$), and the weighted-average earnings ($\mathbb{E}[Y \zeta_{Y, \tau_Y}^U]$) and savings responses ($\mathbb{E}[S \zeta_{S, \tau_Y}^U]$). We calibrate the inputs (and the required surcharge) under several scenarios. First, to examine the role of weighting and income heterogeneity, we report estimates based on unweighted average responses across the entire population. Next, to investigate the role of consumption smoothing, we evaluate the surcharge formula both when individuals perfectly smooth consumption, so there is no effect of the UBI on saving, and when there is an income effect on savings, which is consistent with our empirical evidence in Section 5.³¹ Finally, we evaluate the surcharge formula with and without a pre-existing savings tax.

We empirically implement the sufficient statistics as follows. First, the income effect on earnings is obtained using our estimated MPE in Table 3 scaled by the resource use per dollar

³¹Golosov et al. (2021) evaluate a UBI under a perfect smoothing assumption.

of transfer, $\partial N/\partial B$, which is calibrated using the model-based approach described in Appendix Section C.3. When we impose the perfect smoothing assumption, $\mathbb{E}[\partial S/\partial B] = 0$ and therefore $\mathbb{E}[\partial Y/\partial B]$ is equal to the MPE. Next, we compute mean labor income ($\mathbb{E}[Y]$) for each income group from Statistics Canada’s Distributions of Household Economic Accounts (DHEA). Third, the earnings response to the net-of-tax rate $1 - \tau_Y$ in Equation (5) depends on the uncompensated earnings elasticity, ζ_{Y,τ_Y}^U , which we recover using the Slutsky equation and compensated elasticity estimates from Gruber and Saez (2002) of 0.18 for low-income individuals and 0.57 for high-income individuals, and our lottery estimates. Calibrating the uncompensated savings response, ζ_{S,τ_Y}^U , is more challenging since there are no external estimates we can rely on. Appendix Section D.2 details how we construct it using the compensated earnings elasticity ζ_{Y,τ_Y}^H , mean labor income $\mathbb{E}[Y]$, and our estimated income effects on earnings ($\mathbb{E}[\partial Y/\partial B]$) and consumption ($\mathbb{E}[\partial C/\partial B]$).

Panel A of Table 5 reports the values of the inputs used to calibrate Equation (5). The first four columns are based on estimates that use unweighted average responses across the entire sample, while the next two columns use weighted-average responses across income groups. The estimates under the perfect smoothing assumption are shown in columns (1) and (2) and the remaining columns (3)-(6) relax this assumption. The distinction between perfect smoothing and front-loading has quantitatively important implications for the income effect on earnings. Under perfect smoothing, a one-dollar increase in the UBI generates a one-dollar increase in unearned income, so our estimated MPE implies an earnings response of $\partial Y/\partial B = -0.42$. When we allow savings to respond to the UBI, the income effect increases in magnitude to -0.72 . On the other hand, allowing for income heterogeneity has a more muted effect on the average responses.

Panel B of Table 5 reports the fiscal and behavioral effects of the UBI. It shows that relaxing perfect smoothing substantially increases the surcharge required to finance the UBI, which is a direct consequence of the larger income effect on earnings. Under perfect smoothing, the required surcharge is 2.76 percentage points per \$1,000 of transfer and earnings decline by \$0.60 per dollar of transfer. Allowing for a faster drawdown of resources raises the surcharge to 2.99

percentage points (without a savings tax) and increases the earnings decline to \$0.92 per dollar of transfer. Introducing a savings tax modestly reduces both effects, lowering the required surcharge to 2.88 percentage points and the earnings decline to \$0.91, while generating an increase in savings of \$0.17 per dollar of transfer. Allowing for heterogeneity in responses across the income distribution produces similar aggregate effects, with surcharges of 2.94 and 2.89 percentage points without and with a savings tax, respectively. Taken together, our results indicate that the perfect smoothing assumption has a much bigger impact on the UBI surcharge than whether or not we model heterogeneity.

6.3 Optimal Top Tax Policy

This section uses our estimates to calibrate optimal top tax rates on income and savings using the model of Hellwig and Werquin (2026), which extends the tax formula for labor income at the top (Saez, 2001) to a dynamic setting allowing for heterogeneous preferences for savings.³² Their tax formulas for labor income and savings depend on three additional sets of sufficient statistics for top earners: the Pareto coefficients for consumption ($\bar{\rho}_C$) and income ($\bar{\rho}_Y$) at the top, the compensated elasticities of consumption ($\bar{\zeta}_{C,\tau_Y}^H$) and earnings ($\bar{\zeta}_{Y,\tau_Y}^H$) with respect to the net-of-tax rate on labor income, and the elasticity of intertemporal substitution (EIS). They show that a planner should distort the consumption-savings decision when individuals differ in their preferences for savings. In the empirically relevant case in which consumption is less concentrated at the top than income, the optimal top tax rates are:

$$\bar{\tau}_Y = \bar{\tau}_Y^{\text{Saez}} \left[1 - \Lambda \left(1 - \frac{\bar{\rho}_Y / \bar{\rho}_C}{\bar{\zeta}_{C,\tau_Y}^H / \bar{\zeta}_{Y,\tau_Y}^H} \right) \right] \quad (6) \quad \bar{\tau}_S = \frac{\bar{\tau}_Y^{\text{Saez}}}{1 - \bar{\tau}_Y^{\text{Saez}}} \Lambda \left(1 - \frac{\bar{\rho}_Y / \bar{\rho}_C}{\bar{\zeta}_{C,\tau_Y}^H / \bar{\zeta}_{Y,\tau_Y}^H} \right) \quad (7)$$

$$\text{where} \quad \Lambda \equiv \frac{\bar{\zeta}_{C,\tau_Y}^H}{\bar{\zeta}_{C,\tau_Y}^H + \bar{\zeta}_{C,\tau_S}^H} = \frac{\bar{\zeta}_{C,\tau_Y}^H}{EIS} \quad (8) \quad \bar{\tau}_Y^{\text{Saez}} = \frac{1}{1 - \bar{\zeta}_Y^I + \bar{\rho}_Y \bar{\zeta}_{Y,\tau_Y}^H} \quad (9)$$

and $\bar{\tau}_Y^{\text{Saez}}$ in (9) is the Saez (2001) top tax rate, with $\bar{\zeta}_Y^I$ the income effect on earnings.

³²A related contribution is Ferey, Lockwood, and Taubinsky (2024) who also consider heterogeneous preferences and derive sufficient statistics formulas for optimal income and savings taxes.

Interestingly, they show that one can identify this preference heterogeneity in the presence of non-homothetic preferences. The main idea can be understood as follows. First, $\bar{\rho}_Y/\bar{\rho}_C$ reflects the concentration of consumption relative to income at the top. If consumption is less concentrated than income, this indicates that the rich save relatively more of their income than the poor. This can happen either because of non-homothetic, identical preferences (i.e., savings is a luxury good) or because higher-productivity individuals have a stronger preference for savings relative to consumption, even conditional on resources. In the latter case, savings reveals information about earning ability and a savings tax becomes a screening device, which leads the planner to deviate from the Saez benchmark.

In general, the observed concentration gap can reflect either channel. The key insight is that the elasticity $\bar{\zeta}_{C,\tau_Y}^H/\bar{\zeta}_{Y,\tau_Y}^H$, which represents the within-type causal effect of income on consumption, captures non-homotheticities.³³ When this elasticity is equal to one, consumption scales proportionally with exogenous changes in income, and preferences are homothetic. The observed concentration of consumption relative to income at the top is fully explained by preference heterogeneity. Conversely, if this elasticity is not equal to one, preferences are non-homothetic. Thus, the relevant sufficient statistic that enters the optimal tax formulas is $\frac{\bar{\rho}_Y/\bar{\rho}_C}{\bar{\zeta}_{C,\tau_Y}^H/\bar{\zeta}_{Y,\tau_Y}^H}$ which effectively rescales concentration at the top by the strength of non-homotheticities. If $\bar{\rho}_Y/\bar{\rho}_C = \bar{\zeta}_{C,\tau_Y}^H/\bar{\zeta}_{Y,\tau_Y}^H$, then the observed concentration of consumption relative to income at the top is fully explained by non-homothetic preferences and the optimal tax rate on labor income is the Saez (2001) benchmark; i.e., $\bar{\tau}_Y = \bar{\tau}_Y^{\text{Saez}}$, and $\bar{\tau}_S = 0$. Otherwise, the data reveal that individuals have heterogeneous preferences for savings and the optimal tax mix changes, with the direction of the change depending on the magnitude of $\bar{\rho}_Y/\bar{\rho}_C$ relative to $\bar{\zeta}_{C,\tau_Y}^H/\bar{\zeta}_{Y,\tau_Y}^H$.³⁴

Hellwig and Werquin (2026) calibrate their formulas using consumption-based moments from the literature. A central challenge in this exercise is that the relevant moments should apply to

³³The ratio is an IV estimator (instrument = taxes), where the numerator is the reduced-form effect of taxes on consumption and the denominator is the first-stage of taxes on labor income.

³⁴The term Λ scales the tax shift but is not a preference-heterogeneity statistic. If preferences are separable, then $\Lambda = \bar{\zeta}_Y^I$ and this term can be identified by the income effect on earnings.

high-income individuals, for whom direct evidence on consumption responses is scarce. In the absence of such estimates, they rely on consumption responses from Bartik et al. (2025), which are estimated for low-income households. Another challenge is that calibrating income effects on earnings and consumption depends on assumptions about how individuals smooth consumption.³⁵ Our contribution in this section is to instead calibrate optimal tax formulas using our lottery estimates for high-income earners and compare the resulting optimal policies with those obtained using population-level estimates. As in Section 6.2, we also consider the effect of imposing perfect smoothing.

Equations (6) and (7) identify the sufficient statistics required for the calibration. The Saez benchmark, $\bar{\tau}_Y^{\text{Saez}}$, is determined by the earnings Pareto coefficient, $\bar{\rho}_Y$, the compensated earnings elasticity, $\bar{\zeta}_{Y,\tau_Y}^H$, and the income effect on earnings, $\bar{\zeta}_Y^I$. The remaining sufficient statistics are the consumption Pareto coefficient, $\bar{\rho}_C$, the ratio of compensated tax elasticities, $\bar{\zeta}_{C,\tau_Y}^H / \bar{\zeta}_{Y,\tau_Y}^H$, and Λ . For the calibration, we assume that preferences are separable in consumption, savings, and earnings. Under separability, the ratio of compensated elasticities and Λ can be expressed purely in terms of income effects:

$$\frac{\bar{\zeta}_{C,\tau_Y}^H}{\bar{\zeta}_{Y,\tau_Y}^H} = \frac{\bar{\zeta}_C^I}{1 - \bar{\zeta}_Y^I}, \quad \Lambda = \bar{\zeta}_Y^I. \quad (10)$$

Our calibration approach is implemented as follows. First, as described in Appendix Section E, we estimate the Pareto coefficients directly from our data, using tax records for earnings and credit bureau data for consumption, and obtain $\bar{\rho}_Y = 2.30$ and $\bar{\rho}_C = 3.50$, implying $\bar{\rho}_Y / \bar{\rho}_C = 0.66$. This implies that consumption is less concentrated than income at the top of the distribution. Next, we measure consumption shares from the DHEA, which give $s_c = 1.00$ for the full population (as in Hellwig and Werquin, 2026) and $s_c = 0.87$ for top earners. Using these consumption shares, our estimated MPC and MPE, along with the model-based values of dS/dB , we calibrate $\bar{\zeta}_C^I$ and $\bar{\zeta}_Y^I$. The income effects on earnings are 0.29 (population, perfect smoothing), 0.50 (population, front-loading), and 0.55 (top earners, front-loading), and the income effects on consumption are 0.71

³⁵For example, Hellwig and Werquin (2026) directly use the MPE estimates of Golosov et al. (2024) to calibrate $\bar{\zeta}_Y^I$ which implicitly assumes perfect smoothing.

(population, perfect smoothing), 1.21 (population, front-loading), and 0.97 (top earners, front-loading). For the compensated earnings elasticity, $\bar{\xi}_{Y,\tau_Y}^H$, we assume it is equal to 0.40 for the full sample and 0.57 for high earners based on estimates from Gruber and Saez (2002). The inputs for the calibration are reported in Panel A of Table 6.

Panel B of Table 6 reports the optimal top tax rates obtained from Equation (6) and Equation (7). We express the savings tax as a share of gross savings, $\bar{\tau}_S/(1 + \bar{\tau}_S)$, to make it comparable with the labor income tax rate, $\bar{\tau}_Y$. The table shows that front-loading materially changes the optimal top tax rates. Imposing perfect smoothing and using estimates for the population (column (1)), the optimal labor income tax rate is 0.55, below the Saez benchmark of 0.61, while the optimal savings tax rate is 0.14. Relaxing perfect smoothing (column (2)) increases the income effect on earnings and, thus, the Saez benchmark to 0.71. It also increases Λ (from 0.29 to 0.50) and the ratio of compensated elasticities (from 1.00 to 2.45), which shifts more of the tax burden from labor income to savings, so the optimal labor income tax rate falls to 0.45 and the optimal savings tax rate rises to 0.47. Allowing behavioral responses to vary with income and using those estimated for top earners further lowers the optimal labor income tax rate to 0.35 and the optimal savings tax rate to 0.33, compared to a Saez benchmark of 0.57. Thus, in all three calibrations the optimal labor income tax is below the Saez benchmark and savings are taxed. Allowing for imperfect smoothing widens the gap with the Saez benchmark (from 0.06 to 0.26 for the population) and raises the savings tax from 0.14 to 0.47.

7 Conclusion

We use newly linked administrative data on a large sample of Canadian lottery winners to estimate the causal effects of a financial windfall on consumption, savings, and wage earnings. Using income tax records and credit bureau data, we measure consumption through credit card spending, savings through debt repayment and financial asset accumulation, and wage earnings from tax records. We then examine how each of these responses varies across the income distribution and find that high-income individuals allocated a larger share of their lottery prizes to savings and

leisure, while low-income individuals spent a larger share on consumption. We then translate the estimated effects of a one-time wealth shock to a response to a recurring flow of unearned income. We do this using a novel two-sample instrumental variables approach that combines the estimated consumption responses and the household budget constraint to estimate income effects.

Using the framework in Hellwig and Werquin (2026), we combine our estimated income effects with new estimates of the distributions of income and consumption from tax and credit report data, respectively. Under separability assumptions, the income effects allow us to characterize non-homotheticities in preferences; combining these estimates with the Pareto coefficients of income and consumption then allows us to assess heterogeneity in preferences for saving across the population. We find an important role for such preference heterogeneity in explaining differences in consumption behavior across the income distribution. This distinction matters for policy: when behavioral responses reflect differences in preferences across individuals rather than income alone, responses estimated for one part of the income distribution need not extrapolate to another. Researchers should therefore exercise caution when applying behavioral estimates across income groups, particularly when evaluating policies targeted to specific parts of the distribution, such as transfers to low-income households or changes in top marginal tax rates.

We illustrate the importance of translating from wealth effects to income effects and accounting for heterogeneity in behavioral responses across the income distribution by using our empirical estimates in three different policy counterfactuals: designing fiscal stimulus programs, evaluating a Universal Basic Income (UBI), and determining optimal top tax rates on wage earnings and savings. In each case, we find that the heterogeneous responses that we estimate matter for policy. Specifically, we find that policymakers can obtain a greater “bang for the buck” by targeting stimulus payments to low-income individuals with larger estimated MPCs. We find that a UBI requires a larger surcharge when accounting for how households allocate resources generated by a persistent policy change. Lastly, using our empirical estimates for high-income earners rather than the full population yields a lower calibrated optimal top labor income tax in part because of a substitution from income taxation to a tax on savings.

We conclude by briefly mentioning several areas of future work. First, we do not analyze why some individuals save part of their lottery prize by investing while others pay down pre-existing debt balances. Some research suggests that individuals' debt repayment decisions can be understood as solutions to an optimal portfolio problem (as in Telyukova, 2013), while other research suggests that individuals' debt repayment decisions may instead come from simpler balance payment heuristics (as in Gathergood et al., 2019). Future work could evaluate whether this type of heterogeneity in savings behavior is related to the preference heterogeneity we document in this paper.

Second, we show that our credit card spending measure is a very accurate proxy for changes in overall consumption (both in the full sample and within income groups). This suggests that our credit card spending data may be useful for estimating the consumption consequences of a wide range of public policies. For example, data limitations often make it difficult to measure changes in local consumption at fine geographic and temporal resolution, but our validation exercises suggest that credit card data may be useful for evaluating the economic effects of local policies such as place-based policies.

Third, our test of non-homothetic preferences using our income effect estimates relies on a separability assumption that can be relaxed by combining our income effect estimates with estimates of the consumption responses to permanent changes in labor taxes ($dc/d\tau$). This builds on the modeling approach in Graber et al. (2026) and motivates future work that directly estimates $dc/d\tau$. In addition to being useful for implementing a more robust test of non-homothetic preferences, these consumption response estimates may also be useful as an input into calculating the elasticity of intertemporal substitution (EIS). This is conceptually related to the insight in Chetty (2006) that labor supply responses can be used to estimate the coefficient of relative risk aversion; similarly, the consumption responses to financial windfalls and changes in labor taxes can be combined to estimate the EIS, which has been a notoriously difficult parameter to estimate credibly.

Lastly, while we find clear evidence of heterogeneous responses to financial windfalls by income and conclude that this primarily reflects differences in preferences, we do not investigate

what leads to these differences. For example, does preference heterogeneity stem from differences in self-control or time preference? This would be consistent with prior work documenting correlations between income, education, and measures of impatience (Cadena and Keys, 2015). Alternatively, differences in preferences by income might instead be due to differences in childhood environments, peer effects, or neighborhood effects. This would be consistent with recent work showing a large role for early-life circumstances in consumer financial decision-making (Malmendier and Nagel, 2011; Bakker et al., 2025). The relative importance of these different explanations remains unknown, leaving open the question of why income groups respond so differently to financial windfalls.

References

- Acharya, Sushant, Edouard Challe, and Keshav Dogra (2023). “Optimal Monetary Policy According to HANK.” *American Economic Review* 113(7), 1741–1782. <https://pubs.aeaweb.org/doi/10.1257/aer.20200239> (visited on 08/13/2025).
- Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles (2007). “The Reaction of Consumer Spending and Debt to Tax Rebates—Evidence from Consumer Credit Data.” *Journal of Political Economy* 115(6), 986–1019. <https://www.journals.uchicago.edu/doi/full/10.1086/528721> (visited on 08/13/2025).
- Agarwal, Sumit and Wenlan Qian (2014). “Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore.” *American Economic Review* 104(12), 4205–4230. <https://www.aeaweb.org/articles?id=10.1257/aer.104.12.4205> (visited on 08/13/2025).
- Agarwal, Sumit, Wenlan Qian, and Xin Zou (2021). “Thy Neighbor’s Misfortune: Peer Effect on Consumption.” *American Economic Journal: Economic Policy* 13(2), 1–25. <https://www.aeaweb.org/articles?id=10.1257/pol.20170634> (visited on 08/13/2025).
- Allen, Jason, Michael Boutros, and Benedict Guttman-Kenney (2026). “Paying Off: Evaluating Credit Card Minimum Payment Requirements.” Working Paper. <https://bguttmankenney.github.io/Public/quebec.pdf>.
- Ampudia, Miguel, Russell Cooper, Julia Le Blanc, and Guozhong Zhu (2024). “MPC Heterogeneity and the Dynamic Response of Consumption to Monetary Policy.” *American Economic Journal: Macroeconomics* 16(3), 343–388. <https://pubs.aeaweb.org/doi/10.1257/mac.20210277> (visited on 08/13/2025).
- Auclert, Adrien (2019). “Monetary Policy and the Redistribution Channel.” *American Economic Review* 109(6), 2333–2367. <https://pubs.aeaweb.org/doi/10.1257/aer.20160137> (visited on 08/13/2025).
- Baker, Scott R. (2018). “Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data.” *Journal of Political Economy* 126(4), 1504–1557. <https://www.journals.uchicago.edu/doi/10.1086/698106> (visited on 09/11/2026).
- Bakker, Trevor, Stefanie DeLuca, Eric English, Jamie Fogel, Nathaniel Hendren, and Daniel Herbst (2025). “Credit Access in the United States.” w34053. Cambridge, MA: National Bureau of Economic Research, w34053. <http://www.nber.org/papers/w34053.pdf> (visited on 09/11/2026).
- Bartik, Alexander W., Elizabeth Rhodes, David E. Broockman, Patrick K. Krause, Sarah Miller, and Eva Vivalt (2025). “The Impact of Unconditional Cash Transfers on Consumption and Household Balance Sheets: Experimental Evidence from Two US States.” <https://www.nber.org/papers/w32784> (visited on 08/13/2025). Pre-published.
- Boutros, Michael and Andrej Mijakovic (2024). “The Macroeconomic Implications of Coholding.” <https://www.bankofcanada.ca/2024/05/staff-working-paper-2024-16/> (visited on 08/13/2025).
- Briggs, Joseph, David Cesarini, Sean Chanwook Lee, Erik Lindqvist, and Robert Östling (2023). “Financial Windfalls, Portfolio Allocations, and Risk Preferences.” w31864. Cambridge, MA: National Bureau of Economic Research, w31864. <http://www.nber.org/papers/w31864.pdf> (visited on 09/11/2026).

- Brown, Jason P. (2021). “Response of Consumer Debt to Income Shocks: The Case of Energy Booms and Busts.” *Journal of Money, Credit and Banking* 53(7), 1629–1675. <https://onlinelibrary.wiley.com/doi/abs/10.1111/jmcb.12842> (visited on 08/13/2025).
- Bulman, George, Robert Fairlie, Sarena Goodman, and Adam Isen (2021). “Parental Resources and College Attendance: Evidence from Lottery Wins.” *American Economic Review* 111(4), 1201–1240. <https://www.aeaweb.org/articles?id=10.1257/aer.20171272> (visited on 08/27/2026).
- Cadena, Brian C. and Benjamin J. Keys (2015). “Human Capital and the Lifetime Costs of Impatience.” *American Economic Journal: Economic Policy* 7(3), 126–153. <https://pubs.aeaweb.org/doi/10.1257/pol.20130081> (visited on 09/11/2026).
- Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H C Sant’Anna (2024). “Difference-in-Differences with a Continuous Treatment.”
- Canada, Department of Finance (2025). “Delivering a Middle-Class Tax Cut.” <https://www.canada.ca/en/departement-finance/news/2025/05/delivering-a-middle-class-tax-cut.html> (visited on 03/30/2026).
- Cesarini, David, Erik Lindqvist, Matthew J. Notowidigdo, and Robert Östling (2017). “The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries.” *American Economic Review* 107(12), 3917–3946. <https://www.aeaweb.org/articles?id=10.1257/aer.20151589> (visited on 08/13/2025).
- Chetty, Raj (2006). “A New Method of Estimating Risk Aversion.” *American Economic Review* 96(5), 1821–1834. <https://www.aeaweb.org/articles?id=10.1257/aer.96.5.1821> (visited on 08/27/2026).
- Chetty, Raj et al. (2024). “The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data.” *The Quarterly Journal of Economics* 139(2), 829–889. <https://academic.oup.com/qje/article/139/2/829/7289247> (visited on 03/30/2026).
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber (2020). “How Did U.S. Consumers Use Their Stimulus Payments?” <https://www.nber.org/papers/w27693> (visited on 08/13/2025). Pre-published.
- Cookson, J. Anthony, Erik P. Gilje, and Rawley Z. Heimer (2022). “Shale Shocked: Cash Windfalls and Household Debt Repayment.” *Journal of Financial Economics* 146(3), 905–931. <https://www.sciencedirect.com/science/article/pii/S0304405X22002069> (visited on 08/13/2025).
- Crawley, Edmund and Andreas Kuchler (2023). “Consumption Heterogeneity: Micro Drivers and Macro Implications.”
- Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao (2017). “Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging.” *American Economic Review* 107(11), 3550–3588. <https://www.aeaweb.org/articles?id=10.1257/aer.20141313> (visited on 07/31/2026).
- Druehl, Jeppe and Alessandro Martinello (2022). “Long-Run Saving Dynamics: Evidence from Unexpected Inheritances.” *The Review of Economics and Statistics* 104(5), 1079–1095. <https://direct.mit.edu/rest/article/104/5/1079/97709/Long-Run-Saving-Dynamics-Evidence-from-Unexpected> (visited on 08/19/2026).
- Eika, Lasse, Magne Mogstad, and Ola Vestad (2020). “What Can We Learn about Household Consumption Expenditure from Data on Income and Assets?” *Journal of Public Economics* 189. <https://www.sciencedirect.com/science/article/abs/pii/S004727272030027X>.

- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik (2021). “MPC Heterogeneity and Household Balance Sheets.” *American Economic Journal: Macroeconomics* 13(4), 1–54. <https://pubs.aeaweb.org/doi/10.1257/mac.20190211> (visited on 08/13/2025).
- Feinberg, Robert M. and Daniel Kuehn (2018). “Guaranteed Nonlabor Income and Labor Supply: The Effect of the Alaska Permanent Fund Dividend.” *The B.E. Journal of Economic Analysis & Policy* 18(3), 20180042. <https://www.degruyterbrill.com/document/doi/10.1515/bejeap-2018-0042/html> (visited on 09/11/2026).
- Ferey, Antoine, Benjamin B. Lockwood, and Dmitry Taubinsky (2024). “Sufficient Statistics for Nonlinear Tax Systems with General Across-Income Heterogeneity.” *American Economic Review* 114(10), 3206–3249. <https://pubs.aeaweb.org/doi/10.1257/aer.20221053> (visited on 08/13/2025).
- Fuster, Andreas, Greg Kaplan, and Basit Zafar (2021). “What Would You Do with \$500? Spending Responses to Gains, Losses, News, and Loans.” *The Review of Economic Studies* 88(4), 1760–1795. <https://academic.oup.com/restud/article/88/4/1760/5962017> (visited on 03/30/2026).
- Ganong, Peter and Pascal Noel (2019). “Consumer Spending during Unemployment: Positive and Normative Implications.” *American Economic Review* 109(7), 2383–2424.
- Ganong, Peter and Pascal Noel (2020). “Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession.” *American Economic Review* 110(10), 3100–3138. <https://www.aeaweb.org/articles?id=10.1257/aer.20181243> (visited on 08/13/2025).
- Gathergood, John, Neale Mahoney, Neil Stewart, and Jörg Weber (2019). “How Do Individuals Repay Their Debt? The Balance-Matching Heuristic.” *American Economic Review* 109(3), 844–875. <https://pubs.aeaweb.org/doi/10.1257/aer.20180288> (visited on 09/11/2026).
- Georgarakos, Dimitris, Tullio Jappelli, Geoff Kenny, and Luigi Pistaferri (2025). “Labor Supply Response to Windfall Gains.” *Journal of Public Economics* 250, 105476. <https://www.sciencedirect.com/science/article/pii/S0047272725001744> (visited on 03/30/2026).
- Gibbs, Christa N, Benedict Guttman-Kenney, Donghoon Lee, and Scott T Nelson (2024). “Consumer Credit Reporting Data.” NBER Working Paper 32791. National Bureau of Economic Research.
- Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky (2021). “How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income.” <https://www.nber.org/papers/w29000> (visited on 08/27/2026). Pre-published.
- Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky (2024). “How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income.” *The Quarterly Journal of Economics* 139(2), 1321–1395. <https://academic.oup.com/qje/article/139/2/1321/7328960> (visited on 08/13/2025).
- Government of Canada, Statistics Canada (2017). “Household Spending by Age of Reference Person.” <https://www150.statcan.gc.ca/t1/tb11/en/tv.action?pid=1110022701> (visited on 03/30/2026).
- Government of Canada, Statistics Canada (2021). “Dictionary, Census of Population, 2021 - Complete A to Z Index.” <https://www12.statcan.gc.ca/census-recensement/2021/ref/dict/az/index-eng.cfm> (visited on 08/13/2025).


- Grabner, Michael, Morten Håvarstein, Magne Mogstad, Gaute Torsvik, and Ola L. Vestad (2026). “Substitution and Income Effects of Labor Income Taxation.” NBER Working Paper 34987. National Bureau of Economic Research. <https://www.nber.org/papers/w34987>.
- Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang (2020). “The Marginal Propensity to Consume over the Business Cycle.” *American Economic Journal: Macroeconomics* 12(2), 351–384. <https://pubs.aeaweb.org/doi/10.1257/mac.20160287> (visited on 03/30/2026).
- Gruber, Jon and Emmanuel Saez (2002). “The Elasticity of Taxable Income: Evidence and Implications.” *Journal of Public Economics* 84(1), 1–32. <https://www.sciencedirect.com/science/article/pii/S0047272701000858> (visited on 03/30/2026).
- Guttman-Kenney, Benedict and Andrés Shahidinejad (2025). “Unraveling Information Sharing in Consumer Credit Markets.” Working Paper. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4629496.
- Hellwig, Christian and Nicolas Werquin (2026). “Using Consumption Data to Derive Top Optimal Income and Capital Tax Rates.” *Journal of Political Economy*.
- Hempel, Alexander (2025). “New Estimates of Wealth Inequality in Canada.” <https://www.ssrn.com/abstract=5151980> (visited on 08/07/2025). Pre-published.
- Henry, Christopher, Matthew Shimoda, and Julia Zhu (2022). “2021 Methods-of-Payment Survey Report.” Staff Discussion Paper 2022–23. Bank of Canada. <https://www.bankofcanada.ca/2022/12/staff-discussion-paper-2022-23/>.
- Henry, Christopher S., Kim Huynh, and Angelika Welte (2018). “2017 Methods-of-Payment Survey Report.” <https://www.bankofcanada.ca/2018/12/staff-discussion-paper-2018-17/> (visited on 08/13/2025).
- Hisnanick, John J. and Andreas Kern (2018). “The 2008 Tax Rebate and US Household Debt.” *Applied Economics Letters* 25(9), 592–596. <https://doi.org/10.1080/13504851.2017.1349283> (visited on 08/13/2025).
- Huynh, Kim, Gradon Nicholls, and Mitchell Nicholson (2019). “2018 Merchant Acceptance Survey.” Staff Analytical Note 2019–31. Bank of Canada. <https://www.bankofcanada.ca/2019/12/staff-analytical-note-2019-31/>.
- Imbens, Guido W., Donald B. Rubin, and Bruce I. Sacerdote (2001). “Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players.” *American Economic Review* 91(4), 778–794. <https://www.aeaweb.org/articles?id=10.1257/aer.91.4.778> (visited on 08/13/2025).
- Jacob, Brian A and Jens Ludwig (2012). “The Effects of Housing Assistance on Labor Supply: Evidence from a Voucher Lottery.” *American Economic Review* 102(1), 272–304. <https://pubs.aeaweb.org/doi/10.1257/aer.102.1.272> (visited on 09/11/2026).
- Jappelli, Tullio and Luigi Pistaferri (2010). “The consumption response to income changes.” *Annual Review of Economics* 2(1), 479–506.
- Jappelli, Tullio and Luigi Pistaferri (2014). “Fiscal Policy and MPC Heterogeneity.” *American Economic Journal: Macroeconomics* 6(4), 107–136. <https://pubs.aeaweb.org/doi/10.1257/mac.6.4.107> (visited on 03/30/2026).
- Jappelli, Tullio and Luigi Pistaferri (2020). “Reported MPC and unobserved heterogeneity.” *American Economic Journal: Economic Policy* 12(4), 275–297.
- Jones, Damon and Ioana Marinescu (2022). “The Labor Market Impacts of Universal and Permanent Cash Transfers: Evidence from the Alaska Permanent Fund.” *American Economic*

- Journal: Economic Policy* 14(2), 315–340. <https://pubs.aeaweb.org/doi/10.1257/pol.20190299> (visited on 09/11/2026).
- Keynes, John Maynard (1936–2018). *The General Theory of Employment, Interest, and Money*. Springer. 430 pp.
- Koşar, Gizem, Davide Melcangi, Laura Pilossoph, and David Wiczer (2023). “Stimulus through Insurance: The Marginal Propensity to Repay Debt.” IZA - Institute of Labor Economics. <https://www.jstor.org/stable/resrep67531> (visited on 08/14/2025).
- Kuhn, Peter, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn (2011). “The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery.” *The American Economic Review* 101(5), 2226–2247. <https://www.jstor.org/stable/23045637> (visited on 08/13/2025).
- Laibson, David, Peter Maxted, and Benjamin Moll (2022). “A Simple Mapping from MPCs to MPXs.” NBER Working Paper 29664. National Bureau of Economic Research. <https://www.nber.org/papers/w29664>.
- Lewis, Daniel, Davide Melcangi, and Laura Pilossoph (2026). “Latent Heterogeneity in the Marginal Propensity to Consume.” *The Review of Economic Studies*, rdaf102. <https://doi.org/10.1093/restud/rdaf102> (visited on 03/26/2026).
- Malmendier, Ulrike and Stefan Nagel (2011). “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?” *The Quarterly Journal of Economics* 126(1), 373–416. <https://academic.oup.com/qje/article-lookup/doi/10.1093/qje/qjq004> (visited on 09/11/2026).
- Malthus, Thomas (1798). “An Essay on the Principle of Population, as It Affects the Future Improvement of Society, with Remarks on the Speculations of Mr. Godwin, M. Condorcet, and Other Writers.” Morrison, Kevin A. *Charity and Philanthropy in Nineteenth-Century Britain*. 1st ed. London: Routledge, 5–7. <https://www.taylorfrancis.com/books/9781003056508/chapters/10.4324/9781003056508-2> (visited on 07/30/2026).
- Marshall, Katherine (2011). “Gambling 2011.” <https://www.safec.ca/userfiles/file/%C3%89conomie/EN/04%20Studies/43%20Work-Revenue-Workforce/431%20Employement%20and%20revenue/2011-PerspectivesLabourIncome-Gambling.pdf> (visited on 08/13/2025).
- Olafsson, Arna and Michaela Pagel (2018). “The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software.” *The Review of Financial Studies* 31(11), 4398–4446. <https://academic.oup.com/rfs/article/31/11/4398/4986471> (visited on 09/11/2026).
- Parker, Jonathan A., Jake Schild, Laura Erhard, and David Johnson (2022). “Household Spending Responses to the Economic Impact Payments of 2020: Evidence from the Consumer Expenditure Survey.” <https://www.nber.org/papers/w29648> (visited on 08/13/2025). Pre-published.
- Patterson, Christina (2023). “The Matching Multiplier and the Amplification of Recessions.” *American Economic Review* 113(4), 982–1012. <https://pubs.aeaweb.org/doi/10.1257/aer.20210254> (visited on 08/13/2025).
- Picchio, Matteo, Sigrid Suetens, and Jan C. van Ours (2018). “Labour Supply Effects of Winning a Lottery.” *The Economic Journal* 128(611), 1700–1729. <https://doi.org/10.1111/eoj.12465> (visited on 08/13/2025).
- Rotermann, Michelle and Heather Gilmour (2022). “Who Gambles and Who Experiences Gambling Problems in Canada.”
- Saez, Emmanuel (2001). “Using Elasticities to Derive Optimal Income Tax Rates.” *The Review of Economic Studies* 68(1), 205–229. <https://doi.org/10.1111/1467-937X.00166> (visited on 03/26/2026).

- Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod (2010). “3 Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications.” *Tax Policy and the Economy* 24(1), 69–110. <https://www.journals.uchicago.edu/doi/full/10.1086/649829> (visited on 08/13/2025).
- Statistics Canada (2026). *Distributions of Household Economic Accounts, Income, Consumption and Saving, by Characteristic*. Government of Canada. <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3610058701> (visited on 09/11/2026).
- Telyukova, Irina A. (2013). “Household Need for Liquidity and the Credit Card Debt Puzzle.” *The Review of Economic Studies* 80(3), 1148–1177. <https://academic.oup.com/restud/article-lookup/doi/10.1093/restud/rdt001> (visited on 09/11/2026).
- Tompkins, Michael and Viktoria Galociova (2019). “Canadian Payment Methods and Trends: 2019.” Report. Payments Canada. https://www.payments.ca/sites/default/files/2022-08/PaymentsCanada_2019CanadianPaymentsMethodsAndTrendsReport_En.pdf.
- Vivalt, Eva, Elizabeth Rhodes, Alexander Bartik, David Broockman, and Sarah Miller (2025). “The Employment Effects of a Guaranteed Income: Experimental Evidence from Two U.S. States.”

Figures and Tables


Figure 1: Validating Credit Card Spending as a Consumption Proxy




Notes: This figure compares annual growth in our measure of consumption (blue circles), constructed from credit card spending in our Equifax data, with a benchmark from Statistics Canada (orange crosses), overall and by income group, expressed relative to 2016. The Equifax data cover a 0.1% sample of all Canadians. Credit card spending and its mapping to overall consumption are described in Section 2.2 and Appendix Section B. The benchmark combines data from Statistics Canada's Survey of Household Spending (SHS) and Distribution of Household Economic Accounts, as detailed in Appendix Section B. The SHS has been conducted biennially since 2017, so the benchmark is unavailable for 2018 and 2020. Because income is not observed in the credit bureau data, household income is proxied by the forward sortation area-level average from the closest preceding census (forward sortation areas are comparable in size to five-digit U.S. ZIP codes).

Figure 2: Wealth Effects on Consumption out of \$1 of Lottery Winnings (Credit Bureau Data)

(a) Quarterly wealth effects on credit card spending and overall consumption




(b) 2-year wealth effect on overall consumption, by income




Notes: This figure presents event-study estimates of Equation (1) for the effect of \$1 of lottery winnings on consumption. Panel (a) plots quarterly wealth effects on credit card spending (dashed line) and overall consumption (solid line), normalized to zero in the quarter preceding the win. Panel (b) sums the quarterly effects on overall consumption over the two years before (pre-period) and after (post-period) the win, overall and by income group. We include the partially treated quarter of the win in the post-period to capture the on-impact effect. We report the simple average of the quarterly coefficients within each period (panel a) and the post-period sums (panel b), with standard errors in parentheses. Income groups are estimated separately for winners below and above median pre-win income. The full-sample estimates in panel (b) are constructed as equally weighted averages of the below- and above-median estimates. The sample comprises 11,703 winners from 2015Q1 to 2020Q1, observed for the eight quarters before and after the win. Credit card spending and its mapping to overall consumption are described in Section 2.2 and Appendix Section B. Because income is not observed in the credit bureau data, household income is proxied by the dissemination area-level average from the closest preceding census (dissemination areas are comparable in size to nine-digit U.S. ZIP codes). Shaded regions (panel a) and whiskers (panel b) are 95% confidence intervals from standard errors clustered at the individual level.

Figure 3: Wealth Effect on Saving via Debt Repayment out of \$1 of Lottery Winnings (Credit Bureau Data)




Notes: This figure presents event-study estimates of Equation (1) for the effect of \$1 of lottery winnings on total debt balance. Negative values are reductions in total debt balance and therefore reflect saving via debt repayment. Panel (a) plots quarterly wealth effects on saving via debt repayment, normalized to zero in the quarter preceding the win. Panel (b) averages the quarterly effects on saving via debt repayment over the two years before (pre-period) and after (post-period) the win, overall and by income group. We report the simple average of the quarterly coefficients within each period (panel a) and the post-period averages (panel b), with standard errors in parentheses. Income groups are estimated separately for winners below and above median pre-win income. The full-sample estimates in panel (b) are constructed as equally weighted averages of the below- and above-median estimates. The sample comprises 11,703 winners from 2015Q1 to 2020Q1, observed for the eight quarters before and after the win. The construction of total debt balance is described in Section 2.2. Because income is not observed in the credit bureau data, household income is proxied by the dissemination area-level average from the closest preceding census (dissemination areas are comparable in size to nine-digit U.S. ZIP codes). Shaded regions (panel a) and whiskers (panel b) are 95% confidence intervals from standard errors clustered at the individual level.

Figure 4: Wealth Effect on Saving via Investment from \$1 of Lottery Winnings (Tax Data)




Notes: This figure presents event-study estimates of Equation (1) for the effect of \$1 of lottery winnings on saving via investment. Panel (a) plots annual wealth effects on saving via investment, normalized to zero in the year preceding the win. Post-event values are cumulative effects on savings via investment from the year of the win through year t . Panel (b) reports the simple average of the annual pre-win effects (pre-period) and the two-year savings response (post-period), overall and by income group. The post-period estimate aggregates each component as described in Section 2.3. We report the simple average of the plotted coefficients within each period (panel a) and the post-period savings response (panel b), with standard errors in parentheses. Post-period averages and savings responses include only fully treated years, and thus omit the year of the win. Income groups are estimated separately for winners below and above median pre-win income. The full-sample estimates in panel (b) are constructed as equally weighted averages of the below- and above-median estimates. The sample comprises 30,000 winners from 2004 to 2017. The savings response combines taxable capital gains, interest and other investment income, the fair market value of Tax-Free Savings Accounts (TFsAs), and net contributions to Registered Retirement Savings Plans (RRSPs); its construction is described in Section 2.3. Figure A.3 reports the event-study coefficients for the individual elements of the savings response. Shaded regions (panel a) and whiskers (panel b) are 95% confidence intervals from standard errors clustered at the individual level.

Figure 5: Wealth Effect on Earnings out of \$1 of Lottery Winnings (Tax Data)




Notes: This figure presents event-study estimates of Equation (1) for the effect of \$1 of lottery winnings on total wage earnings. Panel (a) plots annual wealth effects on total wage earnings, normalized to zero in the year preceding the win. Panel (b) sums the effects on total wage earnings over the pre-event years (pre-period) and over the two years following the win (post-period), overall and by income group. We report the simple average of the annual coefficients within each period (panel a) and the post-period sums (panel b), with standard errors in parentheses. Post-period averages and sums include only fully treated years, and thus omit the year of the win. Income groups are estimated separately for winners below and above median pre-win income. The full-sample estimates in panel (b) are constructed as equally weighted averages of the below- and above-median estimates. The sample comprises 30,000 winners from 2004 to 2017. The construction of total wage earnings is described in Section 2.3. Shaded regions (panel a) and whiskers (panel b) are 95% confidence intervals from standard errors clustered at the individual level.

Figure 6: Accounting for \$1 of Lottery Winnings after 5 Years




Notes: This figure decomposes the use of \$1 of lottery winnings over the five years following the win, based on estimates of Equation (1). Each bar shows the 5-year consumption (Figure 2), post-tax wage earnings (Figures 5 and A.5), and savings (Figures 3 and 4) effects out of lottery winnings, for the full sample and separately for winners with below-median and above-median pre-win income. The post-tax earnings effect nets the reduction in taxes paid out of the reduction in wage earnings; we attribute the entire tax response to wage earnings because the effect on capital taxes is negligible (Table A.4). The savings effect out of lottery winnings combines debt repayment and investment. The 5-year effect is the sum of the period coefficients for flows (consumption, wage earnings, and taxes) and the average of the coefficients for stocks (debt balances and investment). For outcomes using the credit bureau data (consumption and savings via debt repayment), we estimate effects for only eight quarters post-win due to the shorter sample window and we extrapolate to 5-year effects by holding the final quarterly coefficient constant over the remaining quarters. Outcomes measured in the tax data (wage earnings, taxes, and savings via investment) are estimated for each of the five years following the win. Because income is not observed in the credit bureau data, household income is proxied by the dissemination area-level average from the closest preceding census (dissemination areas are comparable in size to nine-digit U.S. ZIP codes). The value to the right of each bar reports the total share accounted for. Table 2 reports the underlying point estimates by income group. Figure 7 repeats the full-sample accounting using expenditure in place of consumption, which accounts for 89 cents of the dollar.

Figure 7: Accounting for \$1 of Lottery Winnings after 5 Years: Consumption vs. Expenditure




Notes: This figure decomposes the use of \$1 of lottery winnings over the five years following the win, based on estimates of Equation (1). Each bar shows the 5-year post-tax wage earnings (Figures 5 and A.5), and savings effects out of lottery winnings (Figures 3 and 4); the top bar adds the 5-year effect on consumption (Figure 2) to these effects, while the bottom bar adds the 5-year effect on expenditure (Figure B.2(b)). The post-tax earnings effect nets the reduction in taxes paid out of the reduction in wage earnings; we attribute the entire tax response to wage earnings because the effect on capital taxes is negligible (Table A.4). The savings effect out of lottery winnings combines debt repayment and investment. The 5-year effect is the sum of the period coefficients for flows (consumption, expenditure, wage earnings, and taxes) and the average of the coefficients for stocks (debt balances and investment). For outcomes using the credit bureau data (consumption, expenditure, and savings via debt repayment), we estimate effects for only eight quarters post-win due to the shorter sample window and we extrapolate to 5-year effects by holding the final quarterly coefficient constant over the remaining quarters. Outcomes measured in the tax data (wage earnings, taxes, and savings via investment) are estimated for each of the five years following the win. The value to the right of each bar reports the total share accounted for and Table 2 reports individual point estimates for consumption, wage earnings, and savings via debt repayment and investment.

Figure 8: Marginal Propensities to Consume (MPC), Earn (MPE), and Pay Taxes (MPT) out of Unearned Income (5-Year Horizon)



Notes: This figure presents the effect of lottery winnings on unearned income and the resulting MPC, MPE, and MPT out of unearned income, for the full sample and separately for winners with below-median and above-median pre-win income. We construct these estimates using the two-sample instrumental variables (IV) estimator described in Section 5, which takes as inputs the wealth effects on consumption (Figure 2), wage earnings (Figure 5), and labor taxes (Figure A.5 and Table A.4), each averaged into a mean annual effect over the five years following the win. The household budget constraint recovers the effect of winnings on unearned income from these three responses, and each response is then rescaled by it. The left panel plots this first stage, the mean annual effect of lottery winnings on unearned income per \$1 won. The right panel plots the resulting marginal propensities out of \$1 of unearned income. Full-sample estimates in both panels are constructed as equally weighted averages of the below- and above-median income estimates. For consumption, which uses the credit bureau data, we estimate effects for only eight quarters post-win due to the shorter sample window and we extrapolate to a 5-year effect by holding the final quarterly coefficient constant over the remaining quarters. Outcomes measured in the tax data (wage earnings and taxes) are estimated for each of the five years following the win. Because income is not observed in the credit bureau data, household income is proxied by the dissemination area-level average from the closest preceding census (dissemination areas are comparable in size to nine-digit U.S. ZIP codes). Whiskers are 95% confidence intervals from delta-method standard errors, which treat the wealth effects on consumption, earnings, and taxes as independent. Table 3 reports individual estimates for each marginal propensity out of unearned income.

Figure 9: Comparing Two-Sample IV Estimator to the Annuitization Approach (5-Year Horizon)



Notes: This figure compares the MPC, MPE, and MPT out of unearned income constructed using the two-sample instrumental variables (IV) estimator described in Section 5 (light bars) with those constructed using the annuitization approach common in the literature, described in Appendix Section C.1 (dark bars). Both approaches take as inputs the wealth effects on wage earnings (Figure 5) and labor taxes (Figure A.5 and Table A.4), each averaged into a mean annual effect over the five years following the win. The annuitization approach assumes perfect consumption smoothing, so that winners spread each dollar of lottery wealth evenly over their remaining lifetime. It therefore converts the responses to lottery wealth into responses to a permanent income flow using the annuity value of a dollar of wealth, and recovers the implied MPC as the residual of the budget-constraint identity $MPC - MPE + MPT = 1$. For annuitization, we assume an interest rate of 3%, close to the average yield on long-term Government of Canada bonds over our sample period, and a life expectancy of 80 years, approximately the Canadian life expectancy over the sample period. The two-sample IV estimator instead estimates the effect of winnings on unearned income directly, without imposing any behavioral assumptions. It recovers this first stage from the household budget constraint, using the wealth effect on consumption (Figure 2) as an additional input, and rescales each wealth effect by it. Whiskers are 95% confidence intervals from delta-method standard errors, which treat the wealth effects on consumption, earnings, and taxes as independent. The two-sample IV estimates are the same as those in Figure 8.

Table 1: Summary Statistics

Outcome	Sample			
	Lottery Winners		General Population*	
	Mean (1)	Std. Dev. (2)	Mean (3)	Std. Dev. (4)
<i>Panel A: Credit Bureau Data (quarter prior to winning)</i>				
Lottery winnings (\$)	8,363.23	24,808.19		
<i>Demographics</i>				
Age (years)	48.02	14.36	48.20	14.02
<i>Credit Portfolio</i>				
Total quarterly credit card spending (\$)	2,959.85	5,855.20	3,614.90	6,461.68
Share with a credit card	0.70		0.83	
Share with a mortgage	0.27		0.32	
Total outstanding debt balance (\$)	107,485.06	169,383.31	103,435.20	171,714.70
Credit card outstanding balance (\$)	4,775.74	9,207.63	3,511.81	6,800.48
Mortgage outstanding balance (\$)	72,617.95	145,198.94	75,197.68	152,180.00
<i>Credit Worthiness / Financial Well-being</i>				
Risk score	726.80	109.72	754.95	99.00
Share in delinquency	0.07		0.04	
Share in collection	0.01		0.00	
Share in bankruptcy	0.02		0.01	
Unique winners	11,703			
Winner cohorts (quarter of win)	2015Q1 – 2020Q1			
Sample period	2013Q1 – 2022Q1			
<i>Panel B: Tax Data (year prior to winning)</i>				
<i>Demographics</i>				
Age (years)	49.63	15.04	45.08	21.54
Male	0.64		0.50	
<i>Employment, Income, and Taxes</i>				
Employed	0.75		0.70	
Total taxes paid (\$)	10,793.90	20,302.80	8,728.02	39,394.24
Wage earnings (\$)	43,036.24	47,369.52	33,416.77	44,719.54
RRSP contributions (\$)	1,930.31	4,729.51	1,493.67	4,163.12
Interest and other investment income (\$)	312.09	1,363.15	484.96	2,126.40
Taxable capital gains (\$)	294.93	2,243.60	439.36	3,228.29
Unique winners	30,000			
Winner cohorts (year of win)	2004 – 2017			
Sample period	2001 – 2018			

Notes: This table reports summary statistics for lottery winners and comparison samples from the general population, measured in the period prior to winning. *Panel A* draws on credit bureau data for winners in the quarter prior to winning and a 0.1% random sample of the Canadian population; variable definitions are provided in Section 2.2. *Panel B* draws on Statistics Canada tax records for winners in the calendar year prior to winning; variable definitions are provided in Section 2.3. All dollar variables in *Panel B* are measured unconditionally, with individuals not reporting any amount of a given type entering as zeroes. The number of unique winners in *Panel B* is rounded to the nearest 1,000 to comply with Statistics Canada confidentiality requirements. *The general population sample for *Panel B* consists of tax filers from the same province as the lottery winners.

Table 2: Wealth Effects out of Lottery Winnings

	Sample		
	Full Sample	Below Median Income	Above Median Income
	(1)	(2)	(3)
<i>Panel A: 2-year Effects</i>			
Overall consumption	0.20 (0.04)	0.25 (0.06)	0.15 (0.04)
Total wage earnings	−0.09 (0.01)	−0.04 (0.01)	−0.14 (0.02)
Total labor taxes	−0.02 (0.01)	−0.01 (0.00)	−0.04 (0.01)
Savings	0.28 (0.05)	0.24 (0.03)	0.32 (0.09)
<i>via debt repayment</i>	0.10 (0.03)	0.07 (0.02)	0.12 (0.05)
<i>via investment</i>	0.18 (0.04)	0.16 (0.02)	0.20 (0.08)
<i>Panel B: 5-year Effects</i>			
Overall consumption	0.38 (0.07)	0.50 (0.12)	0.26 (0.08)
Total wage earnings	−0.21 (0.02)	−0.09 (0.02)	−0.33 (0.04)
Total labor taxes	−0.06 (0.01)	−0.02 (0.00)	−0.10 (0.02)
Savings	0.19 (0.04)	0.18 (0.03)	0.20 (0.07)
<i>via debt repayment</i>	0.09 (0.03)	0.06 (0.02)	0.12 (0.05)
<i>via investment</i>	0.10 (0.03)	0.12 (0.02)	0.08 (0.06)

Notes: This table reports event-study estimates of Equation (1) for the effect of \$1 of lottery winnings on overall consumption, wage earnings, labor taxes, and savings via debt repayment and investment, over the two years (Panel A) and five years (Panel B) following the win, for the full sample and by income group. Consumption, earnings, and taxes are flows, so each entry sums the post-win coefficients over the horizon; debt repayment and investment are stocks, so each entry is the simple average of the post-win coefficients over the horizon. Savings via debt repayment is the reduction in total debt balance, and the Savings row sums its two components. For outcomes using the credit bureau data (consumption and savings via debt repayment), we estimate effects for only eight quarters post-win due to the shorter sample window and, in Panel B, we extrapolate to 5-year effects by holding the final quarterly coefficient constant over the remaining quarters. Outcomes measured in the tax data (wage earnings, taxes, and savings via investment) are estimated for each year following the win. Income groups are estimated separately for winners below and above median pre-win income. The full-sample estimates are constructed as equally-weighted averages of the below- and above-median estimates. The underlying event studies are shown in Figure 2 (consumption), Figures 5 and A.5, and Table A.4 (wage earnings and labor taxes), and Figures 3 and 4 (savings via debt repayment and via investment). Credit card spending and its mapping to overall consumption are described in Section 2.2 and Appendix Section B. Because income is not observed in the credit bureau data, household income is proxied by the dissemination area-level average from the closest preceding census (dissemination areas are comparable in size to nine-digit U.S. ZIP codes). Standard errors, clustered at the individual level, are in parentheses. The standard error of the “Savings” rows assumes independence between the estimates for savings via debt repayment and via investment, which are estimated in different datasets.

Table 3: Marginal Propensities out of Unearned Income

	Sample		
	Full Sample	Below Median Income	Above Median Income
	(1)	(2)	(3)
<i>Panel A: 2-year Horizon</i>			
Unearned income (first stage)	0.13 (0.02)	0.14 (0.03)	0.12 (0.02)
Overall consumption	0.75 (0.05)	0.89 (0.04)	0.61 (0.09)
Total wage earnings	−0.35 (0.05)	−0.14 (0.04)	−0.56 (0.10)
Total labor taxes	−0.10 (0.04)	−0.03 (0.01)	−0.17 (0.07)
<i>Panel B: 5-year Horizon</i>			
Unearned income (first stage)	0.10 (0.02)	0.11 (0.02)	0.10 (0.02)
Overall consumption	0.71 (0.05)	0.88 (0.04)	0.54 (0.09)
Total wage earnings	−0.42 (0.06)	−0.16 (0.04)	−0.68 (0.12)
Total labor taxes	−0.13 (0.04)	−0.03 (0.01)	−0.22 (0.07)

Notes: This table reports the effect of lottery winnings on unearned income and the resulting MPC, MPE, and MPT out of unearned income, for the full sample and separately for winners with below and above median pre-win income, over the two years (Panel A) and five years (Panel B) following the win. We construct these estimates using the two-sample instrumental variables (IV) estimator described in Section 5, which takes as inputs the wealth effects on consumption (Figure 2), wage earnings (Figure 5), and labor taxes (Figure A.5 and Table A.4), each averaged into a mean annual effect over the horizon of the panel. The household budget constraint recovers the effect of winnings on unearned income from these three responses, and each response is then rescaled by it. The first row of each panel reports this first stage, the mean annual effect of lottery winnings on unearned income per \$1 won. The remaining rows report the resulting marginal propensities out of \$1 of unearned income. Full-sample estimates are constructed as equally-weighted averages of the below- and above-median income estimates. For consumption, which uses the credit bureau data, we estimate effects for only eight quarters post-win due to the shorter sample window and, in Panel B, we extrapolate to a 5-year effect by holding the final quarterly coefficient constant over the remaining quarters. Outcomes measured in the tax data (wage earnings and taxes) are estimated for each year following the win. Because income is not observed in the credit bureau data, household income is proxied by the dissemination area-level average from the closest preceding census (dissemination areas are comparable in size to nine-digit U.S. ZIP codes). Standard errors, in parentheses, are computed by the delta method and treat the wealth effects on consumption, earnings, and taxes as independent. Figure 8 plots the 5-year horizon estimates.

Table 4: Fiscal Stimulus Simulation: 2-Year Consumption Impact of a \$1,000 Stimulus Package

Income Cutoff	Wealth Effect	Consumption from \$1,000	Eligible Population	Change in Consumption by Response			Fiscal Cost Per \$1 of Consumption by Response		
				Homogeneous	Heterogeneous	% Diff.	Homogeneous	Heterogeneous	% Diff.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
None	0.20	\$201.04	34.01 mil.	\$6.84 bil.	\$6.84 bil.	0%		\$4.97	0%
90th pctl.	0.21	\$206.94	30.61 mil.	\$6.15 bil.	\$6.33 bil.	3%		\$4.83	−3%
75th pctl.	0.22	\$218.73	25.51 mil.	\$5.13 bil.	\$5.58 bil.	9%	\$4.97	\$4.57	−8%
50th pctl.	0.25	\$254.12	17.00 mil.	\$3.42 bil.	\$4.32 bil.	26%		\$3.94	−21%
25th pctl.	0.35	\$353.28	8.50 mil.	\$1.71 bil.	\$3.00 bil.	76%		\$2.83	−43%

Notes: This table simulates the 2-year aggregate consumption response to a hypothetical \$1,000 stimulus payment distributed to Canadian adults under various income eligibility thresholds. Column (1) is the 2-year wealth effect of \$1 of lottery winnings on the consumption of the eligible population, constructed from the estimates in Table 2, and column (2) applies this wealth effect to a \$1,000 stimulus payment. Column (3) is the eligible population. Columns (4) and (7) apply the population-average wealth effect, calculated as the equally-weighted average of the estimates for winners with below and above median pre-win income, thereby ignoring heterogeneity across the income distribution. By contrast, columns (5) and (8) apply the average wealth effect of the population eligible under each threshold, reported in column (1). For the 25th percentile threshold, we take the below-median wealth effect on credit card spending and rescale it by the bottom quartile's credit card share of total consumption (Figure B.1), as we cannot reject equality of the wealth effects on credit card spending for winners with below median income and winners in the bottom income quartile. Columns (7) and (8) report the fiscal cost per \$1 of consumption, the total government outlay ($\$1,000 \times$ column (3)) divided by the change in aggregate consumption in columns (4) and (5), respectively. Columns (6) and (9) report the percentage difference between the heterogeneous and homogeneous responses. Because income is not observed in the credit bureau data, household income is proxied by the dissemination area-level average from the closest preceding census (dissemination areas are comparable in size to nine-digit U.S. ZIP codes).

Table 5: Calibration of the UBI Surcharge Formula

	No heterogeneity				Heterogeneity	
	Perfect smoothing $\partial N/\partial B = 1$		Front-loaded $\partial N/\partial B > 1$		Front-loaded $\partial N/\partial B > 1$	
	No savings tax (1)	Savings tax (2)	No savings tax (3)	Savings tax (4)	No savings tax (5)	Savings tax (6)
<i>Panel A: Inputs</i>						
Mean labor income, $\mathbb{E}[Y]$	42,704	42,704	42,704	42,704	42,704	42,704
Labor income tax rate, τ_Y	0.30	0.30	0.30	0.30	0.30	0.30
Savings tax rate, τ_S	0.00	0.25	0.00	0.25	0.00	0.25
Earnings response to transfer, $\mathbb{E}[\partial Y/\partial B]$	-0.42	-0.42	-0.72	-0.72	-0.69	-0.69
Below-median income					-0.29	-0.29
Above-median income					-1.09	-1.09
Savings response to transfer, $\mathbb{E}[\partial S/\partial B]$	0.00	0.00	-0.71	-0.71	-0.72	-0.72
Below-median income					-0.86	-0.86
Above-median income					-0.58	-0.58
Weighted earnings response, $\mathbb{E}[Y \zeta_{Y, \tau_Y}^U]$	4,616	4,616	4,616	4,616	3,842	3,842
Weighted savings response, $\mathbb{E}[S \zeta_{S, \tau_Y}^U]$	0	0	-21,315	-21,315	-19,530	-19,470
<i>Panel B: Surcharge and Implied Responses</i>						
Surcharge (p.p. per \$1,000 of transfer), $\partial \delta/\partial B$	2.76	2.76	2.99	2.88	2.94	2.89
Δ wage earnings (\$ per \$1 of transfer), dY/dB	-0.60	-0.60	-0.92	-0.91	-0.85	-0.85
Δ savings (\$ per \$1 of transfer), dS/dB	0.00	0.00	0.20	0.17	0.10	0.08

Notes: This table calibrates the surcharge formula in Equation (5) at $B = 0$, where $\delta(0) = 0$. Panel B reports the surcharge in percentage points per \$1,000 of transfer ($10^5 \times \partial \delta/\partial B$) and the total changes in wage earnings and savings per dollar of transfer from Equations (21) and (22), which add the response to the surcharge to the response to the transfer (Appendix Section D.1). Columns (1) to (4) apply full-sample responses to the whole population. Columns (5) and (6) use separate responses for winners with below- and above-median pre-win income and weight the two groups equally. Columns (1), (3), and (5) set $\tau_S = 0$, the case in Golosov et al. (2021). Mean labor income is compensation of employees per consumption unit from Statistics Canada’s Distributions of Household Economic Accounts. Columns (1) and (2) assume perfect smoothing: the transfer is used as it arrives, $\partial N/\partial B = 1$, so the earnings response is the five-year marginal propensity to earn (MPE) in Table 3 and the savings response is zero. Columns (3) to (6) assume front-loaded resource use: households discount the future faster than the interest rate and can borrow against future transfers, so in the first years they use more than the transfer, $\partial N/\partial B > 1$. We take $\partial N_t/\partial B$ from Appendix Section C.3, using the discount rates in Appendix Section C.2 (18.4 percent for the full sample, 22.8 percent below the median, and 15.1 percent above it). The earnings response is $\mathbb{E}[\text{MPE}_t \times \partial N_t/\partial B]$ and the savings response is $1 - \mathbb{E}[\partial N_t/\partial B]$, averaged over post-win years $t = 1$ to 5 and over ages 50 to 80, weighting ages by Statistics Canada’s Life Expectancy Table (Table 13-10-0114-01) for 2022–2024. The age range starts at 50, the mean age at win. The indented rows report the group-specific responses. The weighted earnings response is mean labor income multiplied by the uncompensated earnings elasticity, which the Slutsky equation gives as the compensated elasticity minus $(1 - \tau_Y)|\text{MPE}|$. We use the five-year MPE in Table 3, not scaled by $\partial N/\partial B$, and compensated elasticities from Gruber and Saez (2002) of 0.40 for the full sample, 0.18 below the median, and 0.57 above it. The implied uncompensated elasticities are 0.11, 0.07, and 0.10. Columns (1) to (4) use the full-sample elasticity. Columns (5) and (6) compute the product separately for each income group, with mean labor income of \$18,540 below the median and \$66,867 above it, and average the two. The weighted savings response follows Appendix Section D.2. For the full sample, the compensated savings response is zero up to rounding, since the five-year MPC and $(1 - \tau_Y)|\text{MPE}|$ sum to one, so we set it to zero. The weighted savings response is therefore zero under perfect smoothing and does not depend on τ_S in columns (1) to (4).

Table 6: Calibration of Optimal Top Tax Formulas


	Population		Top earners
	Perfect smoothing $\partial N/\partial B = 1$ (1)	Front-loaded $\partial N/\partial B > 1$ (2)	Front-loaded $\partial N/\partial B > 1$ (3)
<i>Panel A: Inputs</i>			
Earnings Pareto coefficient, $\bar{\rho}_Y$	2.30	2.30	2.30
Consumption Pareto coefficient, $\bar{\rho}_C$	3.50	3.50	3.50
Earnings income effect, $\bar{\zeta}_Y^I$	0.29	0.50	0.55
Consumption income effect, $\bar{\zeta}_C^I$	0.71	1.21	0.97
Compensated earnings elasticity, $\bar{\zeta}_{Y,\tau_Y}^H$	0.40	0.40	0.57
Compensated consumption elasticity, $\bar{\zeta}_{C,\tau_Y}^H$	0.40	0.98	1.22
Scaling term, $\Lambda = \bar{\zeta}_Y^I$	0.29	0.50	0.55
<i>Panel B: Optimal Top Tax Rates</i>			
Saez benchmark, $\bar{\tau}_Y^{\text{Saez}}$	0.61	0.71	0.57
Labor income tax, $\bar{\tau}_Y$	0.55	0.45	0.35
Savings tax on gross savings, $\bar{\tau}_S/(1 + \bar{\tau}_S)$	0.14	0.47	0.33

Notes: This table calibrates the optimal top tax rates on labor income and savings in Equation (6) and Equation (7), together with the Saez (2001) benchmark in Equation (9). Panel A reports the sufficient statistics that enter the formulas and Panel B reports the implied tax rates. We express the savings tax as a share of gross savings, $\bar{\tau}_S/(1 + \bar{\tau}_S)$, so that it is comparable with $\bar{\tau}_Y$. Column (1) assumes perfect smoothing, $\partial N/\partial B = 1$, so the responses of earnings and consumption to a \$1 per-period transfer are the full-sample five-year marginal propensities to earn and consume in Table 3. Columns (2) and (3) instead scale the annual marginal propensities by $\partial N_t/\partial B$ from the life-cycle model, as in columns (3) to (6) of Table 5. Column (2) uses the full-sample responses, which apply the full-sample propensities and discount rate to the whole population. Column (3) applies the responses of winners with above-median pre-win income to top earners. The Pareto coefficients are estimated in Appendix Section E. The compensated earnings elasticities are from Gruber and Saez (2002): 0.40 for the population and 0.57 for incomes above \$100,000. The income effects are $\bar{\zeta}_Y^I = (1 - \tau_Y)|\partial Y/\partial B|$ and $\bar{\zeta}_C^I = (\partial C/\partial B)/s_c$, where τ_Y is 0.30 for the population and 0.50 for top earners, and the consumption share s_c , from Statistics Canada's Distributions of Household Economic Accounts, is 1.00 for the population and 0.87 for top earners. Under separable preferences (Equation (10)), $\Lambda = \bar{\zeta}_Y^I$ and the compensated consumption elasticity is $\bar{\zeta}_{C,\tau_Y}^H = \bar{\zeta}_{Y,\tau_Y}^H \bar{\zeta}_C^I/(1 - \bar{\zeta}_Y^I)$.

Appendix A Additional Figures and Tables


A.1 Figures

Figure A.1: Empirical Cumulative Distribution Function of Lottery Winnings




Notes: This figure shows the empirical cumulative distribution function of lottery winnings in our credit bureau sample. We include lottery prizes up to \$250,000 in our sample. The mean lottery win amount is \$8,363.23 (SD = \$24,808.19).

Figure A.2: Wealth Effects out of \$1 of Lottery Winnings by Credit Utilization (Credit Bureau Data)




Notes: This figure presents event-study estimates of Equation (1) for the effect of \$1 of lottery winnings on credit card spending and on total debt balance, by pre-win credit utilization. Panel (a) sums the quarterly effects on credit card spending over the two years before (pre-period) and after (post-period) the win, overall and by utilization group. Panel (b) averages the quarterly effects on saving via debt repayment over the same windows, overall and by utilization group; negative values indicate reductions in total debt balance and, therefore, reflect saving via debt repayment. We report the post-period sums (panel a) and post-period averages (panel b), with standard errors in parentheses. Utilization groups are estimated separately for winners below and above median credit utilization in the quarter preceding the win, where utilization is the ratio of revolving balances to credit limits. The full-sample estimates are constructed as equally weighted averages of the below- and above-utilization estimates. The sample comprises 11,703 winners from 2015Q1 to 2020Q1, observed for the eight quarters before and after the win. Credit card spending, total debt balance, and credit utilization are constructed from the credit bureau data as described in Section 2.2 and Appendix Section B. Whiskers are 95% confidence intervals from standard errors clustered at the individual level.

Figure A.3: Wealth Effects on Savings via Asset Accumulation from \$1 of Lottery Winnings (Tax Data)




Notes: This figure presents event-study estimates of Equation (1) for the effect of \$1 of lottery winnings on each component of saving via investment: the fair market value of Tax-Free Savings Accounts (panel a); interest and other investment income (panel b); taxable capital gains (panel c); and net contributions to Registered Retirement Savings Plans (panel d). Each panel plots the annual effect in the corresponding event year, normalized to zero in the year preceding the win. Components are shown as estimated, before the capitalization of income flows and the cumulation of net contributions described in Section 2.3, so panels (b) and (c) are annual income flows rather than asset stocks. The sample comprises 30,000 winners from 2004 to 2017. Figure 4 reports the aggregate saving via investment response. Shaded regions are 95% confidence intervals from standard errors clustered at the individual level.

Figure A.4: Labor Supply Responses from \$1 of Lottery Winnings (Tax Data)



Notes: This figure presents event-study estimates of Equation (1) for the effect of \$1 of lottery winnings on employment and on wage earnings conditional on employment. Panels (a) and (b) plot annual wealth effects, each normalized to zero in the year preceding the win. The reported pre- and post-period averages are simple averages of the annual coefficients over the pre-event years and over years 1 to 5 after the win, with standard errors in parentheses. Panel (c) averages the annual employment effects and panel (d) sums the conditional earnings effects, in each case over the pre-event years (pre-period) and over the two years following the win (post-period), overall and by income group. Post-period averages and sums include only fully treated years, and thus omit the year of the win. Income groups are estimated separately for winners below and above median pre-win income, and the full-sample estimates in panels (c) and (d) are equally weighted averages of the two. We use these estimates to decompose the unconditional two-year earnings response in Figure 5 into intensive- and extensive-margin components as described in Section 4.3. The sample comprises 30,000 winners from 2004 to 2017. Employment and wage earnings are described in Section 2.3. Shaded regions (panels a and b) and whiskers (panels c and d) are 95% confidence intervals from standard errors clustered at the individual level.

Figure A.5: Wealth Effect on Income Taxes out of \$1 of Lottery Winnings (Tax Data)



Notes: This figure presents event-study estimates of Equation (1) for the effect of \$1 of lottery winnings on income taxes. Panel (a) plots annual wealth effects on income taxes, normalized to zero in the year preceding the win. Panel (b) sums the effects on income taxes over the pre-event years (pre-period) and over the two years following the win (post-period), overall and by income group. We report the simple average of the annual coefficients within each period (panel a) and the pre- and post-period sums (panel b), with standard errors in parentheses. Post-period averages and sums include only fully treated years, and thus omit the year of the win. Income groups are estimated separately for winners below and above median pre-win income. The full-sample estimates in panel (b) are constructed as equally weighted averages of the below- and above-median estimates. The sample comprises 30,000 winners from 2004 to 2017. Income taxes are described in Section 2.3. Shaded regions (panel a) and whiskers (panel b) are 95% confidence intervals from standard errors clustered at the individual level.

A.2 Tables

Table A.1: Validating Credit Card Spending as a Consumption Proxy

	Implied consumption (Equifax)	Benchmark consumption (Statistics Canada)	Difference (%)
	(1)	(2)	(3)
Full sample	66,333	67,196	-1.3
Below median income	53,624	50,066	7.1
Above median income	79,042	84,327	-6.3

Notes: This table compares the level of household consumption in 2017 implied by our measure, constructed from credit card spending in our Equifax data, with a benchmark from Statistics Canada, overall and by income group. The Equifax data cover a 0.1% sample of all Canadians. We divide average credit card spending per person in each group by the group’s average credit card share of non-housing spending and then multiply by the group’s average equivalent household size, using equivalence scales as in Statistics Canada (2026). We then add housing consumption from Statistics Canada’s Distribution of Household Economic Accounts. Credit card spending and its mapping to overall consumption are described in Section 2.2 and Appendix Section B. The benchmark combines non-housing spending from Statistics Canada’s Survey of Household Spending (SHS) with housing consumption from the Distribution of Household Economic Accounts, as detailed in Appendix Section B. We report 2017 as it is near the midpoint of our credit bureau sample of 2015Q1 to 2020Q1. Housing consumption enters both measures identically, so the difference reflects the reconstruction of non-housing consumption alone. The “Full sample” is constructed as the equally weighted average of the below- and above-median groups, and the difference is expressed as a percentage of the benchmark. Because income is not observed in the credit bureau data, household income is proxied by the forward sortation area–level average from the closest preceding census (forward sortation areas are comparable in size to five-digit U.S. ZIP codes).

Table A.2: Predictive Characteristics of Lottery Win Amount (Credit Bureau Data)

Credit Characteristic	Coefficient (1)	Credit Characteristic	Coefficient (2)
Auto Loan Balance	−0.05 (0.08)	Installment Loan Utilization	−6.53 (19.95)
Auto Loan Count	77.94 (1729.69)	Mortgage Balance	−0.05 (0.01)
Auto Loan Credit Limit	0.03 (0.07)	Mortgage Count	−467.79 (1272.27)
Auto Loan Delinquency Count	−2659.80 (3707.41)	Mortgage Credit Limit	0.04 (0.01)
Auto Loan Utilization	16.60 (25.78)	Mortgage Delinquency Count	−8695.62 (8648.17)
Credit Card Amt. Past Due	1.46 (0.53)	Mortgage Utilization	0.00 (0.00)
Credit Card Balance	0.02 (0.06)	Moved in Last 2 Months	−1404.62 (1902.50)
Credit Card Balance (spend, quarter)	−0.23 (0.23)	Propensity: Bankruptcy	−1067.45 (2766.43)
Credit Card Balance (spend, quarter-end)	0.27 (0.23)	Propensity: Collection	−2855.53 (5940.15)
Credit Card Count	−154.82 (365.64)	Propensity: Consumer Proposal	−6812.56 (3923.86)
Credit Card Credit Limit	−0.03 (0.03)	Student Loan Balance	0.27 (0.30)
Credit Card Delinquency Count	−2208.36 (1571.16)	Student Loan Count	−2375.23 (2240.50)
Credit Card Spend Amount	−0.02 (0.06)	Student Loan Credit Limit	−0.12 (0.24)
Credit Card Utilization	11.97 (13.06)	Student Loan Delinquency Count	−2257.56 (5804.06)
Risk Score	6.55 (5.14)	Student Loan Utilization	−15.56 (34.37)
HELOC Balance	0.02 (0.01)	Telephone Loan Balance	0.42 (2.38)
HELOC Count	−196.33 (1791.43)	Telephone Loan Count	−841.61 (896.02)
HELOC Credit Limit	−0.00 (0.01)	Telephone Loan Credit Limit	1.62 (1.17)
HELOC Delinquency Count	50706.94 (14222.54)	Telephone Loan Delinquency Count	−2546.58 (2617.23)
HELOC Utilization	−22.71 (29.31)	Telephone Loan Utilization	46.81 (18.38)
Inquiry Count	−110.17 (371.19)	Unsecured LOC Balance	−0.02 (0.06)
Installment Loan Balance	−0.05 (0.06)	Unsecured LOC Count	958.27 (486.43)
Installment Loan Count	−728.80 (1186.98)	Unsecured LOC Credit Limit	−0.03 (0.04)
Installment Loan Credit Limit	0.05 (0.05)	Unsecured LOC Delinquency Count	−892.11 (2790.69)
Installment Loan Delinquency Count	−992.43 (3443.77)	Unsecured LOC Utilization	−6.78 (14.97)
R-squared	0.02	R-squared Adj.	0.01

Notes: This table reports estimates from a regression of the lottery win amount on all credit characteristics provided by the credit bureau, measured in the quarter prior to winning, conditional on win-cohort (quarter of win) fixed effects. Standard errors are reported in parentheses. The inability of credit characteristics to predict the win amount—an adjusted R-squared of 0.01, with 44 of 50 coefficients statistically insignificant at the 5% level—is consistent with random assignment of the lottery win amount. Variable definitions are provided in Section 2.2.

Table A.3: 2-Year Wealth Effects on Consumption out of Lottery Winnings, by Prize Restriction

Maximum prize included	Credit card spending	Overall consumption
	(1)	(2)
\$50,000	0.08	0.18
\$100,000	0.09	0.20
\$250,000	0.08	0.18
\$350,000	0.05	0.11
\$500,000	0.04	0.09
\$1,000,000	0.02	0.05
\$10,000,000	0.02	0.05
\$10,000,000 (TCMC > \$350k)	0.03	0.07
All winners (winsorized at \$350k)	0.10	0.23

Notes: This table reports the 2-year effect of \$1 of lottery winnings on credit card spending (column (1)) and overall consumption (column (2)) under alternative sample restrictions on the prize amount, based on estimates of Equation (1) on the full sample. “Maximum prize included” is the upper limit of prizes included in the sample; winners of prizes above the threshold are dropped. “TCMC” is top cell median coding: all prizes greater than \$350k are replaced with the median prize above \$350k. The \$250,000 threshold (in bold) corresponds to the baseline sample restriction used in Table 2; the consumption estimate differs slightly because that table reports the equally weighted average of the below- and above-median income estimates, whereas this table pools income groups. Credit card spending is measured in the credit bureau data; its construction and mapping to overall consumption are described in Section 2.2 and Appendix Section B.

Table A.4: Decomposition of the Two-Year Wealth Effect on Taxes

	Full Sample (1)	Income	
		Below Median (2)	Above Median (3)
Pre-tax wage earnings	−0.088	−0.040	−0.137
Total income taxes	−0.025	−0.008	−0.041
Labor component	−0.026	−0.009	−0.042
Capital component	0.001	0.001	0.001
Implied marginal tax rate on ordinary income	0.291	0.229	0.308

Notes: This table decomposes the two-year wealth effect on total income taxes into the part due to lower labor income and the part due to higher taxable capital income. All responses are per \$1 of winnings, aggregated over event years 1 and 2, the window used for the wealth effects on earnings in Figure 5(b) and on financial assets in Figure 4(b). Lottery prizes are untaxed in Canada, so the tax response operates entirely through winners’ taxable income. We measure the response of taxable income as the sum of the responses of wage earnings, taxable capital gains, and interest and other investment income, less net RRSP contributions, which are tax deductible. TFSA balances are excluded because they generate no taxable income. A single marginal rate τ_L applies to every component because source-specific tax treatment in Canada operates through the tax base rather than the rate (e.g., only the taxable half of capital gains is reported in the tax data). We back out τ_L as the ratio of the tax response to the taxable-income response, so the labor component is τ_L times the earnings response and the capital component is the remainder. The capital component is \$9 to \$11 of tax per \$10,000 won in every group, offsetting about 4 percent of the labor component, so the wealth effect on taxes is, to a close approximation, a labor income tax response. As a check on the accounting, the implied rates are in line with combined federal and provincial statutory rates in the winners’ province over the sample period. Event-study estimates for each component of taxable capital income are shown in Figure A.3.

Appendix B Credit Card Spending as Consumption

B.1 Correcting for Accrued Interest

Equifax calculates our measure of quarterly q credit card spending as the sum of individual-level i monthly m credit card spending:³⁶

$$\text{Credit card spend}_{iq}^{eqx} = \sum_{m \in q} \max(0, \text{Balance}_{im} - \text{Balance}_{i,m-1} + \text{Payment}_{im}). \quad (11)$$

Eligibility for the monthly spending calculation is determined by card-level data integrity thresholds.³⁷ For “transactors,” who pay their statement balance in full and therefore do not accrue interest, Equation (11) recovers credit card purchases. For individuals who “revolve” balances, accrued interest also increases the outstanding balance, so the same calculation can capture both actual purchases and interest charges. Our reduced-form results show that lottery winners, particularly those with below-median income, use some of their winnings to pay down credit card debt. As balances fall, accrued interest also falls. Because accrued interest is included in the unadjusted Equifax spending measure, this decline can offset part of the increase in actual purchases and attenuate the estimated spending response.³⁸ We therefore identify revolvers, estimate their accrued interest, and subtract it from the Equifax spending measure.

B.1.1 Identifying Revolvers

To identify which individuals revolve balances on their credit cards, we first recover quarterly payments from the balance and spending information observed in the credit bureau data. For all individuals i and quarters q , we observe Balance_{iq} , $\text{Balance}_{i,q-1}$, and $\text{Credit card spend}_{iq}^{eqx}$. For

³⁶We switch from using time subscript t in this appendix to clearly distinguish between months and quarters.

³⁷Specifically, a card is deemed ineligible for a spending calculation if any of the following criteria are met: (i) $\text{Balance}_{m-1} > 0$ but $\text{Payment}_m = 0$ and $\text{Amount past due}_m = 0$, as this indicates a reporting error; (ii) the previous monthly report for the card was not within 27–35 days of the current report; or (iii) any of the required data fields for the calculation are missing or the card had “closed” status in the prior monthly report.

³⁸We are grateful to Stephen Zeldes for identifying this phenomenon.

Credit card spend $_{iq}^{eqx} > 0$, we can rearrange Equation (11) to recover quarterly payments:

$$\text{Payment}_{iq} = \text{Credit card spend}_{iq}^{eqx} - \text{Balance}_{iq} + \text{Balance}_{i,q-1}. \quad (12)$$

An individual is a transactor if they pay off their entire previous statement balance, so that $\text{Payment}_{iq} \geq \text{Balance}_{i,q-1}$. Using Equation (12), this condition is equivalent to

$$\text{Payment}_{iq} \geq \text{Balance}_{i,q-1} \iff \text{Credit card spend}_{iq}^{eqx} \geq \text{Balance}_{iq}.$$

Therefore, an individual is classified as a revolver if $\text{Credit card spend}_{iq}^{eqx} < \text{Balance}_{iq}$.

B.1.2 Recovering Actual Credit Card Spending for Revolvers

For revolvers, the Equifax spending measure includes interest charges:

$$\text{Credit card spend}_{iq}^{eqx} = \text{Credit card spend}_{iq}^{actual} + \text{Interest}_{iq}. \quad (13)$$

We make three assumptions to recover $\text{Credit card spend}_{iq}^{actual}$. First, because we observe only quarter-end balances, we treat each quarter as a single billing cycle and assume that transactor/revolver status is stable within a quarter. This assumption is supported by the persistence of revolving: Xiao (2024) documents that a large share of Canadian credit card holders carry balances across consecutive months.

Second, we assume an annual percentage rate (APR) of 20.00% on all credit cards. The standard APR in Canada has long been approximately 19.99% (Henry, Huynh, and Welte, 2018; BDO Debt Solutions, 2025; Allen, Boutros, and Guttman-Kenney, 2026) and credit card APRs are relatively stable over time and similar across the household income distribution. The latter is consistent with lenders managing credit risk primarily through credit limits rather than interest rates, as Agarwal et al. (2018) document for U.S. lenders across the credit score distribution.

Third, we assume that transactions and payments are distributed uniformly throughout the

billing cycle.³⁹ This approximation allows us to construct the average daily balance (ADB).

With these assumptions, accrued interest is given by:

$$\text{Interest}_{iq} = r_q \times \text{ADB}_{iq},$$

where r_q is the average quarterly interest rate and:

$$\text{ADB}_{iq} = \text{Balance}_{i,q-1} + \frac{\text{Credit card spend}_{iq}^{\text{actual}} - \text{Payment}_{iq}}{2}.$$

Combining Equations (12) and (13) yields

$$\text{Credit card spend}_{iq}^{\text{actual}} - \text{Payment}_{iq} = (\text{Balance}_{iq} - \text{Balance}_{i,q-1}) - \text{Interest}_{iq}.$$

ADB_{iq} therefore collapses to the midpoint of the beginning and end of quarter balances, net of accrued interest:

$$\begin{aligned} \text{ADB}_{iq} &= \text{Balance}_{i,q-1} + \frac{(\text{Balance}_{iq} - \text{Balance}_{i,q-1}) - \text{Interest}_{iq}}{2} \\ &= \frac{\text{Balance}_{i,q-1} + \text{Balance}_{iq}}{2} - \frac{\text{Interest}_{iq}}{2}. \end{aligned}$$

Substituting this expression into $\text{Interest}_{iq} = r_q \times \text{ADB}_{iq}$ and solving for Interest_{iq} gives

$$\begin{aligned} \text{Interest}_{iq} &= r_q \left(\frac{\text{Balance}_{i,q-1} + \text{Balance}_{iq}}{2} - \frac{\text{Interest}_{iq}}{2} \right) \\ &= \frac{r_q}{1 + r_q/2} \left(\frac{\text{Balance}_{i,q-1} + \text{Balance}_{iq}}{2} \right). \end{aligned}$$

This expression depends only on the two observed quarter-end balances and can be combined with Equation (13) to recover $\text{Credit card spend}_{iq}^{\text{actual}}$ for revolvers.

³⁹Cardholders receive an interest-free grace period to pay statement balances between the closing date and the payment due date. If the statement balance is not paid in full by the payment date, interest is charged on the average daily balance retroactively from the date of each underlying transaction.

B.1.3 Validation and Empirical Relevance

To validate our procedure, we compare our classification with an external benchmark for the Canadian population of credit card holders. In the quarter before winning the lottery, we classify 42% of credit card-holding winners as revolvers. This is close to the 46% revolving share documented for Canadian credit card holders by Xiao (2024). Also consistent with Xiao (2024), the revolving share declines with household income.

Having validated our procedure, we assess how much it matters for our wealth effects on credit card spending.⁴⁰ If lottery winners use some of their winnings to repay credit-card debt, accrued interest will fall, attenuating our estimates. Consistent with this mechanism, correcting for accrued interest raises the two-year credit card spending effect by 9% overall and by 13% among below-median-income winners, who are both more likely to revolve and more likely to repay credit-card debt after winning.

B.2 From Credit Card Spending to Consumption

While we aim to estimate the wealth effect of lottery winnings on consumption, we only directly observe credit card spending. To recover the wealth effect on consumption, we rely on scaling factors that are equal to the inverse of the share of total consumption that occurs on credit cards. This section details how we construct these scaling factors.

B.2.1 Data

We combine three data sources to construct our scaling factors. We first use estimates from the Bank of Canada’s Methods-of-Payment Survey reports (Arango and Welte, 2012; Henry, Huynh, and Shen, 2015; Henry, Huynh, and Welte, 2018; Henry, Shimoda, and Zhu, 2022) to measure the share of non-housing consumption paid by credit cards.⁴¹ We pair this with annual data on non-

⁴⁰Our procedure provides an approximation but is imperfect, since we observe only quarter-end balances summed across cards and treat the sum of monthly spending measures as exact. In practice, an individual may revolve on some cards but not others, or in some months of a quarter but not all, and our classification cannot distinguish these cases.

⁴¹Over our sample period, the Bank of Canada released these reports in 2012, 2015, 2018, and 2022; they vary in income granularity and sometimes report estimates for earlier years. Our income gradient in credit card use is taken

housing expenditure from Statistics Canada’s Survey of Household Spending (SHS) and housing consumption from the Distribution of Household Economic Accounts (DHEA).⁴²

B.2.2 Constructing the Credit Card Share of Total Consumption

The credit card share of total consumption is the product of two conditional shares. Let y denote income and t denote year, $f_t^{cc}(y)$ the share of non-housing consumption paid by credit card, and $f_t^h(y)$ the share of total consumption that goes to housing. The credit card share of total consumption at income y in year t is:

$$s_t(y) = f_t^{cc}(y) (1 - f_t^h(y)).$$

We construct $s_t(y)$ from spending moments observed in different parts of the income distribution and in different years, interpolating between them to recover the shares we need. The procedure has three steps: we first calibrate the national income distribution; we then build f_t^{cc} and f_t^h as smooth functions of income in each year; and we finally aggregate the resulting annual curves over the income distribution and the timing of lottery winners in our sample.

We begin by calibrating the national income distribution. We fit a log-normal distribution to the 2017 national income quintiles and splice on a Pareto tail above the 90th percentile, with a tail parameter estimated from our tax microdata. We hold this distribution fixed across years. It serves two purposes: it locates where in the income distribution each survey moment sits, and it supplies the weights under which we average the f_t^{cc} gradient to its national level. Next, we construct f_t^{cc} and f_t^h across income and over time.⁴³ Within each year, we interpolate the survey

from the 2009 wave, which reported the most granular income categories.

⁴²Our event-study estimates are at the individual level, whereas the SHS and DHEA are reported at the household level. Because the scaling factor depends on spending shares rather than levels, we assume that these shares do not differ between per-person and per-household measures. Table A.1 supports this: mapping per-person credit card spending through the scaling factors and adjusting for household size with equivalence scales reproduces household consumption in the Statistics Canada data closely, both overall and by income group.

⁴³As is detailed in Section 2.2.1, the credit card share of non-housing spending rose substantially from 2009 to 2020, reflecting wider card ownership and more intensive use among cardholders. Three developments help explain this shift. First, the rollout of contactless (tap) payment led consumers to use credit cards for smaller transactions that were previously paid in cash. Second, e-commerce, for which credit cards are the primary payment method, grew

moments across income with monotone piecewise-cubic (PCHIP) splines, which pass through every observed moment while preserving the monotonicity of the underlying gradient. Outside the range spanned by the survey bins, we extend each curve with a penalized extrapolation that flattens its slope.⁴⁴ The two functions differ only in how they move over time. We observe the credit card gradient by income in 2009, so we hold its shape fixed and rescale its level each year to match the national credit card share, itself interpolated between the survey years with a monotone spline. Housing shares, by contrast, are observed by income in every year, so we rebuild f_t^h directly from each year's quintile shares.⁴⁵

Taken together, these functions yield the credit card share of total consumption $s_t(y)$ at each income level y and in a given year t . We first average the annual curves across the sample period, weighting years equally because winners are approximately uniformly distributed over time; this gives the single curve plotted in Figure B.1.⁴⁶ We then average this curve within each income group, weighting by the income distribution of our lottery sample, which we calibrate in the same way as the national distribution. The scaling factor for each group is the inverse of the resulting share, and we construct the overall full-sample share as their equal-weighted average.

B.3 Constant Housing Share Assumption

As discussed in Section 2.2.2, our scaling procedure assumes that lottery winnings do not change the share of total consumption devoted to housing. This section details the exercise, discussed in Section 4.1, that assesses this assumption by comparing the housing-consumption response implied by our scaling procedure with an independent estimate based on mortgage originations.


substantially, particularly from 2017 onward. Third, rewards programs increased credit card use (see, e.g., Henry, Shimoda, and Zhu, 2022, among many others).

⁴⁴We place each spline knot at the mean income within its data bin, computed under the calibrated national income distribution.

⁴⁵Since the SHS was not fielded in 2018, we linearly interpolate the missing 2018 non-housing quintile values from the 2017 and 2019 waves.

⁴⁶Specifically, winners in our credit bureau sample receive their prizes between 2015Q1 and 2020Q1; we construct the curve for each complete calendar year in that window, 2015 through 2019.

Figure B.1: Credit Card Share of Total Consumption



Notes: This figure shows the credit card share of total consumption, $s(y)$, averaged over 2015–2019 (solid line). The shaded region is the income density of our lottery sample, shown on an unlabeled secondary axis, which serves as the weight when aggregating within-year credit card shares of total consumption $s_t(y)$. The dashed segments are the resulting group-average shares; their inverses are the scaling factors applied to credit card spending. The vertical line marks the sample median income. Construction is described in Section B.2.2.

B.3.1 Data

We supplement our data above with four inputs. First, the SHS reports utility expenditure, including water, fuel, and electricity, by household income quintile. Second, we use the Equifax microdata to measure mortgage ownership by income quartile. We interpolate both utility expenditure and mortgage ownership across income using the same PCHIP procedure described in Section B.2.2. Third, we use the posted five-year fixed mortgage rate from the Canadian Mortgage and Housing Corporation (2018), which averages 4.9% over our sample period. Finally, we construct mortgage originations from the Equifax microdata. Following Di Maggio et al. (2017), we identify an origination when an individual’s outstanding mortgage balance increases and define the origination amount as the size of that increase.

B.3.2 Implied Housing Consumption Response from Scaling Factors

Our scaling factors imply that a winner who increases total consumption by \$1 raises housing consumption by the population-average housing share of consumption. Housing consumption is measured in the DHEA and comprises rent, imputed for homeowners, together with utilities. Because the direct measure we construct below is most comparable to rent, we use the SHS utility expenditure described above to net utilities out of the housing share, under the additional assumption that winnings leave the rental share of housing consumption unchanged. The implied effect on rent is then this net housing share multiplied by the 2-year wealth effect on overall consumption.

B.3.3 Estimated Housing Consumption Response from Mortgage Originations

We estimate the effect of \$1 of lottery winnings on mortgage payments, which we interpret as a proxy for rent. To do so, we estimate the effect of lottery winnings on mortgage origination balances using Equation (1) (Table B.1).⁴⁷ We then convert the increase in mortgage balances into mortgage payments using a 25-year amortization period and the 4.9% average mortgage interest rate. Under the assumption that mortgage-holding homeowners increase housing consumption to the same extent as renters and outright homeowners, we scale this payment response by the inverse of the mortgage ownership rate in the Equifax data. Because mortgage payments generated by an origination recur in subsequent quarters, we accumulate the annuitized payment effects throughout the two-year period. Figure B.2(a) presents the results of this exercise, which are discussed in Section 4.1. Since this exercise abstracts from the cash portion of housing consumption, the resulting estimate should be interpreted as a lower bound.


⁴⁷We focus on originated mortgages to disentangle consumption from saving.

Table B.1: Wealth Effects out of Lottery Winnings on Loan Origination Balances

	Pre-period ($t = -8$ to -1)	Post-period ($t = 0$ to $+8$)
	(1)	(2)
Total origination balances	0.032 (0.017)	0.075 (0.022)
<i>Mortgages</i>	0.034 (0.016)	0.075 (0.021)
<i>Auto loans</i>	−0.006 (0.002)	−0.006 (0.003)
<i>Installment loans</i>	0.003 (0.004)	0.006 (0.005)

Notes: This table presents event-study estimates of Equation (1) for the effect of \$1 of lottery winnings on loan origination balances, defined as increases in outstanding non-revolving loan balances. Each entry sums the quarterly effects over the two years before the win (column (1)) and over the win quarter and the two years following it (column (2)). Total origination balances comprise mortgages, auto loans, and installment loans. The sample comprises 11,703 winners from 2015Q1 to 2020Q1. Origination balances are constructed from the credit bureau data as described in Section 2.2 and Appendix Section B.3. Standard errors, clustered at the individual level, are in parentheses.

Figure B.2: Housing Consumption and Expenditure Responses out of \$1 of Lottery Winnings (Credit Bureau Data)



Notes: This figure compares expenditure and consumption responses to \$1 of lottery winnings, for housing (panel a) and overall (panel b). Consumption measures the flow of services from a durable good over time, whereas expenditure measures outlays at the time of purchase. Responses implied by our scaling procedure apply the equiproportionality assumption to our estimated overall consumption response; responses estimated from mortgage originations are constructed from the effect on origination balances reported in Table B.1. Details of this exercise are discussed in Section 2.2 and Appendix Section B.3. Whiskers are 95% confidence intervals from standard errors clustered at the individual level.

Appendix C Annuitization Approaches

C.1 Level Annuitization Approach

This appendix describes the annuitization approach used to construct the MPC, MPE and MPT estimates in Figure 9. Consider a k -year-old winner who receives a lottery prize L in year $t = 0$ and has $T - k$ years of life remaining after the win year. The unspent portion of the lottery prize earns an interest rate r . The annuitization approach assumes the winner smooths the prize perfectly, implying that the lump sum is treated as an annuity paid in the year of the win and in each of the subsequent years. Under this assumption, unearned income must increase by the following amount in each period:

$$\frac{r}{1+r} \left[1 - \left(\frac{1}{1+r} \right)^{T-k+1} \right]^{-1} L. \quad (14)$$

To empirically implement this approach, we set the interest rate to $r = 3\%$, approximately the average yield on long-term Government of Canada bonds over our sample period, and use a life expectancy of 80 years, approximately the Canadian life expectancy over the sample period. Evaluating Equation (14) at the mean age at win of approximately $k = 50$ and the remaining horizon $T - k = 30$, each dollar of the lottery prize translates to about 4.9 cents of additional unearned income in each remaining year, or roughly \$485 per year per \$10,000 won.

We use this translation, along with our wealth effects on pre-tax labor earnings (Figure 5) and on labor income taxes (Figure A.5), to construct the annuitized MPE and MPT. The MPC is recovered as a residual using the budget constraint:

$$\text{MPC} = 1 + \text{MPE} - \text{MPT}. \quad (15)$$

The MPC, MPE, and MPT estimated using this approach are presented in Figure 9.

C.2 Life-Cycle Model with Stone-Geary Preferences

Our two-sample IV estimates presented in Figure 8 and Table 3 reveal that the measured annual flow of unearned income over the five years post-win is around 10 cents per dollar of prize, which exceeds the 4.9 cents implied by a level annuity as shown in Appendix Section C.1. This section shows that one can relax the perfect smoothing assumption to rationalize the gap in these estimates.

Assume that individual preferences are time-separable with subjective discount rate ϑ . Per-period utility is given by a Stone-Geary function over consumption and leisure ℓ_t :

$$\sum_{t \geq 0} (1 + \vartheta)^{-t} U(c_t, \ell_t), \quad U(c_t, \ell_t) = \gamma \ln(c_t - \underline{c}) + (1 - \gamma) \ln(\ell_t - \underline{\ell}), \quad (16)$$

where $\underline{c}$ and $\underline{\ell}$ are subsistence levels. Earnings are $y_t = w_t(\bar{H} - \ell_t)$ for a time endowment $\bar{H}$, and the winner faces the lifetime budget constraint

$$\sum_{t=0}^{30} \frac{c_t}{(1+r)^t} = a_0 + L + \sum_{t=0}^{30} \frac{y_t - \mathcal{T}(y_t)}{(1+r)^t}, \quad (17)$$

where a_0 denotes initial assets. It can be shown that a k -year-old household with a remaining lifetime of $T - k$ periods allocates a share of the average flow of resources from the lottery prize over the first five post-win years as:

$$\lambda(r, \vartheta) \equiv \frac{1}{5} \sum_{t=1}^5 \left(\frac{1+r}{1+\vartheta} \right)^t \frac{\vartheta}{1+\vartheta} \left[1 - \left(\frac{1}{1+\vartheta} \right)^{T-k+1} \right]^{-1}. \quad (18)$$

We calibrate ϑ so that $\lambda(r, \vartheta)$ matches the measured average flow of unearned income over the first five years following the win (10.48 cents per dollar of prize). This yields a discount rate of $\vartheta = 18.4\%$. Thus, winners draw down the prize about twice as fast as implied by a level annuity, a pattern that can be rationalized by a subjective discount rate substantially above the interest rate. For high-income individuals, the drawdown estimate gives $\vartheta = 15.1\%$, while for low-income individuals, we get $\vartheta = 22.8\%$.

C.3 Savings Response to a Recurring Transfer

This appendix derives the initial savings response to a recurring transfer as a function of the remaining horizon and evaluates it at the discount rates calibrated in Appendix Section C.2.

Suppose that an individual of age k receives a transfer B in each remaining period through age T , so that the transfer arrives in periods $t = 0, \dots, T - k$. Let a_t denote assets entering period t and define $N_t \equiv c_t - (y_t - T_t)$ as consumption net of after-tax earnings. The period budget constraint is then $a_{t+1} = (1 + r)a_t + B - N_t$. Initial-period saving is $S \equiv a_1$. Because initial assets are predetermined when the transfer is introduced, $\partial a_0 / \partial B = 0$.

With perfect capital markets, the transfer affects consumption and earnings only through lifetime wealth W . A one-dollar increase in B increases lifetime wealth by $\partial W / \partial B = \sum_{s=0}^{T-k} (1 + r)^{-s}$. Under Stone-Geary preferences, the marginal resource flow, $\partial N_t / \partial W$, declines geometrically at rate $(1 + r) / (1 + \vartheta)$, where ϑ is the subjective discount rate. The lifetime budget constraint requires that the marginal dollar be exhausted, $\sum_{t=0}^{T-k} (1 + r)^{-t} \partial N_t / \partial W = 1$, which pins down the initial flow as $\partial N_0 / \partial W = [\sum_{s=0}^{T-k} (1 + \vartheta)^{-s}]^{-1}$. Since $\partial N_0 / \partial B = (\partial N_0 / \partial W)(\partial W / \partial B)$, the initial savings response is


$$\frac{\partial S}{\partial B} = 1 - \frac{\partial N_0}{\partial B} = 1 - \frac{\sum_{s=0}^{T-k} (1 + r)^{-s}}{\sum_{s=0}^{T-k} (1 + \vartheta)^{-s}}. \quad (19)$$

The ratio in Equation (19) equals one when $\vartheta = r$, implying $\partial S / \partial B = 0$. When $\vartheta > r$, as implied by the calibrated drawdown in Appendix Section C.2, the ratio exceeds one and $\partial S / \partial B < 0$. In this case, the individual uses more than one dollar of resources in the initial period for each dollar of transfer, financing the difference by reducing saving or borrowing against future transfers.

Figure C.1 evaluates Equation (19) at $r = 3$ percent for remaining horizons of up to 30 years, using the subjective discount rates calibrated for the full population and for low- and high-income individuals. The savings response is zero when $T - k = 0$ and becomes increasingly negative

as the remaining horizon lengthens. At $T - k = 30$, a one-dollar recurring transfer reduces initial savings by \$2.22 in the full sample, \$1.74 among above-median income winners, and \$2.83 among below-median income winners.

Figure C.1: Savings Response to a Recurring Transfer, by Remaining Horizon



Notes: This figure plots the savings response to a one-dollar recurring transfer in the period it is introduced, $\partial S/\partial B$ in Equation (19), against the remaining horizon $T - k$ at an interest rate of $r = 3$ percent. Each series evaluates Equation (19) at the subjective discount rate ϑ calibrated in Appendix Section C.2 to match the group's five-year average first stage: 0.1048 for the full sample ($\vartheta = 18.4$ percent), 0.0963 for above-median income winners ($\vartheta = 15.1$ percent), and 0.1132 for below-median income winners ($\vartheta = 22.8$ percent). The dashed line is the level-annuity benchmark $\vartheta = r$ at which $\partial S/\partial B = 0$ at every horizon. The labels report values at $T - k = 30$, the horizon used in Appendix Section C.2.

Appendix D UBI Derivations

D.1 Marginal Surcharge Derivation

This appendix derives the surcharge formula, Equation (5) of Section 6.2. Differentiating Equation (4) with respect to B :

$$\frac{d\mathcal{R}}{dB} = -1 + \frac{\partial \delta}{\partial B} \mathbb{E}[Y] + (\tau_Y + \delta(B)) \mathbb{E} \left[\frac{dY}{dB} \right] + \tau_S \mathbb{E} \left[\frac{dS}{dB} \right] = 0, \quad (20)$$

where the total derivatives account for the fact that the transfer changes earnings and savings both directly through the unearned income and indirectly through the change in the surcharge.

Next, define the uncompensated elasticities of earnings and savings with respect to the net-of-tax rate, respectively, as:

$$\zeta_{Y, \tau_Y}^U = \frac{\partial \ln Y}{\partial \ln(1 - \tau_Y - \delta)}, \quad \zeta_{S, \tau_Y}^U = \frac{\partial \ln S}{\partial \ln(1 - \tau_Y - \delta)},$$

Earnings change through two channels when the transfer increases. The direct effect is the response to the additional unearned income, holding tax rates fixed. The indirect effect is the response to the higher tax rate needed to fund the transfer:

$$\frac{dY}{dB} = \frac{\partial Y}{\partial B} - \frac{Y \zeta_{Y, \tau_Y}^U}{1 - \tau_Y} \frac{\partial \delta}{\partial B}, \quad (21)$$

$$\frac{dS}{dB} = \frac{\partial S}{\partial B} - \frac{S \zeta_{S, \tau_Y}^U}{1 - \tau_Y} \frac{\partial \delta}{\partial B}. \quad (22)$$

Taking expectations of Equation (21) and Equation (22) across individuals, the indirect effects are the mean responses of earnings and savings to the net-of-tax rate, $\mathbb{E}[Y \zeta_{Y, \tau_Y}^U]$ and $\mathbb{E}[S \zeta_{S, \tau_Y}^U]$, because $\partial \delta / \partial B$ and τ_Y are common to all individuals. Substituting into Equation (20) and solving

for $\partial\delta/\partial B$ yields Equation (5):

$$\frac{\partial\delta}{\partial B} = \frac{1 - (\tau_Y + \delta(B))\mathbb{E}\left[\frac{\partial Y}{\partial B}\right] - \tau_S\mathbb{E}\left[\frac{\partial S}{\partial B}\right]}{\mathbb{E}[Y] - \frac{\tau_Y + \delta(B)}{1 - \tau_Y}\mathbb{E}\left[Y\zeta_{Y,\tau_Y}^U\right] - \frac{\tau_S}{1 - \tau_Y}\mathbb{E}\left[S\zeta_{S,\tau_Y}^U\right]}.$$

D.2 Calibrating the Savings Response to the Net-of-Tax Rate

This appendix derives the savings response to the net-of-tax rate $\mathbb{E}[S\zeta_{S,\tau_Y}^U]$ in Equation (5). The Slutsky equation for earnings is:

$$\zeta_{Y,\tau_Y}^U = \zeta_{Y,\tau_Y}^H - \zeta_Y^I, \quad \zeta_Y^I = (1 - \tau_Y) \left| \frac{\partial Y}{\partial B} \right|, \quad (23)$$

We calibrate the compensated earnings elasticity using estimates from Gruber and Saez (2002) and the income effect using our lottery estimates. The Slutsky equation for savings can be written as:

$$S\zeta_{S,\tau_Y}^U = S\zeta_{S,\tau_Y}^H + (1 - \tau_Y)Y \frac{\partial S}{\partial B}, \quad (24)$$

The budget constraint can be used to help pin down the compensated savings response as follows:

$$(1 - \tau_Y)Y \zeta_{Y,\tau_Y}^H = C\zeta_{C,\tau_Y}^H + (1 + \tau_S)S\zeta_{S,\tau_Y}^H, \quad (25)$$

Under separability, Hellwig and Werquin (2026) show that:

$$\frac{\zeta_{C,\tau_Y}^H}{\zeta_{Y,\tau_Y}^H} = \frac{\zeta_C^I}{1 - \zeta_Y^I}$$

Rearranging this condition gives $\zeta_{C,\tau_Y}^H = \zeta_{Y,\tau_Y}^H \frac{\zeta_C^I}{1 - \zeta_Y^I}$, where $\zeta_C^I = (1 - \tau_Y)Y \frac{\partial \ln C}{\partial B}$. Thus,

$$C\zeta_{C,\tau_Y}^H = (1 - \tau_Y)Y \zeta_{Y,\tau_Y}^H \frac{\partial C/\partial B}{1 + (1 - \tau_Y)\partial Y/\partial B}. \quad (26)$$

Substituting Equation (26) into Equation (25) yields the compensated savings response:

$$S\zeta_{S,\tau_Y}^H = \frac{1-\tau_Y}{1+\tau_S} Y \zeta_{Y,\tau_Y}^H \left(1 - \frac{\partial C / \partial B}{1 + (1-\tau_Y) \partial Y / \partial B} \right). \quad (27)$$

Appendix E Estimating the Pareto Coefficients of Earnings and Consumption

The optimal top tax formulas of Hellwig and Werquin (2026) depend on the ratio of Pareto coefficients of the earnings and consumption distributions, $\bar{\rho}_Y/\bar{\rho}_C$. This ratio has a simple interpretation: if consumption scales with income as $C \propto Y^\eta$ at the top, then $\bar{\rho}_C = \bar{\rho}_Y/\eta$, so the ratio is the cross-sectional elasticity of consumption with respect to income among top earners. A ratio of one means consumption rises proportionally with income at the top, so the saving rate is constant across top earners. A ratio below one means consumption rises less than proportionally, so the saving rate is increasing in income at the top.

This appendix describes how we estimate the Pareto coefficients using the estimator of Hill (1975). Let $x_{(1)} \geq x_{(2)} \geq \dots \geq x_{(k)}$ denote the k largest observations of a variable x in a sample of size n . The Hill estimator of the Pareto coefficient is given by:

$$\hat{\rho} = \left[\frac{1}{k} \sum_{i=1}^k \ln \frac{x_{(i)}}{x_{(k)}} \right]^{-1}, \quad (28)$$

where $x_{(k)}$ is the threshold that defines the top of the distribution, with asymptotic standard error $\hat{\rho}/\sqrt{k}$. We set k to the top 5% of observations for both series.

For consumption, we use a 0.1% random sample of individuals in the credit bureau data in 2017 and measure C as annual credit card spending. We use raw card spending rather than a consumption measure scaled to national benchmarks, because we lack scaling factors specific to the top of the distribution. A uniform scaling factor would leave the Hill estimate unchanged in any case, since Equation (28) depends only on ratios; rescaling card spending by age, an observable correlated with the share of consumption paid by card, also leaves the estimate essentially unchanged.

For income, we use all Canadian tax filers with positive total income from the tax records in 2017, the same year as the credit bureau sample, and measure Y as total income before tax. The resulting estimates are $\bar{\rho}_Y = 2.30$ and $\bar{\rho}_C = 3.50$, implying $\bar{\rho}_Y/\bar{\rho}_C = 0.66$, so consumption is less

concentrated than income at the top of the distribution. For comparison, Hellwig and Werquin (2026) estimate $\bar{\rho}_Y = 1.8$ and $\bar{\rho}_C = 3.1$ in the Panel Study of Income Dynamics (PSID), implying a ratio of 0.58, while the ratio they use to calibrate their formulas, obtained by regressing average consumption on income across top quantiles, is 0.63. Our estimate of the ratio is close to theirs, although both Pareto coefficients are larger, consistent with thinner upper tails in Canada than in the United States.

Appendix References

- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebe (2018). “Do Banks Pass through Credit Expansions to Consumers Who Want to Borrow?” *The Quarterly Journal of Economics* 133(1), 129–190. <https://doi.org/10.1093/qje/qjx027> (visited on 07/31/2026).
- Allen, Jason, Michael Boutros, and Benedict Guttman-Kenney (2026). “Paying Off: Evaluating Credit Card Minimum Payment Requirements.” Working Paper. <https://bguttmankenney.github.io/Public/quebec.pdf>.
- Arango, Carlos and Angelika Welte (2012). “The Bank of Canada’s 2009 Methods-of-Payment Survey: Methodology and Key Results.”
- BDO Debt Solutions (2025). “Why Credit Card Interest Rates Are Not Going Down.” <https://debtsolutions.bdo.ca/why-credit-card-interest-rates-are-not-going-down/> (visited on 07/31/2026).
- Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao (2017). “Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging.” *American Economic Review* 107(11), 3550–3588. <https://www.aeaweb.org/articles?id=10.1257/aer.20141313> (visited on 07/31/2026).
- Government of Canada, Statistics Canada (2018). “Canada Mortgage and Housing Corporation, Conventional Mortgage Lending Rate, 5-Year Term.” <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3410014501> (visited on 08/17/2026).
- Gruber, Jon and Emmanuel Saez (2002). “The Elasticity of Taxable Income: Evidence and Implications.” *Journal of Public Economics* 84(1), 1–32. <https://www.sciencedirect.com/science/article/pii/S0047272701000858> (visited on 03/30/2026).
- Hellwig, Christian and Nicolas Werquin (2026). “Using Consumption Data to Derive Top Optimal Income and Capital Tax Rates.” *Journal of Political Economy*.
- Henry, Christopher, Kim Huynh, and Rallye Shen (2015). “2013 Methods-of-Payment Survey Results.” Bank of Canada. <https://www.oar-rao.bank-banque-canada.ca/record/6360> (visited on 09/11/2026).
- Henry, Christopher, Matthew Shimoda, and Julia Zhu (2022). “2021 Methods-of-Payment Survey Report.” Staff Discussion Paper 2022–23. Bank of Canada. <https://www.bankofcanada.ca/2022/12/staff-discussion-paper-2022-23/>.
- Henry, Christopher S., Kim Huynh, and Angelika Welte (2018). “2017 Methods-of-Payment Survey Report.” <https://www.bankofcanada.ca/2018/12/staff-discussion-paper-2018-17/> (visited on 08/13/2025).
- Hill, Bruce M. (1975). “A Simple General Approach to Inference About the Tail of a Distribution.” *The Annals of Statistics* 3(5), 1163–1174. <https://www.jstor.org/stable/2958370> (visited on 09/11/2026).
- Statistics Canada (2026). *Distributions of Household Economic Accounts, Income, Consumption and Saving, by Characteristic*. Government of Canada. <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3610058701> (visited on 09/11/2026).
- Xiao, Jia Qi (2024). “The Reliance of Canadians on Credit Card Debt as a Predictor of Financial Stress.” <https://www.bankofcanada.ca/2024/07/staff-analytical-note-2024-18/> (visited on 07/31/2026).