New Estimates of Wealth Inequality in Canada^{*}

Alexander Hempel University of Toronto

February 20, 2025

Click Here for Latest Version

Abstract

Measures of wealth inequality are important indicators, but only exist in a handful of countries. This paper is the first to estimate the distribution of wealth in Canada on a regular basis from 1990-2018. Using the income capitalization method of Saez & Zucman (2016), I find that while the top 1% wealth share rose from 15.3% in 1990 to 19.7% in 2008, it fell back to 17.5% by 2018. These results suggest that Canada has much less wealth inequality compared to the US and is even slightly more equal than France. Using linear decomposition methods, I show that this gap with the US is driven by greater concentration across every asset class and is not driven by a single asset or a different composition of assets held in each country. Then, using *synthetic savings* decompositions, I show that most of the variation in the top 1% wealth share can be explained by the collapse in the top 1%'s savings rate, which, while positive from an inequality perspective, could have important ramifications for future economic growth in Canada.

^{*}I would like to thank Michael Smart for his guidance of this project and for helpful comments from Loren Brandt, Clémentine Van Effenterre, Sebastian Dyrda, Robert McMillan, William Arbour and Asa Motha-Pollock. The Longitudinal Administrative Databank (LAD) analysis presented in this paper was conducted at the Toronto Research Data Centre (RDC), which is part of the Canadian Research Data Centre Network (CRDCN). The views expressed in this paper are those of the authors and do not necessarily reflect those of the CRDCN, the partners of the CRDCN, or Statistics Canada.

1 Introduction

Wealth inequality is a topic of interest for both scholars and the broader public alike, particularly since wealth-to-income ratios in several countries have risen dramatically over the last several decades (Piketty & Zucman, 2014). Despite this interest, relatively little is known about the general trends in wealth inequality because reliable data on wealth at an individual level is hard to come by in the majority of countries. Without this data, it is difficult to determine what drives wealth inequality and what policies should be used to address it.

More recently, new approaches of gathering data on wealth have been developed, allowing researchers to make progress on these key questions. The *capitalization method*, popularized by Saez & Zucman (2016) and applied to the United States (US) context, is one such advancement that combines individual-level, administrative tax data with aggregate wealth data to determine the distribution of wealth. This is done by translating capital income flows to stocks using internally consistent rates of return by asset type. This new approach has inspired a growing international effort to measure wealth inequality in countries around the world (Blanchet & Martínez-Toledano, 2023; Garbinti et al., 2020).

This paper contributes to this global project by developing a novel wealth inequality series for a major, non-European G7 country: Canada. While Canada has seen aggregate wealth quadruple from 1990 to 2018, information on the distribution of wealth remains limited because surveys on assets and debts, such as the Survey of Consumer Finances (SCF) in the US, have been relatively sparse and unreliable in Canada during this period. Beyond domestic concerns, an understanding of wealth inequality in Canada is interesting due to its proximity to the US. With similar cultures and economic integration through free-trade agreements, one might expect the two countries to exhibit similar patterns of inequality as Saez & Veall (2005) found in the case of *income* inequality. On the other hand, differing trends might indicate that the factors driving high and rising US wealth inequality are a distinctly American phenomenon.

To investigate this, I estimate the level of wealth inequality in Canada from 1990-2018 using the capitalization method. I use administrative tax data from the Longitudinal Administrative Databank (LAD), which is a 20% sample of tax-filing Canadian census families, and the National Balance Sheet Accounts (NBSAs) to estimate wealth at the household level. I then compute the top 1% wealth share in Canada on an annual basis with the results presented in Figure 1. I find that the top 1% share grew from 15.3% in 1990 to 19.7% in 2008 before falling back down to 17.5% by 2018. These estimates fall far short of those found in the US (around 35%) with a less pronounced upward trend, suggesting that the American experience is fairly unique and does not extend to Canada.



Figure 1: Top 1% Wealth Share in Canada

This figure plots the share of wealth owned by the wealthiest 1% in Canada from 1990-2018 using the capitalization approach from Saez & Zucman (2016). Wealth is inferred based on capital income flows in administrative tax data, the Longitudinal Administrative Databank (LAD), while aggregate wealth is measured in the National Balance Sheet Accounts (NBSAs). The unit of analysis is the census family and so the top 1% refers to the wealthiest 1% of families.

Looking at the broader distribution, I find two contrasting trends. On the one hand, inequality between the middle and upper-middle classes appears to be shrinking. The Gini coefficient is falling and the ratio of wealth held by the 90th percentile and 50th percentile is decreasing, which tells a story of middle class prosperity. On the other hand, inequality is worsening at the extreme ends of the distribution. The share of wealth going to the top 0.1% has gone up over 40%, from 4.5% to 6.4%, and the threshold to enter the top 1% has grown from \$5 million to \$17 million. At the other end of the distribution, the gap between the median Canadian household and those with no wealth is growing dramatically, as median net worth grew 330% during this period. These divergent findings potentially explain some of the disconnect between the observed trends in the Gini coefficient and the top 1% share and the public perception that inequality is growing rampantly.

I then explore the observed pattern in the Canadian top 1% share using two separate decomposition approaches. First, I use linear asset decompositions to see the role played by different assets in driving the trends. I find that, while equities and other investment assets became more concentrated in this period, which puts upward pressure on the top 1% share, other asset classes became less concentrated, which offsets this effect. As a result, most of the increase in the top 1% share during this period came from an aggregate shift towards more concentrated assets such as equities and other investment vehicles.

Then, I use the *synthetic* savings¹ decomposition from Saez & Zucman (2016) to break changes in wealth down into capital gains, savings rates and income inequality, to determine which factors account for the observed trends. I find that the savings rates behaviour of the top 1% explains most of the observed trend, where high savings prior to the Great Recession was followed by a collapse in the savings rate after. Had the savings rate simply remained constant at the average rate from the period, there would be a steady upward trend in the top 1% share that would be continuing today. This finding also suggests that the top 1% may be playing a disproportionate role in the collapse of Canadian business investment observed during this period. Capital gains appear to have only played a limited role because the gap in the rate of return between wealth groups remained constant throughout this period. This can be explained by the fact that higher capital gains for the top 1% after the Great Recession were matched by skyrocketing housing values for the middle class during this time.

This paper contributes to the literature on wealth inequality in a number of ways. First, it makes an important contribution to our understanding of wealth inequality in Canada. The estimated wealth series in this paper is the first to document trends in wealth shares in Canada on an annual basis over an extended period of time.² This was not possible in prior research because the Survey of Financial Security (SFS), which is the main survey on assets and debts in Canada, is only reliably available for 1999 and then every three years starting in 2012. This is also the first paper to use administrative data to estimate the distribution of wealth in Canada, which is more representative of the population than self-reported surveys. Prior research had attempted to resolve the problems of self-reporting in the SFS by exploring the uncensored survey data (Brzozowski et al., 2010) or by using the Pareto-interpolation approach of Vermeulen (2018), where lists of billionaires are joined to the SFS and fitted with a Pareto-distribution (Davies & Di Matteo, 2020; Wodrich & Worswick, 2020). I find that the top 1% share in Canada is higher than the values found in the raw SFS, but lower than the estimates reported using the Pareto-interpolation approaches.

Second, this paper contributes to the literature related to the capitalization method itself. While other methods for measuring wealth have been applied, such as the use of estate tax records (Kopczuk & Saez, 2004; Alvaredo et al., 2018) or Pareto-interpolations (Vermeulen, 2018), the capitalization method has become a staple of wealth measurement when administrative data is available. However, the method continues to be refined. This paper introduces a new approach for imputing assets with no capital income flows, such as housing and pensions. Rather than creating arbitrary bins as in Garbinti et al. (2020),

¹Savings are called synthetic because it tracks the implied savings rate of the top 1% wealth group even though the specific people may not be the same.

²Statistics Canada's newly developed Distributions of Household Economic Accounts (DHEA) does provide annual data back to 2010, but only reports the wealth of the top quintile (20%) of the distribution.

this paper uses a distribution regression approach (Chernozhukov et al., 2020), where the conditional distribution of assets is estimated in the SFS and then used to predict the range of possible asset holdings of each household in the tax data. Each household's asset holdings are then simulated by drawing from their own predicted distribution. This data-driven approach allows for a continuous distribution of assets, which better preserves the observed variation in asset values. This paper also shows that the recent debates around heterogeneous returns (Smith et al., 2023; Saez & Zucman, 2020) make little difference in the Canadian context as the results change very little, suggesting this debate is primarily an American issue.

Third, this paper provides a new point of reference for cross-country analyses of wealth inequality. What is striking about the Canadian estimates found in this paper is that the top 1% share is fairly small compared to its peer countries. The top 1% wealth share found in this paper is only half the US share in Saez & Zucman (2016) and falls below a host of European countries including France (Garbinti et al., 2020), Spain (Martínez-Toledano, 2019), Italy (Acciari et al., 2024), Denmark (Jakobsen et al., 2020) and Germany (Albers et al., 2022). Some of the closest estimates to the Canadian ones are from other anglophone countries: the UK (Alvaredo et al., 2018) and Australia (Katic & Leigh, 2016), which raises questions about the impact of Commonwealth institutions on wealth inequality.

Finally, by exploring the causes for Canada's low top wealth share, I also contribute to the literature on the determinants of wealth dynamics. Why wealth accrues to those at the top of the distribution is a question that has received substantial attention in the literature (Krusell & Smith, 1998; Benhabib et al., 2011; De Nardi & Fella, 2017; Benhabib et al., 2019; Fagereng et al., 2020; Bach et al., 2020; Hubmer et al., 2020; Kuhn et al., 2020). Using linear asset decompositions, I find that the difference in the top 1% share between Canada and the US is not driven by any one asset, but is driven instead by greater concentration across every asset class. Furthermore, this paper's finding that changes in the savings rate of the top 1% is the primary driver of top wealth share dynamics brings it in line with the experience of France (Blanchet & Martínez-Toledano, 2023), but not with the US, where asset price cycles (Kuhn et al., 2020) and income inequality (Hubmer et al., 2020) were found to play larger roles. This suggests that while a stagnating top 1% wealth share is positive from an inequality perspective, it may also indicate declining business investment and diminished economic growth prospects.

The paper starts with Section 2, which provides an overview of the data and trends related to wealth in Canada. Then, Section 3 covers the capitalization method in detail, Section 4 presents the main results and Section 5 explores the causes of the trends using decompositions before Section 6 concludes.

2 Wealth in Canada

2.1 Defining Wealth

Before going further, it is necessary to define the concept of wealth. Marketable wealth is the current market value of all assets owned by households minus their debts. The international standards of the System of National Accounts (SNA) limits assets to those "which are subject to ownership rights and from which economic benefits may be derived by their owners by holding them or using them in an economic activity" (United Nations, 2010). Some notable potential assets that are omitted from this definition include: promises of future government spending (such as government pensions), unfunded pensions, consumer durables and human capital. Using these criteria, in this chapter, the following assets make up marketable wealth: public and private equity, currency and deposits, bonds and short-term paper, unincorporated business assets, pension assets, principal residences and other real estate properties, while debts include mortgages and non-mortgage loans.³

2.2 Aggregate Wealth in Canada

Data on aggregate wealth in Canada from 1990-2018 comes from the National Balance Sheet Accounts (NBSAs), which record the stock of assets and debts in the economy for a variety of sectors.⁴ In particular, the focus will be on the household and non-profit institutions serving households sector, which aligns with the definition of wealth above. Within each sector, the NBSAs break wealth down further into different instrument types such as residential structures, debt securities and listed shares, which is instrumental for the capitalization approach. Because the NBSAs follow the SNA framework (United Nations, 2010), these estimates of aggregate wealth are also comparable to other countries such as the United States, Great Britain and France.

This period provides an interesting backdrop for the study of wealth inequality since aggregate wealth has exploded over this time. Table 1 shows that aggregate wealth in Canada increased by a factor of four in real terms between 1990 and 2018, when it surpassed \$10 trillion dollars. The growth in aggregate wealth has remained steady for most of the period, with average growth rates above 4% for each of the seven year periods. The growth in wealth has also been disproportionate to income. The capital to income ratio in Canada has grown from almost 300% to over 700% of national income. Naturally, the average family net worth has also risen over this period, with the average currently at \$608,848 per family.

 $^{^{3}}$ For a complete explanation of the assets included in the definition of wealth refer to Appendix A.1

⁴This period is chosen because in 2012, revised estimates of the NBSAs were published going back only to 1990 and in this way, the series can only remain consistent up until then.

			Years		
	1990	1997	2004	2011	2018
Total Net Worth (in Millions)	2,518,539	3,827,947	5,018,988	7,265,858	10,296,541
Average Growth Rate		6.21	4.03	5.56	5.14
Capital to Income Ratio $(\%)$	298	435	472	575	716
Average Net Worth	219,029	296,278	362,285	474,579	608,848
Savings Rate (%)	13.1	4.5	1.8	4.5	1.4
Number Of Families	11,498,655	12,920,130	13,853,690	15,310,120	16,911,500

Dollar variables expressed in 2018 CAD $\$

Table 1: Net Worth Summary Statistics

This table presents summary statistics of household net worth for five individual years. Aggregate net worth is computed using data from the National Balance Sheet Accounts (NBSAs). Aggregate net worth, expressed in millions, totalled \$10.3 trillion CAD in 2018. The average in the growth and savings rate refers to the average over the preceding seven years. The savings rate is presented from the household's current and capital account, which also is part of the system of national accounts. The capital to income ratio (K/Y) is the ratio of total net worth to income. Average net worth is with respect to the family unit.

All of this is happening while the aggregate household savings rate, measured in the Current and Capital Account for households, has been declining. This suggests that rising savings in the household sector as a whole cannot explain the increase in aggregate wealth and that other factors, such as capital gains, may be important.

To better understand where this upward trend in aggregate wealth is coming from, it is useful to look at its component parts. Figure 2 plots aggregate wealth over time, broken up into six assets groupings.⁵ A major takeaway from the figure is the key role played by housing. Net housing wealth has increased from \$950 billion in 1990 to over \$4 trillion in 2018 CAD, which is consistent with the rise of Canadian home values - the Canadian home price index shows that the price of a home has tripled since 2000. However, housing alone is not driving the trend in aggregate wealth. Both housing and pensions have been growing at a similar rate and have maintained a consistent share of aggregate wealth over the period. Housing has remained between 34-40% of aggregate wealth, while pensions have hovered around 30%.

That is not to say that there has not been any change to the aggregate wealth portfolio, as Canadian equities have increased dramatically during this period. In 1990, Canadian equities directly held by households were worth only \$112 billion, while in 2018, this number was \$1.2 trillion - ten times the amount from almost thirty years prior. This change boosted

⁵Details on how assets are divided into component parts are available in Appendix Section A.2.



Figure 2: Aggregate Household Wealth in Canada - 1990-2018

This figure depicts the level and composition of household wealth in Canada from 1990-2018 in 2018 CAD using data from the National Balance Sheet Accounts (NBSAs). Information on pensions comes from the Pension Satellite Accounts (PSAs).

the equity share of aggregate wealth from 4.4% to 11.6%. Most of this came at the expense of unincorporated business assets, which fell from 5.6% to 1.4%, and other investments, which fell from 22.7% to 17.3% of aggregate wealth. The decline in unincorporated business wealth reflects a growing trend towards incorporation among sole proprietors.

The important takeaway from this section is the remarkable increase in wealth from 1990 to 2018, but this aggregate data says nothing about how this wealth was distributed. It is on this front that information is fairly limited. There have been a couple Surveys of Financial Security (SFSs) in 1999, 2012 and 2016, but these surveys can be unreliable in capturing the wealthiest families and the sparse nature of the data - only capturing three years - makes drawing conclusions on trends in the distribution of wealth challenging. This is where this paper tries to fill an important gap in better understanding how equally this explosion of aggregate wealth in Canada was distributed.

2.3 Capital Income in Canada

Unlike wealth which is a stock, capital income is a flow that is reliably reported on tax forms annually. In Canada, this capital income data comes from the Longitudinal Administra-



Figure 3: Top 1% Share of Capital Income

This figure depicts the share of capital income earned by the top 1% of capital earners. Capital income is made up of self-employment income, dividends, capital gains, net rental income and interest and other investment income. Data comes from the Longitudinal Administrative Databank (LAD) for the years 1990-2018.

tive Databank (LAD). The LAD is a 20% sample of the annual T1 Family File (T1FF) of Canadian taxpayers, which amounts to over 5.6 million individual observations in 2018. The T1 Family File covers all taxpayers who have a social insurance number (SIN) and creates Census families that link together parents and children through information provided on the tax form.⁶ Census families are comprised of either unattached individuals or a married couple, both including their unmarried children if any.⁷ I will use the Census family as the unit of analysis because some forms of wealth, like housing, are difficult to allocate to just one individual in a household.

While the concentration of capital income has remained fairly stable in recent years, the nature of capital income has changed dramatically between 1990 and 2018. Figure 3 plots the share of capital income going to the top 1% of capital income earners. Capital income concentration rose in the 1990s, from 32.0% in 1990 to 42.6% in 2000, but it has remained around 45% since then. This relatively stable trend masks substantial changes to the overall composition of capital income. In 1990, in an era when the prime lending rate was over 14%, interest and other investment income (which includes deposits and bonds) made up 42% of all capital income. By 2017, with the prime rate down to 2.7%, investment income made

⁶This means that under the family variable for each individual in the LAD, the family's total income is pulled from the T1FF.

⁷This definition includes unmarried, adult children who continue to reside at the same residence as their parents regardless of age.

up only 6.5% of all capital income. Dividends (from 10.7% to 35%) and capital gains (from 13.6% to 31.9%) have made up the difference.

These large changes to the aggregate composition of capital income mean that we should not take the concentration of capital income shown here as representative of the top 1% wealth share. That is because each type of capital income reflects a different level of wealth. For example, a bond worth \$1,000 that generates \$10 in income at a 1% interest rate implies a higher level of net worth than a stock worth \$200 that pays out a \$10 dividend at a 5% rate of return. That means that even though the top capital income earners' share of capital income is rising, it may not reflect a larger share of assets if that income is generated from higher return assets like stocks. This is the key principle of income capitalization - if one knows the rate of return on different asset classes, they can infer the level of wealth based on the capital income flow. This process will be described in detail in the next section.

3 Method

3.1 Capitalization Method

The income capitalization method is an approach to convert capital income flows into the stock value of an asset - and a form of marketable wealth. Suppose we have an asset j and a distribution of agents $i \in \{1, ..., k\}$. The stock of asset j, held by individual i is W_{ij} . This is connected to the capital income flow from asset j received by individual i, I_{ij} , through the annual return of that individual's asset, r_{ij} . This can be written in the following way:

$$r_{ij}W_{ij} = I_{ij} \Longrightarrow W_{ij} = \frac{1}{r_{ij}}I_{ij} = \beta_{ij}I_{ij}$$

Here, β_{ij} is called the *capitalization factor*, which is the inverse of the annual return for that asset. An individual's total wealth, W_i , is the sum of the holdings of each asset type.

$$W_i = \sum_j \beta_{ij} I_{ij}$$

This equation tells us that we can infer the wealth of an individual based on their capital income if we can estimate these annual returns for different assets. However, the primary challenge when using this method is that it is impossible to know the exact annual return received by each individual and each asset. To deal with this, I follow Saez & Zucman (2016) in making a simplifying assumption: the annual return for each asset is constant across all individuals, $r_{ij} = r_j$.

Annual returns by asset, r_j , are computed using aggregate wealth data from the NBSAs and aggregate capital income flows from the LAD. To do this, capital income flows from the LAD are matched with categories of assets in the NBSAs - a process presented in Appendix Section B.1.1. As an example, for Canadian equity, the relevant NBSA variables are listed and unlisted shares, while the two sources of capital income from owning shares are dividends and capital gains, which are reported in the LAD. Then, to get the rate of return by asset, the aggregate capital income flow is divided by the aggregate of the corresponding stock in the NBSA. In the case of Canadian equity, in 2018, the combined rate of return of dividends and capital gains on listed and unlisted shares is 13.8%. The capitalization factor, which is the inverse, is then 7.2. On the other hand, for the other investments category that comprises mostly of bonds and deposits, the rate of return is only 1% in 2018, down from 9.7% in 1990, which coincides with the secular decline in interest rates over the period.

This approach for estimating returns is useful for two reasons. First, it allows for a consistent approach across assets in estimating returns. Rather than relying on asset-specific estimates from a collection of sources that use different methods, this approach is consistent across all assets. Second, using this approach, the total wealth is going to be consistent with the total wealth in the NBSAs. This is because the constant used to scale income to wealth is exactly the ratio of aggregate wealth from the NBSAs to aggregate income.

3.2 Assets Without Capital Income Flows

The capitalization method only works when there are capital income flows observed in the microdata. In Canada, there are a couple assets that do not show up clearly in the LAD: housing and pensions. Unlike in the United States, there is no mortgage interest deduction or other observable metric of homeownership. For pensions, while there are indicators of pension withdrawals, this does not capture the substantial pension wealth that is owned, but not withdrawn by those who are not yet retired. To address this gap in the data, I use a novel imputation approach for this literature that is based on *distribution regression*.

A more complex imputation approach is important when trying to measure wealth inequality because the objective is to mimic the distribution of assets rather than to minimize the squared error at the individual level. Standard imputation approaches that rely on linear regression lead to estimates of the asset distribution that have too little variance. This arises because these models rarely predict extreme values for an asset, such as zero, which is not realistic given that many households do not own their homes or have a pension. Instead of estimating the conditional expectation of asset values, I estimate the conditional distribution of housing and pensions using distribution regression techniques on the SFS and then impute the value of assets in the LAD by drawing from a family's predicted conditional distribution. Since this approach yields a non-parametric distribution for each family, which can give a high probability of drawing a zero, this approach generates a more realistic distribution of assets across the population.⁸

For the distribution regression, this paper follows the work of Chernozhukov et al. (2020). Distribution regression allows for the generalization from a univariate cumulative distribution function (CDF) to a conditional CDF. We can write the conditional distribution of Y as a function of covariates X as follows:

$$F_{Y|X}(y \mid x) = E[\mathbb{1}(Y \le y) \mid X = x]$$

The distribution regression model can then be written as:

$$F_{Y|X}(y \mid x) = \Lambda \left(x'\beta(y) \right)$$

where $\Lambda(\cdot)$ is a link function - in the case of this chapter, the logit transformation, x is a vector of covariates, and $\beta(y)$ is an unknown vector of coefficients that depends on the value of y.⁹ If we think in terms of a single threshold y, then this is just a binary regression of whether one is above or below that threshold. Doing so for many thresholds yields the distribution regression model.

This object is useful because it can be inverted to yield the conditional quantile function.¹⁰ With the conditional quantile function, I can then predict the value of the asset at each percentile $p \in \{0, 0.01, ..., 0.99, 1\}$ for each family. It is then straightforward to draw from a uniform distribution, $p \sim U[0, 1]$, for each family and assign the value of the asset that corresponds to that percentile drawn. This way, the resulting distribution of the asset will be preserved.

I estimate the conditional quantile function using the SFS. I model net housing and pension wealth as a linear function of both individual and city-level variables that are found in both the SFS and the LAD. After estimating the conditional quantile function for all one hundred quantiles in the SFS, I use the estimated coefficients to predict each quantile for

⁸Non-parametric approaches such as in Garbinti et al. (2020), where households are grouped into income and age bins and then assigned the average value of the bin to the proportion in that bin with a positive value represents a middle ground between this paper and standard prediction methods. While they do produce a more realistic distribution of assets, the within-bin variance is limited and the choice of bins is arbitrary and not data-driven like in this paper's approach.

 $^{^{9}}$ In the context of this chapter, because y is a continuous variable (eg. housing value), it will be approximated using a series of 100 grid points at each percentile of y

¹⁰This approach is theoretically equivalent to estimating a quantile regression directly in large enough samples. However, the distribution regression is faster to run computationally and more flexible (Chernozhukov et al., 2020).

every family in the LAD using the set of common characteristics. Then, from each family's estimated distribution, I draw a value of housing and pension wealth from this estimated distribution. In theory, one can repeat these draws many times and compute the average values each time, but in practice, the variance in the top 1% share across simulations is imperceptibly small, so I use a single simulation as the main result.¹¹ While the SFS has its limitations in terms of measuring wealth at the top of the distribution, it is more reliable for housing and pensions, which makes it suitable for this exercise. This is because housing values are publicly available and well-known to respondents and pensions have annual contribution limits, which means that top-coding is unlikely to be an issue.

3.3 Heterogeneous Returns

The main critique of the capitalization method is the assumption of homogeneous returns across the wealth distribution. Bach et al. (2020) and Fagereng et al. (2020) have shown how rates of return to wealth are higher for those at the top of the distribution than those at the bottom of the distribution. They show this to be the case, not just because different wealth groups hold different assets, but within even narrowly-defined asset classes as well. If this is the case, then the assumption of homogeneous returns would lead to an overestimate of the wealth of those at the top of the distribution since the capitalization factors would be overstated. Smith et al. (2023) argue that correcting for this can dramatically lower the top 1% share of wealth in the United States relative to what Saez & Zucman (2016) estimate in their paper.

How to address this challenge remains a contested topic. Smith et al. (2023) take a stance on the rate of return for the rich in the fixed income claims asset category (which is similar to the interest and other investment income category in Canada). They argue that the rate of return on the Moody's AAA corporate bond, which averaged 6.0% in the 2000s and 4.2% from 2010-2016, is a good proxy for the rate of return of the top 0.1% and the 10-year US treasury bond is a good proxy for the next 0.9%. They then capitalize the income of these groups according to those returns and then compute the residual rate of return for the remaining population, which is close to 0. Saez & Zucman (2020) contest this and argue that this correction is excessive and does not line up with the data. They argue that there is no evidence that suggests an interest rate premium for the wealthiest that is close to the Moody's AAA corporate bond rate. At most, they find that the interest rate of the top 1% wealthiest households is 1.4 times higher than the average after 2008, which is far below the Moody's rate. One reason for the disagreement is that Smith et al. (2023)

¹¹In Appendix Section B.1.2, I discuss the method in greater detail and how it performs relative to some alternative approaches.

apply the Moody's rate to the top 1% of *interest* income earners rather than the wealthiest 1% which may lead to its own downward bias.

To address this issue in this paper, I estimate the wealth distribution using all possible approaches. For the Saez & Zucman (2020) approach, since I cannot perform the same analysis of rates of return by wealth in the SFS as was done in the SCF, I will assume the same interest rate premium for the wealthy, 1.4 times, as was found in the United States. This assumption is reasonable for a couple reasons. First, interest rates have followed a similar path in Canada and the United States since the Great Recession. Second, with global capital markets, the wealthiest Canadians have access to many of the same corporate bonds and financial instruments as their American counterparts and likely share similar portfolio strategies. For the Smith et al. (2023) approach, I also use the Moody's AAA corporate bond rate for the top 0.1%, but use the 10-year Canadian government bond rate for the next 0.9%. In Figure 5, I show that the different approaches do not change the results substantially like they do in the US context. As a result, I use the Saez & Zucman (2020) correction as the baseline results for this paper since the rest of the method closely follows their approach.

4 Results

4.1 Wealth Distribution in Canada

Using the capitalization method and the adjustments described above, I estimate the share of wealth for a number of wealth groups in the population. The results are presented in Figure 4. There are a few important observations. First, top wealth shares have not increased substantially over this period. The top 1% share rose only from 15.3% to 17.5% from 1990 to 2018, a very slight increase. While the top 1% do own a large share of overall household wealth, these results suggest that concerns over dramatically *worsening* wealth inequality in Canada may be overblown.

Second, most of the movement observed in the top 1% over this period is driven by the top 0.1%. The top 0.1% saw an increase in wealth from 4.5% of total wealth in 1990 to 7.8% in 2011 before dropping to 6.4% in 2018. The next 0.9% on the other hand was fairly steady for most of this period, hovering between 10.1% and 11.9%. This suggests that while the top 1% is a popular subgroup, most of the movement is actually happening amongst the very wealthy. The top 0.1% in 2018 was comprised of just under 17,000 families, with a threshold wealth level of \$16,779,000 and average wealth of \$39 million.

Finally, when the top 0.1% share was rising dramatically up to 2011, most of these gains



Figure 4: Share of Wealth Held By Various Wealth Groups

This figure presents the share of wealth held by a variety of wealth groupings in Canada from 1990-2018 using the capitalization method from Saez & Zucman (2016). Wealth is inferred based on capital income flows in administrative tax data, the Longitudinal Administrative Databank (LAD), while aggregate wealth is measured in the National Balance Sheet Accounts (NBSAs). The unit of analysis is the census family and so the top 0.1% refers to the wealthiest 0.1% of families.

came at the expense of those near the top rather than those at the bottom. The 95th-99th percentile group fell from 21.3% to 19.8% during this period, while the 90-95 group fell from 15.7% to 14.4%. The bottom 75% actually rose during this period from 21.3% to 23.1%. This result provides suggestive evidence that the increase in the top 0.1% share may be a result of already wealthy people amassing much larger sums of wealth.

Wealth shares represent one measure of wealth inequality, but there are others. Table 2 presents some important alternative measures of wealth inequality over this period. The Gini coefficient, which ranges from 0 to 1, with 0 being completely equal and 1 meaning one person holds all the wealth, is a more general indicator of inequality that is not solely focused on the very top. According to this Gini coefficient, wealth inequality has been steadily falling over the course of the period. Another measure that can be informative about wealth inequality is the ratio of the 90th percentile threshold to the median. This measure, like the Gini coefficient, has also been declining over time, which suggests there is compression of the wealth distribution over this period.

These results are not necessarily contradicting the trends in wealth shares from Figure 4. The Gini coefficient is capturing the fact that, although the top 0.1% share rose over this period, it rose at the expense of the next 9.9% of the distribution, while the wealth of those in the bottom 75% saw a relative increase over this period. Similarly, the 90-50 ratio is

			Years		
	1990	1997	2004	2011	2018
On Aggregate					
Gini Coefficient	0.711	0.710	0.701	0.698	0.699
90/50 Ratio	7.9	7.3	6.5	6.3	6.1
Median Wealth	$72,\!117$	$104,\!012$	$137,\!692$	$183,\!351$	239,000
Ν	$11,\!498,\!655$	$12,\!920,\!130$	$13,\!853,\!690$	$15,\!310,\!120$	16,911,500
-					
<i>Top</i> 0.1%					
Wealth %	4.5	5.2	6.3	6.3	6.4
Threshold	5,216,488	7,396,588	9,995,430	13,486,392	16,779,000
24	0.055.151	1 - 100 000			20.050.000
Mean	9,855,151	15,400,636	22,821,815	29,898,696	38,958,239

Dollar variables expressed in 2018 CAD \$

Table 2: Wealth Inequality Measures

This table presents some key measures of wealth inequality in Canada for five years from 1990-2018 using the wealth estimates from this paper. The Gini coefficient is measured on a scale of 0 to 1, with higher numbers meaning more inequality. The 90-50 ratio is the ratio between the threshold for the 90th percentile and the median. N represents the number of families.

capturing this compression between the family with the median level of wealth and the 90th percentile. Another way to see this is to look at the growth rates of mean wealth by wealth group. The bottom 50% and the 50-75th percentile group saw their average wealth increase 3.40 and 2.94 times between 1990 and 2018 respectively. The 75-90, 90-95, 95-99th percentile groups all had growth in average wealth between 2.55 and 2.65 times for the period. Lastly, the top 0.1% and the next 0.9% group saw a 3.95 and 2.85 times increase respectively. So, while all groups saw their wealth increase significantly over this period, growth rates were u-shaped across the wealth distribution, with those at either end growing fastest.

4.2 Robustness

These results are not overly sensitive to choices around how to compute wealth. Figure 5 plots a number of alternative measures of the top 1% wealth share. The first alternative is where returns in the interest and other investment category are assumed to be homogeneous. Since the correction proposed by Saez & Zucman (2020) only applies after the Great Recession,



Figure 5: Alternative Estimates of the Top 1% Share

This figure presents the share of wealth held by the wealthiest 1% in Canada from 1990-2018 using different approaches. The first three variations involve the three approaches to addressing heterogeneous returns to the interest and other investment income asset class discussed in Section 3.3. The last two involve alternative approaches to measuring the equities and housing categories, respectively.

the difference only emerges in more recent years. As expected, the top 1% wealth share is higher when all returns are treated equally, but never by more than 2 percentage points and the trend is preserved. On the other side is the correction proposed by Smith et al. (2023), where the Moody's AAA corporate bond rate is used to capitalize the top 0.1% share and the 10-year Canadian government bond rate is used for the next 0.9%. This leads to a lower top 1% share in every year, but again is never more than 3 p.p. below the preferred estimate and the trend remains the same. While these adjustments do impact the share of other investments going to the wealthiest 1%, one reason it does not have a large impact on the share of wealth is because these assets only make up around one-fifth of all assets.

Two other alternatives generate very similar estimates to the preferred estimate. The first involves capitalizing dividends only, which yields very similar results except for in 1994, which was a big year for capital gains due to a policy change. The other is the estimate of wealth that uses census tract housing values as a proxy for housing wealth as described in Appendix Section B.1.1. Since housing wealth is not very prominent in the portfolio of the top 1%, the difference in housing wealth assigned across methods has a minimal effect on the total wealth share. All together, these results suggest that the preferred estimate is fairly robust to alternative specifications.

Finally, in Appendix Figure 1, I show how different definitions of wealth affect the estimated top 1% share of wealth. I find that using the NBSA definition of aggregate pension



Figure 6: Comparing Canadian Estimates of the Top 1% Wealth Share

This figure presents the share of wealth held by the wealthiest 1% in Canada from 1990-2018 from different sources. The preferred estimate comes from using the capitalization method on administrative tax data. The Survey of Financial Security (SFS) data is presented both using the restricted access version that is not top-coded and the publicly available version that is top-coded. Davies & Di Matteo (2020) and the Parliamentary Budget Officer (PBO) (Wodrich & Worswick, 2020) use Pareto-interpolation methods between the SFS and the billionaire "rich lists" from the Canadian Business Magazine and Forbes.

wealth and including consumer credit offset one another and result in a very similar top 1% wealth share. Including an estimate of offshore wealth based on the work of Alstadsæter et al. (2018) increases the top 1% share by around 1-2 p.p. - a meaningful, but ultimately small increase.

4.3 Comparisons to Other Estimates

The estimates of the Canadian wealth distribution stand somewhat in contrast to previous Canadian estimates using alternative methods. Figure 6 plots the top 1% wealth share in Canada from a variety of existing estimates. The first set of estimates come from the Survey of Financial Security (SFS) with both the top-coded public version and the restricted access version that is not top-coded. The second set of estimates are based off the method outlined by Vermeulen (2018), where lists of the richest individuals - Forbes in the United States and the Canadian Business Magazine in Canada - are spliced together with wealth surveys, fitted with a Pareto-interpolation and then used to compute wealth shares. Davies & Di Matteo (2020) and the Parliamentary Budget Officer (Wodrich & Worswick, 2020) do this for Canada, although the definitions of wealth and sources of billionaire wealth differ.

There are a few important takeaways from this figure. First, there is obvious value

to having an annual wealth trend. Since the survey data are only available for three years between 1990-2018, it is difficult to deduce any trend in the survey-based estimates. With the capitalization method, a more complete picture can be formed. Second, the estimates using the capitalization method are much closer in magnitude to the raw survey estimates than the Pareto-interpolated ones. One reason for this is that pensions and principal residences, which make up 70% of aggregate wealth are imputed based off the SFS data. However, these results suggest that perhaps the SFS does not do that poorly at approximating the wealth of the top 1% either. While top-coding appears to have played a role in 1999, as evidenced by the gap between the public use SFS and the restricted access one,¹² there appears to be less of a difference in 2012 and 2016. When applying the capitalization method, which theoretically does a better job of capturing large wealth values, the top 1% wealth share only increases slightly compared to the SFS estimates.

On the other side are the Pareto-interpolated estimates, which estimate much higher levels of wealth inequality relative to the capitalization method. One reason for this is the definition of wealth used in both cases. The definition of wealth used by Davies & Di Matteo (2020) omits employer pension plans, which form a substantial portion of total wealth - 19% in 2016. The PBO estimate from 2021, which does not omit pensions, is 6 p.p. lower. In addition, the billionaire lists seem to use a broader concept of wealth than the one employed in this paper. Forbes says that they include "art, yachts, planes, ranches, vineyards, jewelry, car collections and more" in their definition of wealth for billionaires. However, this paper does not count consumer durables or vehicles in its definition of wealth, nor expensive art or jewelry.¹³

Overall, the capitalization method results fall somewhat in between the two sets of estimates, but follow the raw SFS estimates more closely. This finding suggests that perhaps wealth inequality is not as severe in Canada as previous estimates suggested and that the narrative of runaway wealth inequality may be unfounded.

It is also interesting to compare the results for Canada to some other major countries. Figure 7 plots the top 1% wealth share for Canada alongside the United States (Saez & Zucman, 2016) and France (Garbinti et al., 2020). What this figure shows is that the top 1% share in Canada is lower than other countries even when using the same approach. In 1990, the top 1% share was 28.5% in the US, 15.3% in Canada and 17.3% in France. By 2014, the US sat at 36.6%, Canada at 16.6% and France at 24%. The following section will

 $^{^{12}}$ This gap is discussed in detail by Brzozowski et al. (2010)

¹³While this closes some of the difference between the approaches, there would still be a gap. Some of this remaining discrepancy could be attributable to differences in how the value of equities held by billionaires is calculated. Another plausible explanation is that the assumption of a constant Pareto coefficient across the top of the wealth distribution does not hold perfectly.



Figure 7: International Estimates of the Top 1% Wealth Share

This figure presents the share of wealth held by the wealthiest 1% for Canada, the USA and France between 1990 and 2019. The United States estimates come from Saez & Zucman (2016) and the distributional national accounts they update. The French estimates come from Garbinti et al. (2020). Both these estimates use the capitalization method as well.

explore some explanations for the differences between the countries.

5 Decompositions

The results of the previous section leave some pressing questions. The first question is why wealth inequality, despite rampant discussion of its increased prevalence, does not appear to be increasing in Canada. In fact, the level of wealth inequality appears to be similar compared to 20 years ago. The second question is why the level and trend are so different from other countries like the United States and even France. Despite being neighbours and sharing many cultural similarities, the top 1% wealth share in the US is around double that in Canada. This section seeks to explore these questions in greater detail.

The literature on wealth inequality has focused on different explanations for wealth inequality dynamics. Kuhn et al. (2020) emphasized how asset price cycles played a key role in shaping the trends in wealth inequality in post-war America. When housing prices rise, this disproportionately helps the middle class, where housing is a large share of their wealth portfolio. Conversely, when equity markets are strong, this largely helps the rich, who own the vast majority of equity wealth. Another explanation comes from Hubmer et al. (2020), who argue that rising after-tax income inequality following the introduction of major tax cuts on the rich played a major role in driving up top wealth shares in the US since the 1980s. Blanchet & Martínez-Toledano (2023) find that in France, trends in the savings rate of the wealthiest appear to best explain the observed dynamics of top wealth shares. To investigate which explanation best fits the Canadian context, I employ two decomposition exercises: a linear asset decomposition and a synthetic savings decomposition.

5.1 Linear Asset Decompositions

The linear asset decompositions are based on the observation that the share of wealth owned by a given wealth group, sh_{wg} , can be expressed as a weighted average of the shares of each asset class owned by each wealth group, ψ_{jg} , where the weights are the aggregate share of each asset class, ω_j .

$$sh_{wg} = \frac{W^g}{W} = \sum_j \frac{W_{jg}}{W} = \sum_j \frac{W_j}{W} \frac{W_{jg}}{W_j} = \sum_j \omega_j \psi_{jg} \tag{1}$$

In Figure 8, I plot the values of ψ_{jg} for the top 1% wealth group and the aggregate asset shares, ω_j . The aggregate asset shares show an increasing trend in Canadian equity and a declining share of aggregate other investment wealth, while the remaining categories hold fairly steady throughout this period. Since equities and other investments are both more concentrated than the total wealth share, increases in the aggregate share of these assets would lead to an increase in the total wealth share of the group. However, because the aggregate share of other investments is falling, this will serve to counteract the increase in the aggregate equity share. In terms of the top 1% shares, the top 1% share of other investment wealth was rising throughout this period, with declining shares of secondary residences wealth, unincorporated business wealth and pensions.¹⁴ Decomposing wealth into these components can help in better understanding what is driving the overall trends in wealth inequality.

Equation 1 can be used to help explain exactly how much of the change over time in the share of wealth owned by a given wealth group is driven by changes to the aggregate asset composition versus changes to within-asset concentration. The difference in wealth shares for a given group over time, from t_1 to t_2 , can be written as:

$$sh_{wg}^{t_2} - sh_{wg}^{t_1} = \sum_j \omega_j^{t_2} \psi_{jg}^{t_2} - \sum_j \omega_j^{t_1} \psi_{jg}^{t_1}$$

Adding and subtracting $\sum_{j} \omega_{j}^{t_2} \psi_{jg}^{t_1}$, means the above expression can be re-written in terms of the component that is explained by changing aggregate wealth shares and the component

¹⁴In Appendix Figure 5, I present the asset compositions of other wealth groups over time.





This figure plots the breakdown of the top 1% wealth share by the aggregate asset share (left-panel), ω_j , and the share of each asset held by the top 1% (right-panel), ψ_{jg} . Wealth is estimated using the capitalization method and the unit is the census family. The overall top 1% wealth share is indicated with the dotted line.

that is explained by changing within-asset concentration.

$$sh_{wg}^{t_2} - sh_{wg}^{t_1} = \sum_a \underbrace{\psi_{jg}^{t_1}(\omega_j^{t_2} - \omega_j^{t_1})}_{\text{agg. asset shares}} + \sum_j \underbrace{\omega_j^{t_2}(\psi_{jg}^{t_2} - \psi_{jg}^{t_1})}_{\text{within-asset concentration}}$$
(2)

The first term captures the extent to which wealth shares of a group change when the aggregate asset composition changes, holding the share of wealth held by that wealth group constant at their level in t_1 . It is important to note that the aggregate asset composition should always sum to one, $\sum_j \omega_j = 1$, so changes to aggregate asset shares should be thought of as a change to the entire distribution and not to a single asset. The second term captures the extent to which changes in the concentration of various assets affect the share of wealth held by a group, holding the composition of assets fixed at their shares in t_2 . Computing each of these terms will yield the amount each component contributes to the change in wealth shares over time.

In Table 3, I present the results of the decomposition exercise for a few key time periods. First, I show the period when the top 1% share was increasing the most, rising 3.24 p.p. from 1990-2007. I find that this increase is driven almost equally by the two explanations. This is partly because larger increases in concentration within the equity and other investment categories were offset by declines in the other asset categories. Second, following the Great

	Periods			
	(A) 1990-2007	(B) 2007-2018	(C) 1990-2018	
Δ Total Wealth Share (p.p)	3.24	-0.99	2.26	
Δ Aggregate Wealth Shares, ω_j	1.64	0.48	1.74	
Δ Share of Asset to Top 1%, ψ_j	1.61	-1.47	0.51	
Equity	1.00	-0.45	0.50	
Other Investments	1.65	0.61	2.57	
Unincor. Business	-0.11	-0.08	-0.15	
Pensions	-0.78	-0.23	-0.94	
Principal Residences	-0.00	-0.06	-0.06	
Secondary Residences	-0.15	-1.24	-1.40	

Table 3: Decomposition of Change in Top 1% Share

This table presents the results of the decomposition shown in Equation 2. It tells us the change in top 1% wealth share between 1990 and 2018 that can be explained by changing the aggregate wealth composition and the shares of each asset going to the top 1%. The change in the share of each asset going to the top 1% is further broken down into the role each asset played.

Depression, the share of assets going to the top 1% fell across almost every asset class except the other investments category while the aggregate asset composition continued to move towards more concentrated asset categories. Looking at the period as a whole, Column (C) shows that changes in the aggregate asset composition and the concentration of equity and other investments were pushing up the top 1% wealth share, while the remaining asset categories were becoming less concentrated. If the other asset categories had remained as concentrated as in 1990, the increase in the top 1% wealth share would have been 4.81 p.p, which is more than double what was actually observed. These results indicate that while inequality is rising within some asset classes, it is falling in others, which partly explains why the overall increase in the top 1% share during this period was modest.

This decomposition approach can also be used to compare across countries within a given year. In Appendix Figure 6, I plot the aggregate asset shares and the within-asset shares of the top 1% for three different countries: Canada, the US (Saez & Zucman, 2016) and France (Garbinti et al., 2020). These figures show that business equity in Canada makes up a lower aggregate share than in France and the US, while the US has a lower share of housing wealth and France has a lower share of pension wealth. France's lower share of pension wealth stems

	Cour	ntries
	\mathbf{FRA}	USA
Δ Total Wealth Share (p.p)	3.82	19.85
Δ Aggregate Wealth Shares, ω_j	3.43	3.90
Δ Share of Asset to Top 1%, ψ_j	0.39	15.95
Business Equity	-3.89	3.98
Other Investments	-2.49	5.85
Pensions	2.97	2.14
Housing	3.79	3.98

Table 4: Decomposition of Change in Top 1% Share Between Countries, 2014

This table presents the results of the decomposition shown in Equation 2. It tells us the change in top 1% wealth share between Canada, the United States and France that can be explained by changing the aggregate wealth composition and the shares of each asset going to the top 1%. The change in the share of each asset going to the top 1% is further broken down into the role each asset played.

somewhat from the definition of wealth used in these papers, where publicly funded pensions do not count towards wealth, because France has generous public pension plans that replace the need for lower wealth individuals to have separate retirement savings plans.

In Table 4, I conduct the linear asset decomposition for the year 2014 with each country. I find that almost all of the difference between Canada and France stems from the differences in aggregate asset composition, which is partly driven by the fact that public pensions are not counted as wealth. France has less concentration in business equity and other investments, but is more concentrated in pensions and housing which offsets the within-asset concentration effects. Conversely, the difference between Canada and the US arises almost entirely from greater within-asset concentration across all asset classes. This suggests that it is not a single asset class driving the difference between Canada and the US, but a more unequal distribution of *all* assets.

These linear asset decompositions highlight the role played by different assets in driving the observed trends in Canadian wealth inequality. One factor holding Canadian wealth inequality in check is the fact that some assets, such as housing and pensions, are getting more equally distributed over time. When comparing to the US, it is clear that the differences between the countries are driven by broader societal factors that lead to much greater levels of wealth concentration that are not solely within a single industry or asset class.



Figure 9: Savings Rate and Capital Gains Rate by Wealth Group, 1990-2018 This figure plots the savings rate and capital gains rate over time by wealth group as determined using the synthetic savings approach. The line represents the LOWESS (locally weighted scatterplot smoothing) estimate of the underlying scatterplot values.

5.2 Synthetic Savings Decompositions

While the linear asset decompositions highlight which assets contributed to the observed trends, they do not speak to the behaviour of the different wealth groups. To address this, I employ the *synthetic savings* decompositions from Saez & Zucman (2016), which are based on the following transition equation for wealth or any specific asset:

$$W_{t+1}^g = (1 + q_t^g)(W_t^g + s_t^g Y_t^g)$$
(3)

Here, the wealth of group g in period t + 1, W_{t+1}^g , is equal to the wealth from the previous period, W_t^g , plus any savings that happened during that period, $s_t^g Y_t^g$, where s_t^g is the savings rate of the group and Y_t^g is the income earned, all multiplied by the capital gains of the group in that period, q_t^g . Savings are called *synthetic* because from year-to-year the composition of each wealth group changes, which means the calculated savings will not reflect the same people. From this equation, changes in wealth shares can be decomposed into three components: capital gains, the savings rate and income.

One challenge is that the group-specific capital gains rate and savings rate are not directly observed in the data. This is addressed in two steps. First, the *asset*-specific capital gains rates are computed using data from the Financial Flow Account (FFA). The FFA is the flows compliment to the stocks in the NBSAs and breaks the changes in wealth over time that are observed in the NBSA data into the portion explained by investment flows and

the portion explained by other changes (eg. capital gains).¹⁵ Then, using the asset-specific capital gains, the amount each group invested in each asset can be determined as the residual of the transition equation. Summing across assets will yield the total amount of savings for each wealth group and dividing this by the level of income for that wealth group will yield the savings rate. The average capital gains rate for the wealth group can then be reverse-engineered using the computed savings rate.¹⁶

In Figure 9, I plot the LOWESS (locally weighted scatterplot smoothing) values for the computed synthetic savings rate and the capital gains rate for the top 1% and the remaining 99% of the distribution.¹⁷ I find that the savings rate of the top 1% was rising in the 1990s, from about 20% to over 40% of before-tax income, before falling back down over the 2000s and rising again in recent years. The savings rate of the bottom 99% fluctuates much less and hovers around 10%.¹⁸ I find that the trends in capital gains are quite similar for the top 1% and bottom 99% - between 4 and 6% during this time period. These results suggest that the savings rate may be driving trends in wealth inequality more than capital gains.

In Figure 10, I show how the growth rate in the top 1% wealth share correlates with the growth rates of the three main components of the synthetic savings decomposition. In the first figure, the growth rate of the savings rate tracks the evolution of top wealth shares fairly closely, with increases in the 1990s, a decline in the 2000s and then an upswing in recent years. The capital gains rate on the other hand appears to be negatively correlated with trends in the top wealth share. The capital gains rate was declining when the top 1% share was increasing the most and increasing when the top 1% share was falling. Income growth also appears to be somewhat uncorrelated with trends in the top 1% share.

To explain why the uptick in capital gains from the mid-2000s to mid-2010s did not result in a larger increase in the top 1% wealth share, it is useful to consider how capital gains affect the wealth share. High capital gains will only increase the wealth *share* of a group if their capital gains are substantially larger than the gains of other groups. Although the top 1% experienced high capital gains following the Great Recession, largely driven by a strong equity market, the bottom 99% also had strong capital gains, owing to a robust housing market. As a result, the top 1% wealth share did not increase as a result of these strong capital gains.

A final exercise to explore the role of savings and capital gains in driving the top 1%

¹⁵For details on how the financial flows for the categories are measured see App. Section A.3.

¹⁶For details on how the synthetic savings and capital gains rates are computed see App. Section B.2.2. ¹⁷I use the LOWESS values since the year-to-year values can fluctuate dramatically in some cases (as seen in Appendix Figure 8), which distorts the broader trends.

 $^{^{18}}$ As discussed in Appendix Section A.3, the average savings rate baseline in this exercise is higher than the one reported in the capital account and is around 12%, which lines up with these results.



(c) Income Share

Figure 10: Comparison of Top 1% Wealth Share Growth Rate with Key Factors This figure plots the growth rate of the top 1% share over time in comparison to the growth rate of the three other factors in the synthetic savings decomposition: savings rate (left), capital gains (centre) and income (right). The trends are the LOWESS smoothed values during this period.

wealth share is to simulate what would have happened to the top 1% share without the dynamics of each component. Starting from the initial period, I simulate the changes in the wealth share that would have occurred if each component in turn was fixed at its average for the period. I present the results of this simulation in Figure 11. I find that if the savings rate dynamics were ignored and remained constant throughout this period, then the top 1% share would display a steady upward trend throughout the entire period, reaching over 18% in 2018. Changing the capital gains and income inequality values would not have as significant an effect.¹⁹

An important distinction to draw is that, although this section shows that capital gains

 $^{^{19}}$ In Appendix Figure 10, I show how the top 1% share would evolve had Canada had the trends of each component that were observed in the United States.



Figure 11: Simulations of the Top 1% Wealth Share Using Component Averages This figure shows the simulation results for the top 1% share of replacing the computed values for annual savings rates, capital gains rates and income inequality with their averages for the period. Each simulation is run separately for each component.

play a small role in explaining the dynamics of the top 1% wealth share, they account for much of the aggregate increase in wealth during this time. In Appendix Figure 7, I show that capital gains account for the majority of the cumulative increase in wealth over the last three decades, especially for housing and equity wealth. This helps square the fact that savings rates have been declining over recent years, but wealth has been increasing. The concern, however, is that much of this huge increase is wealth is largely a paper gain and does not reflect an increase in actual investment. As a result, future increases in wealth may be limited unless Canadians increase investment or strong capital gains continue.

These results show how the dynamics in the top 1%'s savings rate account for many of the dynamics of the top 1% wealth share in Canada. The top 1% exhibits a much higher degree of volatility in their savings behaviour and it appears to be highly correlated with their share of wealth. This has some interesting implications for thinking about the trends in wealth inequality in Canada. Although wealth inequality may not be increasing as rapidly compared to previous decades and other countries, this is largely driven by a decline in savings and particularly business investment, which could be concerning for innovation and economic growth moving forward.

6 Conclusion

This paper fills a major gap in our understanding of wealth inequality in Canada. Using the capitalization method from Saez & Zucman (2016), I estimate the share of wealth going to the top 1% in Canada from 1990-2018. I find that although wealth inequality did rise during this time, the gains were modest and inequality actually fell over the last decade of the period in question. These estimates suggest wealth is far more equally distributed in Canada compared to the United States and is even slightly more equal than France. Canada's relatively low level of wealth inequality is due in part to the large role played by housing and pensions in Canada's wealth portfolio and the relative lack of concentration across all asset classes. The modest increase in wealth inequality over this period is best explained by declining savings from the top 1% and strong capital gains for the bottom 99%. These findings raise some interesting questions about how inequality, investment and economic growth are connected.

These results also raise important questions about how the public perceives inequality. These results show that Canada is performing much better than many of its peer countries when it comes to wealth inequality. However, it is common to encounter those in Canada with major concerns about rising inequality and the role of billionaires in society. This has led to calls for price controls on goods sold by major corporations and wealth taxes on billionaires. Future work is needed to better understand the link between the public perception of inequality and actual measures of it. With this paper, the hope is to contribute to this important discussion.

References

- Acciari, P., Alvaredo, F., & Morelli, S. (2024). The Concentration of Personal Wealth in Italy 1995–2016. *Journal of the European Economic Association*, 22(3), 1228–1274.
- Albers, T., Bartels, C., & Schularick, M. (2022). Wealth and its Distribution in Germany, 1895-2018. SSRN Electronic Journal.
- Alstadsæter, A., Johannesen, N., & Zucman, G. (2018). Who owns the wealth in tax havens? Macro evidence and implications for global inequality. *Journal of Public Economics*, 162, 89–100.
- Alvaredo, F., Atkinson, A. B., & Morelli, S. (2018). Top Wealth Shares in the UK Over More Than a Century. *Journal of Public Economics*, 162, 26–47.
- Bach, L., Calvet, L. E., & Sodini, P. (2020). Rich Pickings? Risk, Return, and Skill in Household Wealth. American Economic Review, 110(9), 2703–2747.
- Benhabib, J., Bisin, A., & Luo, M. (2019). Wealth Distribution and Social Mobility in the US: A Quantitative Approach. American Economic Review, 109(5), 1623–1647.
- Benhabib, J., Bisin, A., & Zhu, S. (2011). The Distribution of Wealth and Fiscal Policy in Economies With Finitely Lived Agents. *Econometrica*, 79(1), 123–157.
- Blanchet, T. & Martínez-Toledano, C. (2023). Wealth Inequality Dynamics in Europe and the United States: Understanding the Determinants. *Journal of Monetary Economics*, 133, 25–43.
- Brzozowski, M., Gervais, M., Klein, P., & Suzuki, M. (2010). Consumption, Income, and Wealth Inequality in Canada. *Review of Economic Dynamics*, 13(1), 52–75.
- Chernozhukov, V., Fernandez-Val, I., & Melly, B. (2020). QRPROCESS: Stata module for quantile regression: fast algorithm, pointwise and uniform inference. Statistical Software Components, Boston College Department of Economics.
- Davies, J. B. & Di Matteo, L. (2020). Long Run Canadian Wealth Inequality in International Context. *Review of Income and Wealth*, (pp. roiw.12453).
- De Nardi, M. & Fella, G. (2017). Saving and Wealth Inequality. Review of Economic Dynamics, 26, 280–300.
- Fagereng, A., Guiso, L., Malacrino, D., & Pistaferri, L. (2020). Heterogeneity and Persistence in Returns to Wealth. *Econometrica*, 88(1), 115–170.
- Garbinti, B., Goupille-Lebret, J., & Piketty, T. (2020). Accounting for Wealth-Inequality Dynamics: Methods, Estimates, and Simulations for France. *Journal of the European Economic Association*.
- Hubmer, J., Krusell, P., & Smith Jr., A. A. (2020). Sources of U.S. Wealth Inequality: Past, Present and Future. *NBER Macroeconomics Annual*, 35.

- Jakobsen, K., Jakobsen, K., Kleven, H., & Zucman, G. (2020). Wealth Taxation and Wealth Accumulation: Theory and Evidence From Denmark*. The Quarterly Journal of Economics, 135(1), 329–388.
- Katic, P. & Leigh, A. (2016). Top Wealth Shares in Australia 1915–2012. Review of Income and Wealth, 62(2), 209–222.
- Kopczuk, W. & Saez, E. (2004). Top Wealth Shares in the United States, 1916-2000: Evidence from Estate Tax Returns. National Tax Journal, 57(2, Part 2), 445–487.
- Krusell, P. & Smith, A. A. (1998). Income and Wealth Heterogeneity in the Macroeconomy. https://doi.org/10.1086/250034, 106(5), 867–896.
- Kuhn, M., Schularick, M., & Steins, U. I. (2020). Income and Wealth Inequality in America, 1949–2016. Journal of Political Economy, 128(9), 3469–3519.
- Martínez-Toledano, C. (2019). House Price Cycles, Wealth Inequality and Portfolio Reshuffling. *Working Paper*, (pp. 1–143).
- Piketty, T. & Zucman, G. (2014). Capital is Back: Wealth-Income Ratios in Rich Countries 1700-2010. Quarterly Journal of Economics, 129(3), 1255–1310.
- Saez, E. & Veall, M. R. (2005). The Evolution of High Incomes in Northern America: Lessons from Canadian Evidence. American Economic Review, 95(3), 831–849.
- Saez, E. & Zucman, G. (2016). Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data *. The Quarterly Journal of Economics, 131(2), 519– 578.
- Saez, E. & Zucman, G. (2020). Trends in US Income and Wealth Inequality: Revising After the Revisionists. NBER Working Paper.
- Smith, M., Zidar, O., & Zwick, E. (2023). Top Wealth in America: New Estimates Under Heterogeneous Returns^{*}. The Quarterly Journal of Economics, 138(1), 515–573.
- United Nations (2010). System of National Accounts. United Nations.
- Vermeulen, P. (2018). How Fat is the Top Tail of the Wealth Distribution? Review of Income and Wealth, 64(2), 357–387.
- Wodrich, N. & Worswick, A. (2020). Estimating the Top Tail of the Family Wealth Distribution in Canada. Technical report, Parliamentary Budgetary Officer.

A Appendix: Data

This section describes in detail how wealth in Canada is measured along several dimensions. First, I discuss exactly how aggregate wealth is defined in this paper both conceptually and statistically. Second, I explain how I divide wealth up into different asset classes, which is required for the capitalization method and for understanding overall trends. Finally, I explain how I use the data on financial flows to divide changes in wealth between changes in savings and capital gains to use in the synthetic savings decompositions.

A.1 Aggregate Wealth in Canada

The first step in measuring the distribution of wealth is determining the definition of wealth both conceptually and statistically. In this paper, wealth is defined as marketable wealth, which are assets that are "subject to ownership rights and from which economic benefits may be derived by their owners by holding them" (United Nations, 2010). This definition encompasses assets such as currency and deposits, bonds, equities and real estate, which all are bought and sold on open markets and provide economic benefits from holding them. Notable assets that are not counted in this definition of wealth are promises of future government spending (eg. social security), unfunded pensions and human capital. Many of these assets cannot be sold on an open market and therefore are not marketable wealth.

This definition of wealth aligns closely, but not exactly, with the main source of wealth data in Canada: the National Balance Sheet Accounts (NBSAs). The following paragraphs will explain the key differences between this paper's definition and the NBSAs. In App. Table 1, I present an excerpt of the NBSAs from the first quarter of 2018 to illustrate the different categories that are included. The first difference between the NBSAs and this paper is that this paper does not consider consumer durables as wealth, which is in keeping with the work of Saez & Zucman (2016) and the guidance of the System of National Accounts (SNA) (United Nations, 2010), while the NBSAs do. Correspondingly, I also do not include consumer credit, which is largely used to pay for consumer durables like cars, as a liability.

A second difference is that I use a different definition of pension wealth than the NBSAs that comes from the Pension Satellite Accounts (PSAs). In App. Table 2, I present an excerpt from the PSAs for 2018 highlighting the main categories. Pension wealth in this paper is measured as trusteed, employer pension plans and individual registered retirement savings plans (RRSPs),²⁰ which are documented in the PSAs. This is done to ensure that categories of retirement saving such as social security and unfunded pensions are not included. The NBSAs, by contrast, simply have a "Life Insurance and Pensions" category

 $^{^{20}}$ RRSPs are similar to 401(k)'s in the United States.

Category	Value (\$ M)	% NW	Assigned Asset Class
Total Assets	13,388,991		
Non-Financial Assets	6,393,924		
Residential Structures	2,356,701	21.2	Principal Residences (80%), Secondary Properties (20%)
Non-Residential Structures	60,751	0.5	Unincorporated Business
Machinery And Equipment	28,821	0.3	Unincorporated Business
Intellectual Property Products	3,291	0.0	Unincorporated Business
Consumer Durables	670,674	6.0	Unassigned
Inventories	19,187	0.2	Unincorporated Business
Land	3,254,499	29.3	Principal Residences (80%), Secondary Properties (20%)
Total Financial Assets	6,995,067		
Total Currency And Deposits	1,418,469	12.8	Other Investments (67%), Pensions (33%)
Debt Securities	137,111	1.2	Other Investments (67%), Pensions (33%)
Loans	1,653	0.0	Unassigned
Listed Shares	470,824	4.2	Equities (67%) , Pensions (33%)
Unlisted Shares	643,551	5.8	Equities
Mutual Fund Shares (Units)	1,422,123	12.8	Equities (16%), Other Investments (45%), Pensions (31%), Unassigned (8%)
Foreign Investments: Equity	154,954	1.4	Other Investments (67%), Pensions (33%)
Life Insurance And Pensions	2,589,507	23.3	Replaced by Pension Satellite Account
Other Accounts Receivable	156,875	1.4	Unincorporated Business
Total Financial Liabilities	2,264,710		
Consumer Credit	630,178	-5.7	Unassigned
Non-Mortgage Loans	127,022	-1.1	Unincorporated Business
Mortgages	1,450,398	-13.0	Principal Residences (80%), Secondary
			Properties (20%)
Other Accounts Payable	57,112	-0.5	Unassigned
Net Worth	11,124,281		

Table 1: National Balance Sheet Account, Household Sector, 2018 Q1

This table shows the values found in the National Balance Sheet Accounts for the Household and non-profit institutions serving households sector in the first quarter of 2018. The values are expressed in millions of dollars and each category share is in terms of net worth. I include the asset type that each category in the NBSA is assigned to in the paper. Since the definition of wealth differs slightly in the NBSA from this paper, the total net worth values will also be slightly different as well.

Category	Value (\$ M)	% NW	Assigned Asset Class
Total Plans	3,851,538		
Social Security	443,338	11.5	Unassigned
Employer-Based Pension Plans	$2,\!208,\!030$		
Trusteed Pension Plans	$1,\!818,\!275$	47.2	Pensions
Gov. Consolidated Revenue Arrangements	272,674	7.1	Unassigned
Other Employer-Based Pension Plans	117,081	3.0	Unassigned
Individual Registered Saving Plans (RSP)	1,200,170	31.2	Pensions

Table 2: Pension Satellite Account, 2018

This table shows the values found in the Pension Satellite Account (PSA) for 2018. The PSA provides a more detailed breakdown of assets in Canada's pension system. The values are expressed in millions of dollars and the shares represent the share of total pension wealth. I include which asset class each category is assigned to, where unassigned means it is not included in the definition of wealth in this paper.

that does not provide additional detail. This choice does have the consequence of omitting life insurance policies from the definition of wealth despite being included in other papers on wealth inequality.²¹

The use of the PSAs requires some adjustments of the NBSAs to ensure that double counting does not occur. This stems from the fact that RRSPs, which are retirement investment vehicles, are also counted in the NBSAs as their component parts (ie. equities, deposits). To address this, I reduce each asset that could be included in RRSPs (eg. equities, currency and deposits, mutual funds) by the share that RRSPs take up of those assets combined. For example, in 2018, 33% of these assets are assigned to RRSPs. This assumes that RRSPs are comprised of the same portfolio composition as the NBSAs themselves, which is supported by an annual survey of RRSPs by a major Canadian bank. To see this, mutual funds make up 39.5% of these assets in the NBSAs and 42% in the survey, listed shares and foreign equity make up 17.3% in the NBSAs and 34% in the survey (including ETFs) and currency and deposits are 39.5% in the NBSAs and 34% in the survey (cash and GICs).

A final difference is that some inconsequential categories that do not fit into an obvious asset class are omitted. This includes loans held by households (\$1.6 billion), other accounts payable (\$57 billion) and the value of these assets held by mutual funds. The mutual funds part comes from the fact that I distribute mutual funds according to their holdings into different asset classes and the assets held in these categories are therefore also omitted.

I explore whether these choices around the definition of wealth matter by estimating the top 1% share of wealth had alternative definitions been used. I do this using the linear asset decompositions discussed in Section 5, where I change the aggregate asset shares to

 $^{^{21}}$ The omission of consumer credit and life insurance are the two main differences between this paper and Saez & Zucman (2016) in terms of the definition of wealth.



(a) Individual Changes

(b) Combinations of Changes

Figure 1: Top 1% Wealth Share in Canada Using Different Wealth Definitions These figures show how sensitive the top 1% wealth share is to using different definitions of wealth. I compute the shares using the linear decomposition methods detailed in Section 5. For a greater discussion of the different definitions of wealth, refer to Appendix Section A.1.

reflect different wealth definitions and insert assumed wealth shares as needed. The first alternative is if the NBSA definition of pension wealth is used instead of trusteed pension plans. The second alternative is if consumer credit was included as a liability. I assume that the distribution of consumer credit is similar to what is observed in the SFS. Lastly, I also explore how including an estimate of offshore wealth would affect the top 1% share. Alstadsæter et al. (2018) provides estimates of the share of offshore wealth attributable to each country as a share of GDP. For Canada, they find this number to be 4.64%, which multiplied by Canadian GDP amounts to about \$32 billion in 1990 and \$104 billion in 2018. I assume that 90% of this offshore wealth is distributed to the top 1%.

I present the results of this exercise in App. Figure 1 both with each alternative individually and with different combinations of alternatives. The first alternative, where the definition of pensions in the NBSA is used, reduces the top 1% share by 1-2 p.p. because pensions are fairly equally distributed and this would increase the amount of pension wealth as a share of all wealth. In the second alternative where consumer credit is included, the top 1% share increases by around 1-2 p.p. since consumer debt largely reduces the wealth of those at the bottom of the distribution. Lastly, including an estimate for offshore wealth increases the top 1% share by around 1-2 p.p., which is similar to including consumer debt. Since the pension alternative and consumer credit alternative largely cancel out, the trends using these different approaches are very similar to the main results presented in this paper. Including offshore wealth would increase the estimates by about 1-2 p.p.

A.2 Aggregate Wealth by Asset Class

While the previous section detailed what would be included in the definition of wealth, it is also important to explain how asset categories were determined as this is a crucial detail in how the capitalization method works. In App. Tables 1-3, I detail the different asset classes used in this paper and the corresponding categories from the NBSA that they are comprised of. The determination of asset classes is largely driven by the capital income streams that are available in the Longitudinal Administrative Databank (LAD).

The first major category is that of Canadian equity, which is comprised of listed shares, unlisted shares and mutual fund assets. Although breaking this down further into listed and unlisted shares would provide further detail and accuracy, this information in the LAD only goes back to 2006. Mutual funds are more complex since mutual funds are comprised of a variety of different assets and mutual fund distributions are observed in the LAD only as the component parts. For example if a mutual fund is comprised of equities and bonds, the LAD simply captures the dividends and capital gains from the equities and the interest payments from the bonds rather than as a payout from mutual funds. For this reason, the mutual fund category is split into the Canadian equity and other investment categories according to the share of each asset should also be distributed accordingly. To do this, it is useful to note that the total sum of an asset class (say, equities) held by households through mutual funds takes the form of an infinite geometric sequence.

$$\begin{split} EQ^{M} &= \underbrace{sh_{EQ}^{M}MF^{H}}_{\% \text{ EQ in MF}} + \underbrace{sh_{EQ}^{M}sh_{M}^{M}MF^{H}}_{\% \text{ EQ of \% MF in MF}} + \underbrace{sh_{EQ}^{M}sh_{M}^{M}MF^{H}}_{\% \text{ EQ of \% MF of \% MF in MF}} + \dots \\ EQ^{M} &= sh_{EQ}^{M} \left(MF^{H} + sh_{M}^{M}MF^{H} + sh_{M}^{M^{2}}MF^{H} + \dots\right) \\ EQ^{M} &= sh_{EQ}^{M} \left(\frac{MF^{H}}{1 - sh_{M}^{M}}\right) \\ EQ^{M} &= \frac{EQ^{M}}{\mathcal{P}\mathcal{O}\mathcal{T}^{M}} \frac{MF^{H}\mathcal{P}\mathcal{O}\mathcal{T}^{M}}{TOT^{M} - MF^{M}} \\ EQ^{M} &= \frac{EQ^{M}MF^{H}}{TOT^{M} - MF^{M}} \end{split}$$

After distributing the mutual funds to their corresponding asset classes, I remove the value of assets assumed to be held in RRSPs to avoid double counting with the PSA (RRSPs can include mutual funds). For example, for the equities asset class, I multiply the number of listed and unlisted shares plus the value of mutual funds assigned to equities by the fraction of assets not assigned to RRSPs, $1 - \frac{RSP}{EQ+OI}$. In 2018, this is about two-thirds of the assets

Asset Class	Value (\$ M)	% Total	NBSA Categories
Canadian Equity	1,190,100	11.6	Listed Shares, Unlisted Shares, Mutual Funds
Other Investments	1,785,291	17.3	Currency and Deposits, Debt Securities, Mutual Funds, Foreign Equity
Unincorporated Business	141,903	1.4	Non-Residential Structures, Machinery and Equip- ment, Intellectual Property, Inventories, Other Ac- counts Receivable, Non-Mortgage Loans
Pensions	3,018,445	29.3	Trusteed Pension Plans, Individual RRSPs
Principal Residences	3,328,642	32.3	Residential Structures, Land, Mortgages
Secondary Properties	832,160	8.1	Residential Structures, Land, Mortgages
Total	10,296,541		

Table 3: Total Marketable Wealth by Asset Class, 2018

This table shows the different aggregate asset categories used in this paper, their value in millions for 2018, the share of the total and the NBSA categories that are used to create them. These values reflect the net asset value of each category after subtracting liabilities such as mortgages.

that are held in RRSPs. This also implies that the resulting distribution of the mutual funds category is 16% equities, 45% other investment assets, 31% pensions and 8% unassigned categories.

The next asset class is called "Other Investment" because it refers to all other assets that pay interest and investment income that are not Canadian equities. This includes currency and deposits, debt securities (bonds) and foreign equity. Foreign equity is included because in the LAD it is grouped in with the "Interest and Other Investment" category. For mutual funds holdings of these assets, a similar exercise is conducted as for Canadian equity.

The remaining assets are more straightforward or have been discussed previously. Pensions are measured in the Pension Satellite Accounts as trusteed pension plans and RRSPs. Real estate assets include residential structures, land and mortgages. I split real estate assets into principal and secondary residences by noting that about 80% of real estate assets in the SFS are primary residences. This allows the creation of two categories of real estate and for the rental income variable in the LAD to have a corresponding asset class. Unincorporated business assets is a catch all for the remaining assets that are generally held by self-employed individuals. This used to be its own sector of the NBSAs, but was discontinued. In Canada, self-employment rates have been falling in favour of incorporation so these assets play a very small role in recent years.

A.3 Financial Flows in Canada

In Section 5, I decompose the change in wealth from year to year into different components - capital gains, savings and labour income - for each wealth group. To do this requires an estimate of the average capital gain by asset class, which then can be used to back out the implied savings. I measure capital gains using of the Financial Flow Account (FFA) in the Canadian System of Macroeconomic Accounts. The FFA breaks down the changes in aggregate wealth by category into a portion explained by investment and a portion explained by other changes, which typically means capital gains as part of the revaluation account.

It is briefly useful to note what the FFA represents and its connection to the capital account, which includes a measure of the household savings rate that is widely reported. The FFA is the financial counterpart to the capital account that breaks down the net financial transactions by the same asset categories as found in the NBSAs. The closing line in both the FFA and the capital account is net lending or borrowing, although they are reached using different categories and have some slight discrepancies due to differences in measurement.

To see how these measures compare, I show the household savings rate from the two accounts using two separate measures in App. Figure 2.²² From the capital account, I include household net saving and net lending or borrowing, which includes the depreciation of capital assets and the acquisition of new, non-financial capital assets. From the FFA, I include the net transactions in financial assets (not liabilities) and net financial investment, which includes liabilities (such as mortgages). The figure shows how the two sets of measures from the different accounts are roughly aligned. Since this paper uses a definition of wealth that includes liabilities, the net lending or net financial investment measure is a better baseline to consider for the synthetic savings analysis, but it is also worth noting that this measure is negative for much of the period. This means that the household sector is largely a net borrower from other sectors of the economy using this definition.

The negative value for net lending, however, can be explained by a few factors that I adjust for in the synthetic savings analysis. First, this net financial investment does not account for investment in non-financial assets, which are included in the NBSAs, such as housing and unincorporated business assets. I include investment in housing using the Housing Economic Account (HEA), which includes a measure of investment for residential structures over time. I subtract an estimate of housing depreciation of 1.5% of the residential structures stock, which is taken from the Canadian CPI calculation, from the housing investment value. The capital gain for the real estate assets is then the residual of the change in real estate value over time minus the level of net investment in housing calculated from the HEA. I show in

 $^{^{22}}$ I divide the savings measures by primary household income rather than disposable income to be consistent with the use of market income for determining income shares of each wealth group.



(a) Capital Account vs. FFA

(b) FFA Adjustments

Figure 2: Household Savings Rates from Different Accounts The figure on the left shows the household savings rate across different macroeconomic accounts and using different measures of savings. From the capital account, this includes the net household savings and net lending. From the Financial Flow Account, this includes net transactions in financial assets, which omits liabilities and net financial investment which includes them. The figure on the right shows the household savings rate after making several adjustments to be more consistent with the rest of this paper. This includes omitting certain asset categories and using the Pension Satellite Account for pension assets, including non-financial investment such as housing and including retained earnings from corporations as savings. Net savings are divided by primary household income rather than disposable income.

App. Figure 3, how the estimated capital gains for both listed equities, which are measured directly in the FFA, and housing, which is measured this way, track closely with the Toronto Stock Exchange (TSX) and the Canadian Housing Price Index (HPI) over this period. I include investment in unincorporated business assets by just assuming that half of the change in value of these assets is considered a capital gain and half is investment. Since these assets make up only a small share of all assets, they are unlikely to affect the results.

A second adjustment relates to the inclusion of corporate net savings. Corporations can choose to save money from year to year and not distribute dividends to shareholders. However, since households are the shareholders, they are effectively choosing to save this money through the corporation rather than to distribute and reinvest the income elsewhere. In Canada specifically, where many professionals, such as doctors and lawyers, choose to self-incorporate to take advantage of lower tax rates, including corporate net savings as household savings is important. I assume that half of corporate savings is income that could have been distributed and then saved in a way that would register as savings in the FFA. This also creates a more realistic measure of investment in equities as the cumulative amount of investment in equities when not accounting for corporate net savings is negative as seen in App. Figure 4. To ensure that I am counting corporate net savings that accrue to Canadians only, I take the same share of corporate net savings as the share of dividend



Figure 3: Capital Gains from FFA and HEA versus Major Indicators These figures compare the implied capital gains calculated from the Financial Flow Accounts (equities) and Housing Economic Account (housing) to major indicators: the Toronto Stock Exchange (TSX) and the housing price index (HPI).

payouts to Canadians from corporations.

Finally, I make some further adjustments to the FFA values to be more consistent with the analysis in the paper. Since the financial flows data is reported quarterly and I use the first quarter of the year for asset values, the flows represent the cumulative total from the second quarter of the prior year up the the first quarter of the given year. I also use the PSA financial flows rather than those reported in the FFA and omit consumer credit. Last, due to revisions made to the NBSAs and FFA and the fact that the estimates from the LAD are based on the old version of the NBSA, while only the new version of the FFA is available, there is a slight adjustment made to these values. Since the values for financial flows and other changes will not add up to the total change, I calculate the share of the change in asset value observed in the revised FFA that comes from capital gains and assign this to the observed change in the old NBSAs. In some rare cases when the shares are greater than 1 (or less than 0), such as if one of the components is negative, I assign a value equal to the total change in the NBSA plus the negative difference if more than 1 or the negative amount if less than 0.²³

In App. Figure 2, I present different versions of the aggregate savings rate that incorporate the various adjustments discussed above. Adding non-financial investment, mainly in the form of net housing investment, makes a substantial difference to the savings rate, bringing it up above the rate reported in the capital account. Including corporate savings

²³Simply using the ratio led to some extremely large outlier values in cases where the ratio was large and the old NBSA change in values were much larger than the ones in the FFA.



Figure 4: Equity Asset Growth With and Without Corporate Savings This figure shows the cumulative change in equity wealth over time broken down into savings and capital gains. The figure on the left shows the case where half of corporate savings accruing to Canadians is considered savings, while this is not included in the figure on the right side. Corporate net savings data comes from the Statistics Canada Undistributed Corporation Profits table.

also increases the savings rate marginally. Together, these adjustments result in an average savings rate around 12% that holds fairly steady over the period unlike the reported savings rate, which declines precipitously during this time. These numbers serve as a useful baseline for thinking about savings rates in the synthetic savings rate decompositions.

B Appendix: Methodology

This section provides some additional detail on the methods used in this paper. The first part discusses the capitalization method and the approaches used to fill in some of the gaps related to assets without accurate capital income flows. The second part provides greater detail on the two sets of decompositions used in this paper: the linear asset decompositions and the synthetic savings decompositions.

B.1 Capitalization Method

B.1.1 Details of Capitalization Method

The capitalization method is based on the intuition that if you know the rate of return of an asset, you can estimate the amount of wealth a person holds based solely on the capital income flows they report on their income tax forms. The first step in this process involves estimating these rates of return. The simplifying assumption is that the rates of return will be the same for people across wealth groups. With this assumption, the rates of return can be estimated using aggregate estimates of wealth in the NBSAs and aggregates of the capital income flows in the NBSAs. In Table 4, I outline the asset classes that will be used for the capitalization method and the corresponding asset categories from the NBSAs and LAD that correspond to those asset classes.

The Canadian Equity category captures listed and unlisted shares in the NBSAs plus the value of these assets held in mutual funds, minus the amount held in RRSPs. The corresponding capital income flows in the LAD are eligible and non-eligible Canadian dividends as well as capital gains. Eligible dividends generally correspond to dividends issued by larger corporations to shareholders that are eligible for a larger dividend tax credit because the corporations pay a higher rate of corporate tax. Non-eligible dividends generally correspond to dividends issued by small Canadian controlled private corporations (CCPCs) that already benefited from lower corporate tax rates. Capital gains are a trickier inclusion because, although they do reflect ownership of shares, they are realized in a very lumpy pattern, which can lead to greater variability in share ownership. In my main results, I include capital gains as a capital income flow, but I also present the results without capital gains as a robustness check in Figure 5, which shows that the overall trend is relatively similar.

The Other Investments category is fairly broad and includes currency and deposits, debt securities and foreign equity plus the value of these assets held in mutual funds, minus the amount held in RRSPs. The corresponding capital income flow in the LAD is interest and other investment income. This is a broad category that combines any interest payments

Categories	NBSA Variables	LAD Variables
Canadian Equity	Listed Shares Unlisted Shares	Eligible Canadian Dividends Non-Eligible Canadian Dividends Capital Gains
Other Investments	Currency and Deposits Debt Securities (Bonds) Foreign Equity	Interest and Other Investment Income
Unincorporated Business	Non-Residential Property Machinery Inventories Intellectual Property Other Receivables (Minus) Non-Mortgage Loans	Self-Employment Income
Pensions	Registered Pension Plans Registered Retirement Savings Plans	No Direct Capital Income Flow
Primary Residences	Residential Structures Land	No Direct Capital Income Flow
Other Real Estate	(Minus) Mortgages	Net Rental Income

Table 4: Categorization of Asset Classes in NBSA and LAD

This table shows the different asset categories used in the capitalization method in both the National Balance Sheet Accounts (NBSAs) and the Longitudinal Administrative Databank (LAD).

received with any other payments from assets that are taxed at 100%. This includes income from foreign equity, which is a somewhat awkward inclusion in this category, but also reflects a fairly small share of overall wealth. In 2018, this represented just 1.4% of aggregate net worth in Canada.

Unincorporated business assets are defined narrowly in this exercise because the NBSAs do not separate non-corporate business activity from the household sector. As a result, this category considers capital assets owned by households for the purposes of operating a business. This includes machinery, non-residential structures and intellectual property and subtracts non-mortgage loans, but does not include assets such as bank deposits, land or equities since these are just attributed to the household itself. The capital income flow used is self-employment income, which should capture the income received from these small business ventures.

As mentioned in Section 3.2, there are no capital income flows that correspond to pensions or primary residences and so these values are imputed. Since the LAD does contain information on net rental income, I use this in conjunction with the Secondary Residences category to capture wealth form owning additional properties and renting them out.

B.1.2 Distribution Regression

To impute the values of housing and pension wealth in the microdata, I use the distribution regression techniques from Chernozhukov et al. (2020). To start, I compute 100 quantiles of housing or pension wealth in the full sample of the SFS. Then, for every quantile, I estimate a logit regression of whether an individual is above or below that threshold. For the logit regression, I model net housing prices (value of principal residence - mortgage) as a linear function of market income, age, age-squared, family indicators and city fixed-effects for the 15 largest cities in Canada. For pensions, I compute the family pension value per person and model it as a quadratic function of average age, and a linear function of income, pension contributions and pension income (both employer pensions and RRSPs).

Using the estimated coefficients, I can predict the probability of a given household of being above or below each quantile threshold in the LAD using the common covariates.²⁴ I then repeat this procedure for all 100 quantiles. This will yield a series of probabilities at a hundred different points. Using linear interpolation, you can compute the conditional quantile values at each integer quantile for each individual. That is, suppose a household in the LAD has a 35% predicted probability of being above \$420k in housing wealth and a 37% probability of being above \$400k, where the \$420k and \$400k were consecutive quantiles of housing wealth used in the initial logit regression. Then, by linear interpolation, there is a 36% probability of falling above \$410k. Therefore, the 63rd percentile outcome for the household in the LAD is \$400k, the 64th percentile outcome is \$410k and the 65th percentile outcome is \$420k. If the value drawn from the uniform distribution, $p \sim U[0, 1]$, is 0.64, then the value of housing wealth assigned to the household is \$410k.

This method does a good job replicating the distribution of housing in the LAD that is found in the SFS, especially compared to some more rudimentary alternatives. App. Table 5 reports the share of housing wealth going to the top 1% by total wealth across two different imputation methods and compares the results to the share observed in the SFS. In the first column, housing values are assigned based on the average housing value in the census tract of residence obtained from the Canadian Census Profiles. In Canada, a census tract comprises 2,500 to 8,000 people and there are around 5,000 of them across the country. To deal with the fact that not everyone in a census tract owns a home, the share of home ownership is measured for each census tract and families are sorted based on wealth within a census tract. Then, the share of home ownership is used as the percentile cutoff for assigning the average home value. That is, if the census tract home ownership rate is 65%, then those

 $^{^{24}}$ I use SFS estimates from the nearest year to the LAD year in question. That is, I use the 1999 SFS estimates for all LAD years prior to 2006. I then re-scale housing and pension values so that they add up to the aggregate wealth value in the NBSAs for that year, which functions as a quasi-inflation adjustment.

	Variations				
	Census Profiles	SFS Imputation	SFS Values		
1999	2.90	4.80	5.59		
2012	2.80	4.50	5.63		
2016	3.20	4.80	6.66		

Table 5: Comparing Top 1% Share of Housing Across Imputation Approaches

This table compares tow different approaches to imputing housing wealth to the Survey of Financial Security (SFS) true values (column 3) by comparing the share of total housing wealth owned by the wealthiest 1%. The first method involves assigning the average housing value of a census tract to those that live there (column 1). The second employs the distribution regression approach described in Section 3.2 (column 2). The distribution regression approach generates housing estimates at the top much closer to what is observed in the actual SFS.

above the 35th percentile of non-housing wealth in the census tract are assigned the average census tract value and those below are assigned a value of zero. While this approach can get geographically specific values, the data itself does not capture mortgages (which are larger at the bottom of the distribution) and the method of assigning ownership as well as the lack of variation in imputed values likely leads to an underestimate of top share housing wealth.

The second column, the distribution regression approach described above, does a much better job of matching the true SFS values for housing (reported in column 3). There are a couple reasons this approach is useful compared to more non-parametric approaches such as in column 1 or as employed by Garbinti et al. (2020).²⁵ First, the method does a good job dealing with zero values, which is often a challenge in a linear regression model or even tobit and hurdle models. This is because the conditional quantile function can yield many predicted quantiles of 0 for families, giving a high likelihood of drawing a zero value for some families. Second, in non-parametric approaches that use bins, there is a concern of no within-group variation. This will reduce the level of variation and inequality generated. Last, the model provides the opportunity to include several covariates to increase predictive power - including city fixed-effects. Better prediction may not necessarily change the overall distribution, but does raise confidence in the approach. In addition, the parametric assumptions made for housing values are fairly reasonable: housing wealth is linear in market income and quadratic in age.

 $^{^{25}}$ In Garbinti et al. (2020), they employ a non-parametric approach where people are assigned to around 200 bins based on income and age. Then using the share of people holding an asset in each bin, they randomly draw to see if the average value should be assigned. In their refined method, they do allow values in the bin to vary according to the observed distribution as well.

B.2 Decompositions

B.2.1 Linear Asset Decompositions

In App. Figure 5, I present the portfolio of assets held by different wealth groups over time. There are a few interesting observations. First, the importance of pensions and principal residences for the wealth of those not in the top 1% cannot be overstated. For the bottom 75%, principal residences and pensions make up 84% of their wealth in 2018, while it is just 16% of the portfolio of the top 1%. The reverse is apparent for Canadian equities. For the bottom 75% and 75th-90th percentiles, equities make up only 3% and 5% of their portfolios in 2018 respectively, compared to 33% for the top 1%.



Figure 5: Portfolios of Wealth Groups Over Time

This figure presents the portfolios of different wealth groups over time using wealth calculated with the capitalization method. Each bar's height corresponds to the share of overall wealth held by that wealth group in a given year. Each bar is divided into the components of wealth for each wealth group.

In App. Figure 6, I plot the aggregate and within-asset shares for the top 1% for a number of asset classes. There are some notable patterns that emerge. First, Canada has a low aggregate share of business assets compared to France and the USA, while the USA has a lower aggregate share of housing wealth and France has a low share of pension wealth. It should be noted that this low share of pension wealth in France should be expected given the definition of wealth used in these papers. France has a much higher pension replacement

rate - the government pension entitlement received by someone with average pre-retirement income net of taxes as a share of pre-retirement income - than Canada. France has a 70% replacement rate, while Canada only has a 50% replacement rate. (OECD, 2019). This could result in less pension saving through employer-pension plans or individual retirement vehicles like an RRSP. Second, Canadian assets are less concentrated among the top 1% for all asset classes compared to the US.



(b) Within-Asset Shares By Asset

Figure 6: Comparison of Aggregate and Within-Asset Shares Across Countries This figure plots the aggregate wealth share and within-asset shares of each asset in Canada, the United States (Saez & Zucman, 2016) and France (Garbinti et al., 2020) between 1990 and 2020. Business assets includes equities and unincorporated business wealth.



Figure 7: Cumulative Capital Gains and Investment Flows by Asset Category, 1990-2018 This figure shows the cumulative increase in capital gains and investment flows by asset category from 1990-2018. The breakdown is based off data from the Financial Flow Account (FFA) and the Housing Economic Account (HEA).

B.2.2 Synthetic Savings Decompositions

The synthetic savings decompositions are based on the following transition equation:

$$W_{t+1}^g = (1+q_t^g)(W_t^g + s_t^g Y_t^g)$$
(4)

This equation can be extended to any specific asset on aggregate or asset owned by a specific wealth group. It is assumed that capital gains occur after all savings are invested.

Since the group-specific capital gains rate, q_t^g , and savings rate, s_t^g , are not observed in the data, they have to be computed using related information and some assumptions described here. This is done in two main steps. First, the asset-specific capital gains rates are calculated using data from the Financial Flow Account (FFA).²⁶ The average capital gains rate for each asset, q_t^j , can be computed using the asset-specific transition equation:

$$1 + q_t^j = \frac{W_{t+1}^j}{W_t^j + F_t^j}$$

Here, W_t^j is the stock of the asset and F_t^j is the financial flow observed in the FFA. Under the

²⁶See Appendix Section A.3 for details.



Figure 8: Comparison Between Scatterplot of Savings Rate and LOWESS This figure plots the synthetic savings rate of the top 1% and bottom 99% as a scatterplot and with the smoothed LOWESS values.

assumption that asset-specific capital gains are homogeneous across the distribution, then the group-specific capital gain on each asset is $q_t^{jg} = q_t^j$. In Appendix Figure 7, I present the cumulative total of capital gains and financial flows by asset category from 1990-2018. The figure shows how capital gains account for a large proportion of the overall increase in wealth during this period, particularly for housing and equities.

The second step involves computing the asset-by-group level of savings using the capital gains calculated in the previous step. From the asset-specific transition equation, the amount invested by each group into each asset is:

$$S_t^{jg} = s_t^{jg} Y_t^g = \frac{W_{t+1}^j}{1+q_t^j} - W_t^j$$

Summing across all assets yields the total savings for group g, $S_t^g = \sum_j S_t^{jg}$, and dividing this by their income yields their total savings rate, s_t^g .

To recover the group specific capital gains rate, q_t^g , the group savings total, S_t^g , can be inputted into the group-specific transitional equation:

$$1 + q_t^g = \frac{W_{t+1}^g}{W_t^g + S_t^g}$$

Using these equations, data from the FFA and estimates of the wealth and income held by



Figure 9: Savings Rate and Capital Gains Rate by Wealth Group, 1990-2018 This figure plots the savings rate and capital gains rate over time by wealth group as determined using the synthetic savings approach. The line represents the LOWESS (locally weighted scatterplot smoothing) estimate of the underlying scatterplot values. The figures compare the results with and without corporate savings being included in the savings of households.

each group , I can recover the values of q_t^g and s_t^g for each group in each year.

When plotting the figures of q_t^g and s_t^g , there is significant year-to-year variation in the data, which makes it difficult to determine the trends. To address this, I smooth the trends using a LOWESS (locally weighted scatterplot smoothing) approach. In App. Figure 8, I plot the trends in the savings rate as a scatterplot and with the LOWESS smoothed results. One observation from the figure is that there are certain years with extreme values for the top 1%. In 1995, the year after a major capital gains reform that encouraged people to make use of a capital gains exemption, there was a large jump in equity wealth from the prior year which was very low. Given that the actual rate of return on assets was not atypical, the transition equation rationalizes this with a very high savings rate - above 100% of income, which should not be possible. In 2009, the converse occurs as wealth falls dramatically from Q1 2008 to Q1 2009 as estimated using the capitalization method, which is based on dividend payouts and realized capital gains, but actual, overall capital gains were not negative enough to rationalize the decline in wealth - hence the large negative savings rate. Negative savings rates are possible as it implies money is being borrowed either from other wealth groups (note the jump up for the bottom 99%) or other "sectors" of the economy (eg. corporations, government etc.).

As discussed in detail in App. Section A.3, the inclusion of corporate savings plays a major role in the savings of those at the top of the wealth distribution. In App. Figure 4, I plot the synthetic savings rate if corporate savings is included and if it is not. Naturally, this decreases the savings rate of the top 1% and increases the capital gains since more of



Figure 10: Simulations of Top 1% Wealth Share Using USA Component Averages This figure shows the simulation results for the top 1% share of replacing the computed values for annual savings rates, capital gains rates and income inequality with the USA averages for the period. Each simulation is run separately for each component.

the change in wealth is ascribed to larger capital gains. Despite this, the trend is largely preserved, which means that while corporate savings affect the aggregate savings rate, they do not seem to change the finding that savings and not capital gains are the main driver of wealth inequality dynamics in Canada during this period.

Finally, I re-run the simulation for how wealth would have evolved in Canada had it followed trends from the United States. I present the results in App. Figure 10. I find that had Canada had the capital gains experience of the US during the 1990s, where the US outperformed Canada dramatically, the top 1% share would have been increasing more rapidly. However, the decline during the Great Recession was also far more extreme as well.