

The heterogeneous effects of COVID-19 on Canadian household consumption, debt and savings

James MacGee
Bank of Canada

Thomas Michael Pugh
Bank of Canada

Kurt See
Bank of Canada

Abstract. This paper develops an agent-based model to quantify the impact of COVID-19 on household debt and savings. To build a representative cross-section of households that vary by income, debt portfolios and consumption baskets, we merge data from the Survey of Household Spending and the Survey of Financial Security. We construct paths for consumption and employment over the crisis, accounting for heterogeneous risk of unemployment across demographics, government transfers, and substitution between expenditure categories that vary in contact intensity. Our model simulations yield a heterogeneous effect of COVID-19 across the income distribution. Low-income households face the highest risk of unemployment, but transfers provide generous income replacement. Middle-income job losers see the fastest rise in debt because transfers only partially replace lost income. Most unplanned savings are accumulated by high-income households that face lower risk of unemployment and larger declines in hard-to-distance spending. We find the rise in savings could generate a brief jump of nearly 6% of monthly consumption.

Résumé. Effets hétérogènes de la COVID-19 sur la consommation, la dette et l'épargne des foyers canadiens. Dans cet article, nous développons un modèle multiagent afin de chiffrer l'impact de la COVID-19 sur la dette et l'épargne des ménages. Compte tenu des disparités en matière de revenus, de portefeuilles de dettes ou de paniers de consommation, et afin d'obtenir un échantillonnage représentatif des foyers, nous avons fusionné les données de l'Enquête sur les dépenses des ménages et celles de l'Enquête

Corresponding author: Kurt See, seek@bankofcanada.ca

The authors would like to thank staff members of the Bank of Canada as well as seminar participants from the Arab Monetary Fund, Bank of Mexico and Federal Reserve Board of Governors for insightful comments. We also thank two anonymous referees and our editor for their guidance and advice. We thank Colin Grimes for excellent research assistance. An earlier version of this paper was published as Bank of Canada Staff Working Paper 2020-51. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Bank of Canada.

Canadian Journal of Economics / *Revue canadienne d'économie* 2022 55(S1)
February 2022. Printed in Canada / *Février 2022. Imprimé au Canada*

ISSN: 0008-4085 / 22 / pp. 54–87 / DOI: 10.1111/caje.12546

© Canadian Economics Association

sur la sécurité financière. Nous développons ensuite les trajectoires de la consommation et de l'emploi pendant la crise en tenant compte à la fois du risque hétérogène de chômage parmi les groupes démographiques, des transferts gouvernementaux et de la substitution parmi les catégories de dépenses nécessitant plus ou moins de distanciation physique. Les simulations de notre modèle montrent que la COVID-19 produit un effet hétérogène sur la distribution des revenus. Les ménages à faible revenu connaissent le plus grand risque de chômage, mais les transferts constituent de généreux revenus de remplacement. Les ménages à revenu moyen ayant subi des pertes d'emploi voient leur dette augmenter le plus rapidement, les transferts ayant seulement compensé une partie de leur baisse de revenus. Quant aux ménages à revenu élevé, ils ont accumulé la plus grande partie de l'épargne non planifiée, car ils sont moins susceptibles que les autres d'être touchés par des mises à pied et font face à des baisses de dépense plus importantes en matière de biens et services restreints par la distanciation physique (*hard-to-distance spending*). Nous constatons que la hausse de l'épargne pourrait générer un bref regain de la consommation mensuelle de près de 6 %.

JEL classification: E21, E24, G51

1. Introduction

THE COVID-19 PANDEMIC has seen an unprecedented decline in Canadian employment because public health measures and consumers' health concerns led to a contraction in spending, especially on contact-intensive goods and services. In this paper, we examine how the joint dynamics of household income and consumption during the pandemic shaped the buildup of debt and unplanned savings. We find that the concentration of employment losses within specific demographics (low income, young and less educated), the design of fiscal transfers to the unemployed, and differences in consumption baskets across the income distribution led to heterogeneous changes in households' balance sheets.

This paper makes three main contributions. First, we develop a novel approach to studying consumption dynamics that captures the unique impact of the COVID-19 crisis on consumption. We partition consumption into baskets of essential, hard-to-distance and luxury goods on the basis of the differential impact of the pandemic on expenditures. We integrate this approach with microdata on household expenditures to construct household-specific consumption portfolios across these expenditure classes. Second, we study the evolution of labour markets and consumption over the 2020 waves of the pandemic. Our approach combines aggregate national accounts statistics with high-frequency spending data to quantify the differential impacts of lockdown measures on each class of expenditure. Finally, we incorporate these elements into an agent-based model (ABM) that we use to simulate the evolution of household savings and debt and identify differential outcomes across the income distribution.

Our analysis begins with the construction of a representative cross-section of Canadian households' demographics, disaggregated consumption

expenditures, income and balance sheets. Because no individual dataset contains all of this information, we merge data from two household surveys: the 2016 Survey of Financial Security (SFS) and the 2017 Survey of Household Spending (SHS). Specifically, we use variables common to both surveys to impute household consumption expenditures in the SFS. A key feature of our approach is the allocation of household expenditures into four classes: essentials/easy-to-distance (e.g., groceries), hard-to-distance (e.g., travel, dining), luxuries (e.g., jewelry, fashion apparel) and shelter.

We use the Labour Force Survey (LFS) to guide and discipline the evolution of shocks to employment across age–education groups. A novel element of our approach is the inclusion of “COVID-reduced hours” as an employment state in addition to employed and unemployed states. Our definition of unemployment includes COVID-related non-participation and absences to capture how severe lockdowns and virus transmission risk depressed job-finding rates and discouraged laid-off workers from actively looking for jobs. Finally, we incorporate transfers to households, such as the Canada Emergency Response Benefit (CERB), that were introduced in response to the rise in unemployment. Our simulation thus features shifts in household income because of unemployment, reduced hours worked and changing government transfer programs.

To track the evolution of consumption, we map consumption categories in the 2019 and 2020 National Accounts into our SHS spending classifications. To construct time-varying and spending-class-specific paths of the evolution of essential, hard-to-distance, luxury and shelter spending, we combine this with high-frequency data on disaggregated consumption expenditure categories.

Our analysis of the data yields three main findings. First, consumption portfolios vary systematically with income. Of particular relevance is that high-income households account for a disproportionately large fraction of spending on hard-to-distance and luxury goods and services. Second, the pandemic differentially impacted spending, with hard-to-distance spending experiencing the largest contraction and slowest recovery. Our analysis also points towards consumer substitution towards luxury goods over the second half of 2020. Third, while job losses have been more severe for low-income, young and less-educated individuals, fiscal transfers such as CERB introduced a short-lived but sizeable increase in the income of workers with low earnings.

We develop an agent-based model that takes as inputs the initial distribution of agents and exogenous shocks to consumption and income. This allows us to simulate the evolution of savings and debt over the crisis in response to shocks to income and consumption and the introduction of mortgage payment deferrals. We find that the effects of the crisis on households’ balance sheets are heterogeneous. Our model simulations generate a buildup of excess savings driven by restrictions on consumption and the introduction of transfers during the crisis. While some households in all income quintiles accumulate excess savings, high-income households account for the majority of the rise in

savings. The top income quintile alone accounts for nearly 40% of the savings attributable to the crisis by the end of 2020. The larger decline in spending of higher-income earners plays a key role in these dynamics because the rise in savings is particularly pronounced among middle-aged and older homeowners whose fall in consumption is large relative to their income.

Our simulations point to a rise in the number of households with high debt payments relative to their income. This is driven largely by middle-income renters and homeowners with mortgages who experience an unemployment spell and for whom CERB only partially replaces lost income. Although most households have sufficient unused lines of credit to meet the gap between income and expenditures, for a small subgroup—largely renters, access to credit could become an issue once CERB ends. While low-income households experience the highest risk of unemployment, CERB provides a relatively high replacement (or increase upon) of previous income and thus limits the rise in their debt. Higher-income households saw a slower initial rise in debt during the crisis because of less exposure to employment losses and larger declines in consumption expenditures. However, we find that debt begins to rise after consumption habits adjust and a substitution from hard-to-distance towards luxury spending drives increased spending. Despite these forces, the availability of mortgage deferrals combined with the relatively rapid recovery in employment limits both debt accumulation after income losses and the buildup of households with high DRSS.

We adapt our model to consider the implications of higher unplanned savings and the buildup of debt for post-COVID consumption. Our approach builds on estimates of the marginal propensity to consume out of transitory income shocks by Fagereng et al. (2021). They find that the fraction of lottery winnings spent depends on its magnitude and household wealth. In our experiment, we assume that post-COVID households treat spending of their excess savings akin to lottery winnings, while households whose debt grows reduce consumption because of higher debt service payments. We find a short-lived but significant bounceback in consumption of approximately 4%.

Our paper builds on a growing empirical literature that examines the linkages between consumption and income over the pandemic, including Achou et al. (2020), who surveyed Quebec households and asked how their spending and income had changed; Cox et al. (2020), who use US bank account data to investigate saving and spending over the income distribution; Chetty et al. (2020), who study the effects of the pandemic on household and firms across US localities; Hacıoğlu-Hoke et al. (2021), who use UK fintech app data to examine types of spending and the development of consumption and income inequality; and Coibion et al. (2020), who examined how households' income losses, expectations and spending were impacted by social restrictions. A common message from these papers is that the initial pandemic income and job losses were concentrated among low-wage workers, whilst the fall in spending

was concentrated among high-wage households and in goods and services with greater personal interaction. We share with Achou et al. (2020) a Canadian focus and a focus on the heterogeneous changes in spending and income across households similar to Cox et al. (2020) and Hacıoğlu-Hoke et al. (2021). We differ in our approach to these questions because we use an agent-based model to account for the impact of COVID-19 on household balance sheets over 2020. Our paper complements these analyses by developing a methodology that uses a time-varying allocation of expenditures into classes of spending to capture the pandemic's impact on consumption as well as a model simulation strategy to map shifts in consumption and income into changes in household saving and debt.

These analyses overlap with work on the lessons from the pandemic experience for the effect of fiscal transfers on consumption. In a recent survey of this work, Hackethal and Weber (2020) find that the poor often spend transfers whilst richer households cut spending by more. Our work is also related to Carroll et al. (2020), who model the US CARES act and recovery in an environment where some income losses are temporary but others are persistent and Baker et al. (2020) who note both the high spending of fiscal transfers by low-income households and the quick rebound of sales of essential goods. Our analysis complements these papers by examining how transfers impacted Canadian consumption and shaped the evolution of household saving and debt over the pandemic.

We also study the heterogeneous build up of debt across households as a result of the effects of the pandemic on incomes and consumption. Our work is related to Bilyk et al. (2020), who find that only a fifth of Canadian households with mortgages have sufficient liquid assets to cover two months of mortgage payments. We build upon this insight by constructing a cross-section of households to evaluate which households will use credit lines. Similar to Kaplan et al. (2020), who use a HANK model to study the heterogeneous impact of COVID-19 on US households, we focus on the heterogeneous impact of COVID-19 on Canadian households. Our ABM setup incorporates a novel approach to categorizing the time-varying impact of social distancing restrictions on consumption goods and services.

2. Start point: Financial distribution of households in Canada

The ideal dataset to study the evolution of household finances and consumption over the pandemic would include household demographics and the composition of wealth, income and consumption. Currently, there is no dataset that satisfy all of these criteria. We address this deficiency by linking the 2016 Survey of Financial Security (SFS) (Statistics Canada 2016) which reports household balance sheet information with the 2017 Survey of Household

TABLE 1

Wealth distribution by income quintile

Variable	Inc. 1	Inc. 2	Inc. 3	Inc. 4	Inc. 5
Mean gross income (\$/month)	941	2,356	4,367	7,272	16,286
Mean net income (\$/month)	1,550	3,222	4,757	6,907	13,165
Liquid assets < \$1,000	0.40	0.27	0.21	0.12	0.06
Mean positive liquid assets \$	2,605	44,517	49,634	75,465	180,645
Proportion with auto debt	0.09	0.20	0.31	0.44	0.43
Mean auto debt \$	15,627	15,981	18,910	22,344	26,148
Proportion renting	0.67	0.44	0.31	0.18	0.08
Mean rent (\$/month) ^a	793	1,038	1,122	1,251	1,400
Proportion with mortgage	0.11	0.23	0.30	0.54	0.60
Mean mortgage payment (\$/month)	1,055	1,000	1,109	1,319	1,752
Proportion with line of credit (LOC) debt	0.07	0.16	0.21	0.30	0.35
Mean LOC balance \$	31,009	24,554	31,238	35,228	64,387
Proportion with unsecured debt	0.40	0.45	0.51	0.53	0.47
Mean unsecured debt \$	10,082	11,128	13,091	14,897	18,965
Liquid assets < 1 month mortgage payments	0.18	0.16	0.16	0.13	0.08
Liquid assets < 1 month income	0.37	0.39	0.38	0.40	0.35

NOTES: Cross-section of households based on the Survey of Financial Security (SFS) and the Survey of Household Spending (SHS). Respondents are grouped by income quintile, net of taxes. Liquid assets is the sum of cash, deposits, stocks, tax-free savings accounts, bonds and mutual funds. Unsecured debt includes student loans, credit card balances and any remaining other loans. Dollar values in 2016 CAD.

^aSHS values. Other table values from SFS.

Spending (SHS) (Statistics Canada 2017), which reports disaggregated household expenditures.¹

We first use the SFS to construct key summary statistics of the Canadian wealth, debt and income distribution by post-tax household income quintiles (see table 1).² We define liquid assets as the sum of cash, deposits, stocks, tax-free savings accounts, bonds and mutual funds. Despite this broad definition there are many households with low liquid assets even in the middle of the

1 These are, respectively, the major surveys used by Statistics Canada for the distribution of household wealth and spending. See Statistics Canada table 11-10-0223-01, available at www150.statcan.gc.ca/t1/tb11/en/tv.action?pid=1110022301, Norris and Pendakur (2015) or Gellatly and Richards (2019). The SFS and SHS also provide the distributional basis for the Statistics Canada Distribution of Household Economic Accounts Social Policy Simulation Database/Model (Statistics Canada 2020a), which we find our results broadly consistent with (Statistics Canada 2021).

2 For comparison, note that the unemployment-triggered CERB payments during the crisis are approximately \$2,000 (pre-tax) per month.

income distribution, suggesting that some households could be vulnerable to transitory declines in income.³

The SFS details household balances and payments on mortgages, auto loans, lines of credit, student loans, credit card debt as well as monthly rent (for renters). Few high-income households rent, while roughly two thirds of the lowest income quintile are renters. Not surprisingly, monthly rent rises less quickly with income than mortgage payments. Line-of-credit debt in Canada is associated with higher income (likely because of income- and wealth-based eligibility for personal credit lines) and the presence of a mortgage (because of the prevalence of home equity lines of credit). However, the proportion holding unsecured debt is more uniform across incomes.

The SHS reports households expenditures by detailed, granular categories as well as characteristics of the household and its members. Given the prominent role of voluntary and mandated social distancing restrictions during the pandemic, we aggregate expenditures into classes that reflect the ease with which their purchase/consumption can be “distanced” as well as whether it is essential (e.g., food, shelter). Specifically, we group expenditures into one of four spending classifications to understand the impact of the pandemic:

1. **Easy-to-distance essentials (ETD)**: Goods that households must continue to consume (e.g., groceries, medications) or are easy to consume during social/physical distancing (e.g., deliverable electronic goods).
2. **Easy-to-distance luxuries (L)**: Discretionary goods on which spending often declines when income falls (e.g., jewellery) but are easy to consume during a pandemic.
3. **Hard-to-distance goods (HTD)**: Goods for which regulations, increased aversion to risk of infection or guidelines prevent or reduce consumption (e.g., overseas travel, drinks at bars, gym classes).
4. **Shelter costs (S)**: Rent payments, property taxes, key maintenance and so on (but not including mortgages, which are part of our dynamic debt modelling).

We develop a schedule for the division of SHS consumption goods into these *spending classes* using Health Canada Guidelines.⁴ Our procedure is not discrete for each good because we divide certain goods between spending classes. For example, travel is proportioned between essential ETD, hard-to-distance and luxury. This procedure yields the benchmark consumption groups for each household ($C_{i,s}$). We detail this schedule in appendix A1.

To merge the surveys, we use the multiple imputation by chained equations (MICE) procedure of van Buuren and Groothuis-Oudshoorn (2011) to

3 For low liquidity “wealthy hand-to-mouth” households, see Kaplan et al. (2014).

4 Available at www.canada.ca/en/public-health/services/publications/diseases-conditions/covid-19-going-out-safely.html.

impute consumption for households in the SFS using expenditures reported in the SHS. Our imputation employs demographic and income variables in both the SFS and SHS to predict household expenditures.⁵ The procedure adds a random residual to this conditional expenditure prediction, drawn from the estimated residual distribution after accounting for these controls. We use a non-linear random forest machine-learning algorithm to generate conditional expectations of expenditures. To ensure that the distribution of consumption matches across income and housing status (renter, mortgagor, owner), we apply an adjustment procedure on the algorithm-imputed consumption values. We group SFS households by income ventiles and housing status and align median consumption within each group to their SHS equivalent.⁶ Comparing classes of spending across income (the top and bottom panels in table 2) shows that the results of the linking procedure align closely with the SHS data.

There are significant differences in expenditure shares on goods across households. While high-income households naturally account for a larger share of total consumption expenditures, their shares of total expenditure on hard-to-distance and luxury items are relatively larger. For example, the top 20% of households (by income) account for 39% of aggregate expenditure on luxuries and 39% of hard-to-distance goods. In comparison, the lowest 20% of earners account for less than 8% of either spending class. This pattern of increased prominence of luxury and HTD holds with individual consumer budgets because the top 20% earners spend on average 30% of their expenditure on non-essentials (luxury plus HTD) whilst the lowest 20% spend less than 20% of their consumption basket on non-essentials. The distribution of consumption expenditures plays an important role in our simulations because it results in different impacts on consumption expenditures across households from COVID-necessitated restrictions on travel and hard-to-distance consumption.

In summary, our linking procedure yields a sample of households with information on balance sheets, income and consumption by spending class. We use this distribution as the starting point for our simulations of the COVID-19 crisis. To simulate the evolution of household balance sheets over the crisis, we combine this initial distribution with time-varying shocks to consumption spending classes and employment/income shocks. The next steps, discussed in sections 3 and 4, involve the construction of a path for consumption across spending classes and employment/income shocks during the crisis.

5 We use: household net income, transfer/government income, key demographics (head of household age, sex, education, location, household composition—number and ages of children, household size, vehicle leasing or owning, type of dwelling, rent/owner status and size of mortgage payments for mortgagors).

6 For this procedure, we use the SHS subset including diary observations so as to capture all consumption items.

TABLE 2

Expenditure shares by income quintile

Spending class	Inc. 1	Inc. 2	Inc. 3	Inc. 4	Inc. 5
<i>Shares of spending class (SFS-imputed)</i>					
Essentials-ETD	0.10	0.14	0.19	0.24	0.33
Luxuries	0.07	0.13	0.19	0.23	0.39
HTD	0.07	0.12	0.17	0.25	0.39
Shelter	0.16	0.19	0.22	0.20	0.23
<i>Shares of spending class (SHS)</i>					
Essentials-ETD	0.10	0.15	0.20	0.23	0.32
Luxuries	0.08	0.13	0.19	0.22	0.37
HTD	0.08	0.13	0.18	0.23	0.38
Shelter	0.16	0.20	0.20	0.20	0.23

NOTES: Cross-section of households based on the Survey of Household Spending (SHS) and the Survey of Financial Security (SFS). Share of spending of each income quintile in overall spending within a spending class. ETD = easy-to-distance. HTD = hard-to-distance.

3. Consumption in a crisis

In this section, we describe our construction of the path of consumption for each spending class over the simulation period. This allows us to speak to two key characteristics of the COVID-19 pandemic. The first is the combined effect of containment measures, behavioural changes and increased precautionary saving in depressing overall consumption. The second is substitution between spending classes that have been differentially affected by COVID-19. For example, the decreased spending in hard-to-distance activities such as dining out or travelling may have engendered increased demand for other goods such as home entertainment or personal fitness and sporting equipment.

We adopt a two-stage procedure to trace out a consumption path for each spending class, starting with the individual distribution of consumption $C_{i,s,t=0}$ across spending classes $s \in \{E, L, HTD, S\}$ described in section 2.

In the first stage, we construct a path for each spending class $\{C_{i,s,t}\}_{t=1}^T$, where T is the simulation horizon. We discipline this exercise by combining data from Canada's 2019 and 2020 National Accounts with high-frequency year-on-year spending growth from the RBC spending tracker.⁷ In the second stage, we adjust the implied path of imputed aggregate consumption to match the observed drop in consumption from the 2020 National Accounts. We outline each stage in detail below.

Our imputed database from section 2 forms our initial distribution. Next, we construct adjustment factors $\{\gamma_{s,t}\}_{t=1}^T$ to obtain a time path for

⁷ See table 36-10-0124-01 (formerly CANSIM 380-0085) for the National Accounts. RBC spending tracker reports are available at <https://thoughtleadership.rbc.com/covid-consumer-spending-tracker/>. The combination of sources is key because of greater granularity, and frequency in RBC, which enables us to effectively track the impact of lockdowns and re-openings.

consumption over the simulation horizon. That is, aggregate consumption for time $t > 0$ for spending class s is $C_{s,t} = \gamma_{s,t} C_{s,0}$.

We use the National Accounts to obtain consumption levels for a detailed list of expenditure categories in 2019.⁸ We then use RBC's high-frequency year-on-year growth figures to arrive at an estimate of 2020 monthly consumption levels across these categories. This procedure relies on the mapping of common major categories in the National Accounts and the high-frequency spending data.⁹ This provides a monthly estimate of each major consumption category's consumption level in 2020. We emphasize that these major consumption *categories* (Dining, Transport, Recreation, etc.) should not be confused with the four spending *classes* we define earlier. Below, we detail how we break down each consumption category (e.g., Transport) into different spending classes (E, L, HTD, S).

Because the National Accounts and the SHS share similar categorizations of consumption, we adopt the same apportionment of consumption categories into spending classes. For example, the National Accounts expenditure category "Transport" contains subcategories such as "Air transport," which we assign fully (100%) to HTD, but also "Taxi and limousine," which we assign as 75% Essential and 25% HTD. Given these weights and the reported expenditures of each Transport subcategory, we aggregate the subcategories to obtain a breakdown of Transport into spending classes E–L–HTD–S. This enables us to account for shifts in consumption across subcategories (e.g., from air travel or public transport to taxis or private cars) and thus how the E–L–HTD–S composition of each major category shifted over 2020.¹⁰ For example, this procedure implies that in 2020 Q1, Transport was 36.4% essential, 1.5% luxury and 62.1% hard-to-distance. By 2020 Q2, this had shifted to 49.4% essential, 1.7% luxury and 48.9% hard-to-distance, reflecting that some transport activity moved away from HTD modes (air transport, public transport) to modes that were considered essential and relatively easy to distance (e.g., non-contact-intensive taxis and private cars).

To take stock: we now have: (i) aggregate consumption for each major category by month in 2020 and (ii) how each major category is broken down into the four spending classes E–L–HTD–S over 2020. We now construct a path for each of the four spending classes. Let $\phi_{r,s,t}$ be the weight of major consumption category r in spending class s in time t and $X_{r,t}$ be the derived value

8 This data is from www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3610012401.

9 Supplementary files that describe this mapping can be found in our online repository (see appendix A1). For categories in the National Accounts that are absent from the high-frequency spending data, we use less granular information from the 2020 National Accounts instead.

10 For consumption in 2021, we assume the continuation of the composition in 2020 Q4.

of major consumption category r in time t . Then, the adjustment factor $\gamma_{s,t}$ for spending class s (e.g., Essentials) is given by $\gamma_{s,t} = \frac{\sum_{r \in R} \phi_{r,s,t} X_{r,t}}{\sum_{r \in R} \phi_{r,s, \text{Feb 2020}} X_{r, \text{Feb 2020}}}$ wherein we normalize to the level of consumption in February 2020.

We combine the constructed time series for each $\gamma_{s,t}$ with our initial distribution of consumption (from section 2) $C_{i,s,t=0} \forall i$ to trace out a time path of consumption for each agent i at time t : $C_{i,s,t} = \gamma_{s,t} C_{i,s,0}$. Further, we assume that expenditure declines by 10% whilst a household is unemployed, in line with Jappelli and Pistaferri (2010) and Christelis et al. (2015). Finally, we use changes in aggregate consumption from the 2020 National Accounts relative to 2019 Q4 to adjust our consumption paths so that the model's aggregated quarterly consumption series aligns with aggregate data.

Figure 1 plots the resulting paths of $\{\gamma_{s,t}\}_{t=0}^T$ (left panel) and consumption $\{C_{s,t}\}_{t=0}^T$ for $s \in \{E, L, HTD, S\}$ (right panel). This shows that hard-to-distance spending experienced the largest drop and slowest recovery during both lockdowns. Essential and shelter spending experienced more modest declines and largely recovered by the end of 2020. In contrast, luxury (but non-contact-intensive) spending experienced an initial decline in March but then substantially overshot its pre-pandemic levels because households substituted away from hard-to-distance spending.

Differentiating spending between classes allows us to capture substitution, as reflected by the shifting weights of spending from hard-to-distance goods towards luxuries over time. Meanwhile, our aggregate adjustment to match the aggregate decline in consumption relative to 2019 levels allows us to, albeit in a rudimentary way, capture precautionary saving motives that our model will not be able to account for. In this sense, this procedure captures both

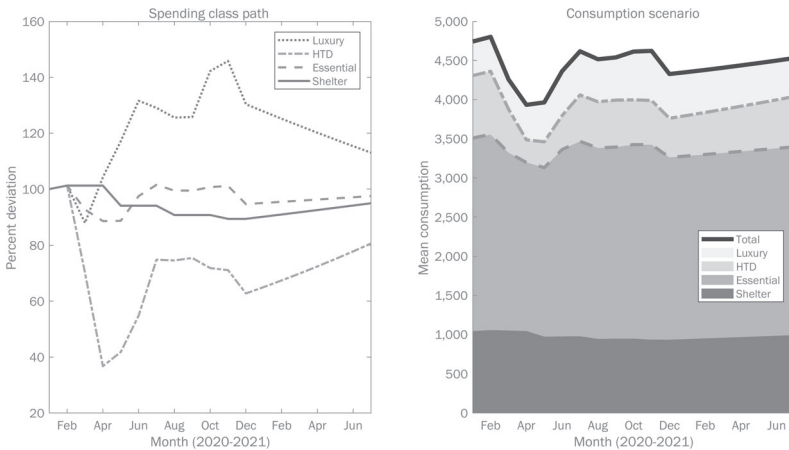


FIGURE 1 Baseline consumption scenario
NOTES: This figure plots the path of mean consumption and its subcategories over the simulation period. Consumption categories in the Survey of Household Spending (SHS) are here broken down into four categories: essential, luxury, hard-to-distance and shelter. Detailed definitions found in text.

compositional shifts in household consumption bundles as well as observed level declines.

4. Pandemic employment risk

We use publicly available monthly Labour Force Survey (LFS) (Statistics Canada 2020b) microdata to estimate a probabilistic employment status process of households' transitions over the simulation period.

To capture the unprecedented disruptions caused by COVID-19, we construct COVID-adjusted employment measures that track the rise in reduced hours and unemployment. Figure 2, which compares 2020 to 2019 and the financial crisis years of 2008 and 2009, shows that the COVID-19 crisis is unique and requires different employment measures. We focus on two major deviations for understanding unemployment—the startling rise of “absence” from work in spring (much of which was unpaid) and the drop in the participation rate. These two features drive our decision to adopt a broader measure of unemployment, which we detail below.¹¹ Also shown in the subfigure

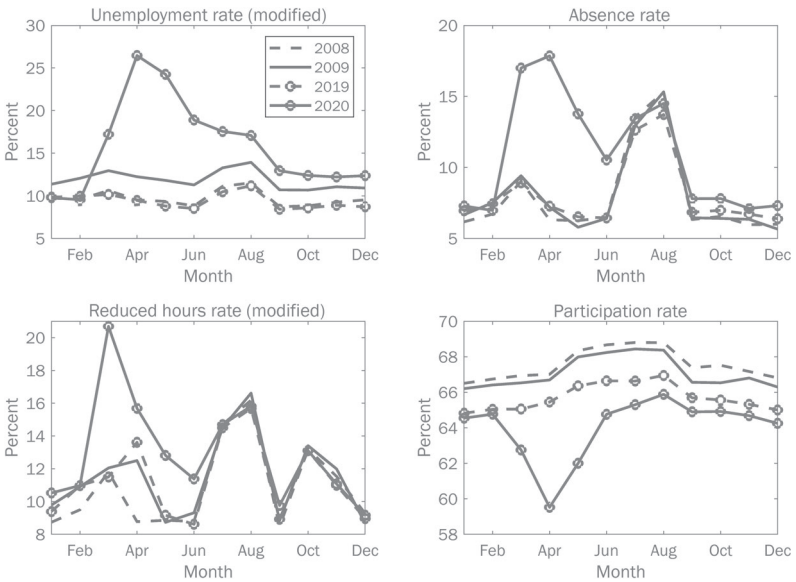


FIGURE 2 Labour markets during the Great Recession vs. the COVID-19 pandemic
NOTES: This figure compares the path of unemployment (including absent without pay), reduced hours (less than 80% of usual hours or absent with pay), absence rate (paid and unpaid) and participation rate during the Great Recession and the COVID-19 crisis. Employment information is obtained from the Labour Force Survey.

11 We use the same COVID-19 adjusted definition for all years of comparison.

reporting reduced hours is the jump in employed workers reporting significantly less than usual hours ($< 80\%$) in spring of 2020, which we model as an additional labour market state.

We consider a respondent fully employed (E) if they report being employed, non-absent and working hours above 80% of usual hours. The reduced hours (RH) category includes two groups: (i) workers who report being employed but with actual working hours less than 80% of their usual hours and (ii) workers who report being absent from work but receiving pay. Finally, because lockdowns may have discouraged laid-off workers from actively looking for jobs by depressing job-finding rates, we modify our definition of unemployment (U) to capture COVID-related non-participation and absences. Thus, we broaden the standard definition of unemployment to include: (i) workers displaced between March and May who report being out of the labour force but would like to work and (ii) workers who report being absent from work without pay.

The COVID-19 crisis has disproportionately affected younger, less-educated and lower-wage workers.¹² To incorporate the heterogeneous risk of unemployment across households, we use the LFS to estimate age- and education-specific month-to-month employment status transition matrices. Because the SFS (and SHS) reports household income while the LFS public-use files report only individual income, we assume that income shocks within households are perfectly correlated, i.e., that in households with two (or more) workers, both (or all) become unemployed or employed.

We divide LFS labour force participants into six groups on the basis of education and age—two education groups (college, non-college) and three age groups (0–39, 40–54, 55+). We estimate a time-varying (Markov) transition matrix for each age–education group, represented in table 3.¹³

We match the LFS densities of each of the three states with data and restrictions on the transition matrix. We begin with proxy measurements of these flows then apply matrix restrictions and adjustments to ensure the flows generate the observed rates of unemployment and reduced hours households. We use the duration of unemployment to identify flows of newly unemployed for each month and thus the probability of moving from E_{t-1} to U_t (position (1,3), where row 1 is the state in $t - 1$ and column 3 is the state in t in table 3). We perform the same exercise for full time job finding rates U_{t-1} to E_t and allocate finders of part-time jobs to U_{t-1} to RH_t . The flows into reduced hours from employment E_{t-1} to RH_t are assumed to be the same as those from employment into unemployment E_{t-1} to U_t , denoted F .

12 For a more detailed discussion of the initial impact of COVID-19 on Canadian labour markets, see Lemieux et al. (2020).

13 One may think of this as being allocated to a 2-D slice of a 5-D array of {Previous employment state, Current employment state, Age, Education, Time}.

TABLE 3

Employment transition matrix

$t - 1$ to t	E_t	RH_t	U_t
E_{t-1}	$1 - RU - F$	ER	F
RH_{t-1}	RE	$1 - ER - F$	F
U_{t-1}	UE	UR	$1 - UE - UR$

NOTES: This table illustrates a representative employment-status transition matrix for a given age–education group and a given month t . E = employed. RU = reduced hours to unemployed. F = employment or reduced hours to unemployment. ER = employed to reduced hours. RH = reduced hours. RE = reduced hours to employed. U = unemployed. UE = unemployed to employed. UR = unemployed to reduced hours.

Given these initial probabilities, we apply a single-direction exclusion restriction on transitions in and out of RH. If the RH group is increasing in mass, we set the employment outflow RH_{t-1} to E_t to zero and recalculate the probability of E_{t-1} to RH_t that is required to create the the RH population in the data at t from the previous state vector $(E_{t-1}, RH_{t-1}, U_{t-1})$. This transition matrix for each group is adjusted to ensure the flows reproduce the rates of unemployment and reduced hours. We change the probability of E_{t-1} to U_t to generate correct U_t . To generate correct RH_t , we modify E_{t-1} to RH_t or RH_{t-1} to E_t and U_{t-1} to RH_t .

Our estimation accounts for the compositional difference between LFS individuals and SFS households. We use simulated method of moments targeting the aggregate unemployment and reduced hours rates in the LFS with two common shift parameters which raise or lower $P(U_t | E_{t-1})$ and $P(RH_t | E_{t-1})$ for all household age–education groups in the SFS. The LFS-comparable individual unemployment and reduced hours rates are computed from the simulated SFS under the assumption that intra-household employment shocks are perfectly correlated.

For simulation periods in 2021 (after the latest data available as of writing), our estimation of transition probabilities follows the same structure but uses data from the 2019 LFS mapped to our future employment scenario. Our COVID-adjusted unemployment measure is assumed to converge gradually down to 10% (roughly equivalent to pre-pandemic levels) while excess reduced hours return to zero. The appendix includes the underlying calibrated probabilities in figure A1.

Figure 3 compares our model’s sequence of employment, reduced hours and unemployment rates (solid) with the data (dashed) that we target in our simulation. Our COVID–RH series matches 2020 reduced hours in excess of 2019 values (month-on-month). By construction, this RH rate is close to zero prior to March and converges to zero after the pandemic’s first wave.

Underlying the aggregate employment paths are the heterogeneous labour market outcomes depicted in figure 4. All age–education groups experienced a rise in unemployment during the initial spring lockdown months, but younger workers without degrees were particularly affected. Reduced hours were

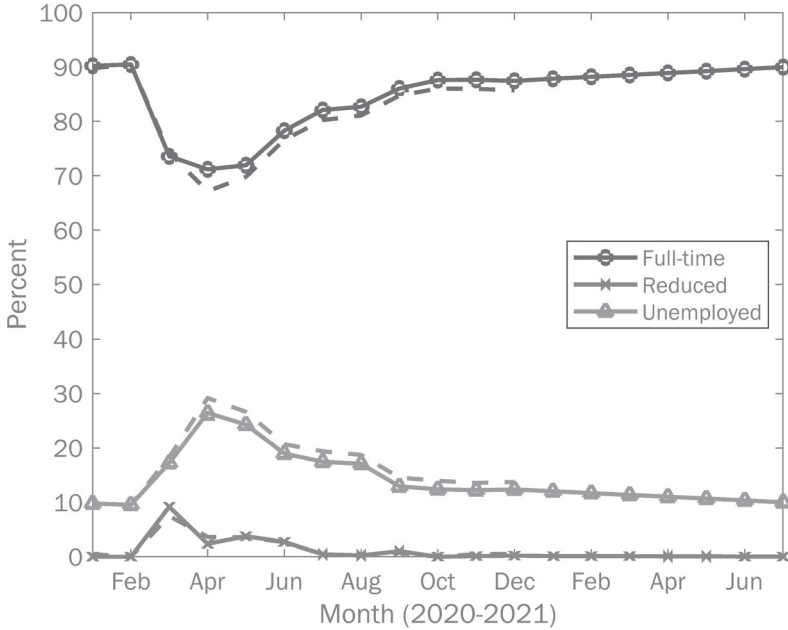


FIGURE 3 Baseline sequence of employment, reduced hours and unemployment
NOTE: This figure plots the simulated path of employment, reduced hours and unemployment rates in the simulation together with comparable measures computed from the latest available Labour Force Survey data (dashed).

initially equally distributed, but older workers lingered in this status longer than the young. The decline in unemployment slows by the fall, with unemployment well above pre-pandemic levels by the end of 2020. While those with degrees had lower rates of unemployment, the persistence of elevated unemployment levels was higher.¹⁴

Heterogeneous impact on household income

We now explain how employment status translates into income in our model. Employed households are assumed to earn their reference income reported in

14 The estimated employment process matches the patterns of the aggregate unemployment and reduced hours rates, but through the different age and education demographics of heads of households in SFS versus the population in the LFS, and the correlated income assumption, our estimation process will not fit exactly for each age-education group. In particular, the matching requires slightly higher unemployment rates amongst older and more-educated households because these groups have greater representation amongst the population of SFS households than that of LFS individuals. We do not see this as problematic because these households, on average, contain some lower-education and younger members.

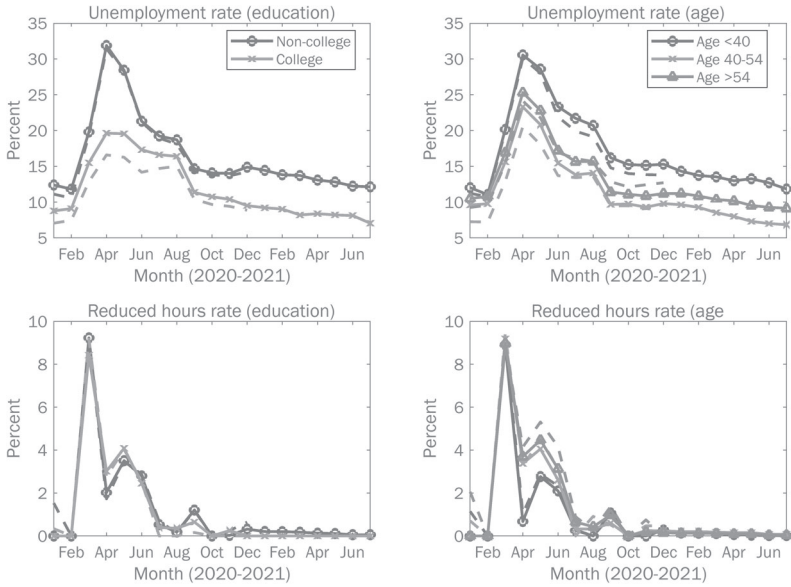


FIGURE 4 Unemployment and reduced hours rates by demographic groups
NOTES: The figure compares the model-generated unemployment and reduced-hours rates with the series constructed from the Labour Force Survey (LFS) (dashed). The LFS unemployment rate is adjusted for COVID-induced non-participation while reduced hours are in excess of 2019 levels (see section 4 for details). Respondents are classified as college or non-college graduates, and the three age groups are under 40, between 40 and 54 and above 54.

the SFS.¹⁵ We assume that households experiencing COVID-reduced hours receive only 75% of their reference income.¹⁶ Finally, the unemployed receive appropriate government transfers active during the pandemic. Qualifying unemployed households receive CERB payments (\$500 weekly per adult household member) from April 2020 to September 2021. Outside of the CERB program, the unemployed receive Employment Insurance (EI) or Social

15 Because we lack reference employment income for unemployed households in the SFS, we regress income on household features plus a random residual to impute a reference income. We follow a similar process to the consumption imputation in section 2.

16 In the data, this proportion did not vary substantially with demographics, income or by month during the crisis and was tightly distributed around 75%. The CEWS employment subsidy provides many workers on reduced hours with 75% of income, which theoretically censors the lower tail of net income for such households to 75%.

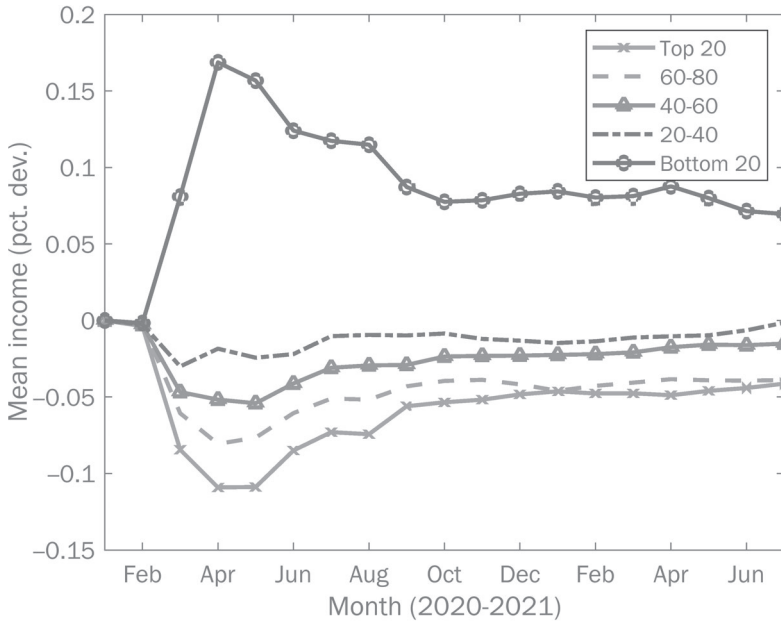


FIGURE 5 Shifts in household income by quintile of pre-COVID income
NOTES: This figure plots the percent deviation of simulated average household income from pre-crisis levels for employed households in the Survey of Financial Security (SFS). The quintiles are based on pre-crisis household income. The plotted series describes the simulated path of income once the employment shocks and accompanying Canadian Emergency Response Benefit (CERB) transfers are introduced.

Assistance (SA) payments.¹⁷ In the simulation, if households would receive more income from government benefits than working, we allow them to choose the higher income.

The heterogeneous rise in unemployment and the structure of fiscal transfers have different impacts across household income groups. To illustrate these differences, we examine the path of income for households grouped by quintiles of pre-crisis income in figure 5. The bottom 20% of earners (on average) saw a temporary rise in income because of CERB exceeding the pre-pandemic earnings for many of these households. Because CERB pays a flat dollar amount, the fall in average income for higher-income quintile groups is larger because of a smaller replacement rate.

After the recovery in employment in June, income for the upper income groups recovers substantially versus other groups. This group also experiences a slight depression in income during the second wave at the end of the year.

¹⁷ After September 2020, CERB was replaced by the Canada Recovery Benefit (CRB) program which pays a similar biweekly amount, which we model. EI features a 55% replacement rate up to a maximum benefit level while SA pays close to \$1000 per adult.

Meanwhile, the gradual loss of CERB payments causes the bottom 20% to lose around half of the income boost by year end.

5. Agent-based simulation: Savings and debt dynamics

Agent-based models specify rules for how agents act on the basis of their individual inputs in order to simulate the behaviour of the system.¹⁸ Our agent-based modelling approach begins with the initial distribution of households (the agents) across their demographics as well as initial wealth, income and consumption obtained from section 2. The evolution of household debt and savings is determined by a simple budget constraint rule on the basis of heterogeneous and time-varying inputs of consumption spending classes (from section 3), income shocks (from section 4) and endogenous previous wealth.

We incorporate information on each household's debt payments. For mortgages, the SFS reports the interest rate and monthly payments in addition to balance. We use this information to estimate a term length assuming that no refinancing occurs. For other debt categories, we construct an amortization schedule by assuming fixed terms and rates, as per table A1.

Households are allowed to defer their mortgage payments between March and August 2020, in line with deferral options during the crisis.¹⁹ Deferred interest payments increase the outstanding balance and imply that mortgagors who defer face higher payments post-deferral. We calibrate a rule on the basis of liquid assets and loan-to-value ratio (LTV) required for a mortgagor to qualify for a deferral such that 15% of mortgagors defer, and 20% of those who defer have an LTV of less than 50%, consistent with Allen et al. (2021).

Two key assumptions in this exercise are no new loans for durable goods (e.g., home purchases) and no default.²⁰ We also abstract from asset price changes because liquid wealth is a single variable in our model. This is motivated by evidence suggesting that much of the change in savings remained in liquid assets as deposits and other near monetary objects (Statistics Canada 2020c).

Simulation

We use the SFS to construct households' initial portfolios. In our simulation, we track households' liquid savings net of their total line-of-credit balance B . We initialize B with the household's reported liquid savings (cash; non-registered mutual funds, other investments, bonds, stocks and shares; tax-free

18 See www.oecd.org/naec/Agent-based_models_background.pdf for the principles of agent-based modelling.

19 Many households opted for a deferral during the early months of the pandemic. See www.cmhc-schl.gc.ca/en/about-cmhc/corporate-reporting/mortgage-deferral-numbers for more details.

20 Personal insolvency filings remained well below pre-pandemic levels during 2020. See MacGee (2012) regarding personal insolvency in Canada.

savings account) net of their line-of-credit balance. Thus, we impose that agents first tap their liquid savings (if $B > 0$) before drawing on credit to finance any shortfall in their income. Given initial balance B as well as paths for individual income Y , consumption C and debt service obligations, we simulate the evolution of debt with the dynamic budget constraint:

$$B_{i,t+1} - B_{it} = Y_{it} - C_{it} - \sum_i DS_{it}^j + I\{B_{it} > 0\}B_{it}r, \quad (1)$$

where $j \in \{\text{Mortgage, Auto, Line of credit, Credit card, Instalment, Student, Other}\}$, I is an indicator function and $DS^{LOC} = 0$ when $B > 0$.

For households with $B_t < 0$, i.e., borrowing on a line of credit, we specify an amortization schedule to determine the monthly payments associated with B . We assume a 15-year term and interest rate equal to $i^{\text{mortgage}} + 3\%$ to determine payments on lines of credit.²¹ When $B_t > 0$ households receive interest r , which we set equal to an annualized value of 2%.

Households whose income Y is less than the sum of consumption expenditures C and debt service obligations from debt $\sum_j (DS^j)$ see their debt rise. Conversely, households whose income exceeds consumption and debt payments see (net) liquid savings rise.

6. Results

Our simulation delivers several insights into the heterogeneous impact of COVID-19 on Canadian household finances. First, some households within each income group experience a rise in savings, albeit for different reasons. Some lower-income households saw their earnings rise because of CERB transfers, which more than replaced 100% of lost income and thus contributed to higher savings. However, the majority of the rise in savings during the lockdowns in March to May and later in the year was by higher-income households who saw large declines in consumption expenditures because of the prominence of hard-to-distance goods in their regular consumption bundle—although the substitution to luxury goods discussed in section 3 partially mitigates this.

Despite the rise in total savings, debt rose for some households. This results in an increase in the number of borrowers with high (above 40%) debt service ratios (DSRs). The majority of the increase in debt is concentrated among middle and higher-income homeowners with a mortgage for whom CERB only partially replaces lost income. Meanwhile, low earners experience an initial increase in debt that is reversed by the introduction of CERB payments.

Heterogeneous impact on household consumption

21 This is broadly consistent with the structure and rate of home equity line of credit (HELOCs) which are a common form of line of non-mortgage credit in Canada.

A variety of data sources from several countries has been used to study the impact of COVID-19 on household consumption over 2020, as discussed in section 1. Consistent with our findings in section 3, the decline in consumption across developed countries has been concentrated in goods and services that are hard-to-distance. Here, we connect those results to household income and finances through the simulation, as shown in figure 6 which depicts changes in non-shelter consumption for different income groups.

Households with higher pre-crisis earnings saw larger falls in consumption in response to the April lockdown, and higher earners' consumption recovered by less. This slower recovery is driven by the higher consumption of hard-to-distance goods and the relatively slower job recovery for high-income households that lost jobs (despite there being relatively few in number). This finding is similar to those of Chetty et al. (2020), who use US data and estimate that top quartile households accounted for 39% of the fall in total

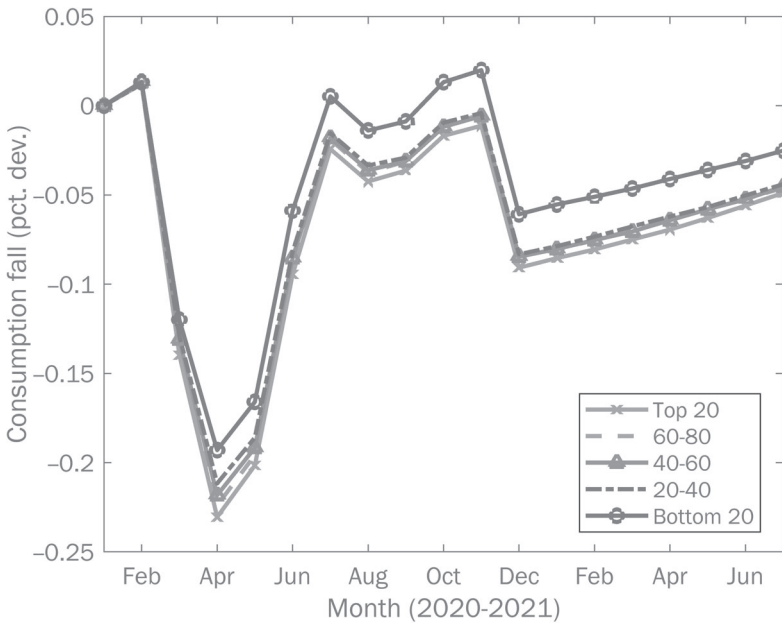


FIGURE 6 Consumption drop by income quintile

NOTES: Plotted are the percent deviation of simulated non-shelter consumption from its pre-crisis levels among employed households in the SFS. The sample is divided into quintiles of initial household income. Starting from the imputed consumption levels from the SHS, we apply the path of consumption adjustments shown in figure 1. While each spending class (essential, luxury, hard-to-distance, shelter) is subject to the same path of adjustment factors, the varying paths of consumption drop across income quintiles reflect heterogeneous consumption bundles across income groups reported in table 2.

spending while the bottom quartile accounted for only 13%, whilst we find 33% and 17% for those groups in Canada.²²

The second wave of COVID-19 also saw a universal and rapid decline in consumption. However, the gap between high- and low-income groups was maintained because of continued disproportionate effects of lockdown regulations on higher-income groups, as indicated earlier from table 2 and figure 1.

Heterogeneous buildup of savings

We estimate the rise in savings *in excess* of what would have occurred without the restrictions on consumption and fiscal transfers adopted in response to the pandemic. Specifically, we construct excess savings as the difference between household savings in our baseline and those in a counterfactual economy where there were no pandemic-driven shocks to consumption and employment.

The left panel of figure 7 plots the average amount of excess savings accumulated by income quintile, conditioned on having positive excess savings. The buildup of savings is driven largely by high-income households. The top income quintile accounts for roughly 40% of the stock of excess savings in December 2020, with accumulated savings averaging roughly \$2,500. Nevertheless, each income quintile contains some households who accumulated savings. In the simulations, bottom-quintile households whose savings rose accumulated roughly \$450 in excess savings by the end of 2020. This arises from their relatively low pre-crisis spending levels as well as CERB transfers.

Why do high-income households account for a large fraction of excess savings despite larger falls in income? This somewhat paradoxical result is driven by two factors. First, the expenditure share of higher-income households on hard-to-distance goods and services such as air travel and restaurants is higher (see table 2). With our methodology, this implies a relatively larger decline in their consumption expenditures during the lockdown (see figure 6). The top income quintile observes an average consumption drop of close to 23% while the bottom quintile observes only a 19% drop at the depth of the crisis. Importantly, this gap is sustained throughout the crisis. Because a larger share of aggregate consumption is attributable to top earners, their decline in spending accounts for a larger share of aggregate consumption decline. During the depth of the crisis (April), around 27% of the total decline in consumption is attributable to the top quintile of earners, while the bottom quintile accounts for only 13%. In our model, this reduction in consumption results in significant excess savings for high-income groups from the restrictions on hard-to-distance consumption.

The rate at which savings accumulate reflects the severity of containment measures and the resulting declines in consumption and employment. For example, the rise in savings slowed following the relaxation of restrictions in

22 Our measure restricts the sample of households to those that indicate attachment to the labour force.

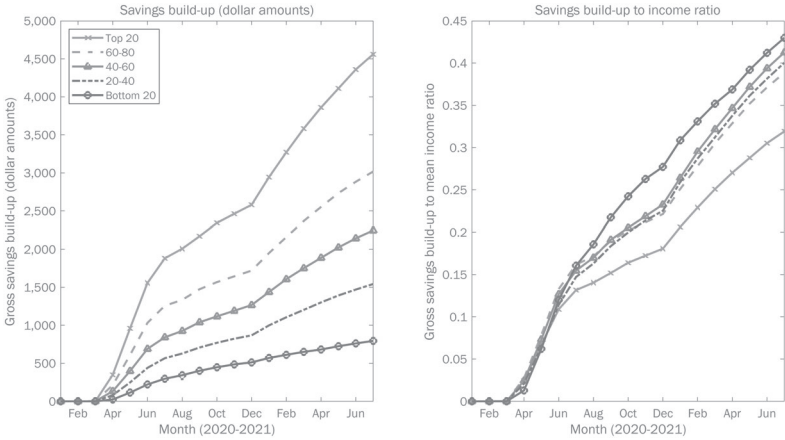


FIGURE 7 Gross savings by pre-COVID earnings quintiles
NOTES: This figure plots the simulated buildup of gross savings in dollar amounts (left panel) and as a fraction of average monthly income (right panel) among households that report to be employed in the Survey of Financial Security (SFS). The initial sample is first divided into quintiles of initial household income. For each group and time period, we condition on households that have positive period savings and report group averages. Reported values are net of savings that would have been accumulated under a no-crisis scenario in order to isolate savings buildup attributable to the COVID-19 crisis.

July 2020, but rose again during the second-wave lockdown that occurred around December 2020. Because of consumption baskets of higher-income households being more impacted by such measures, evidence of these changes in the slope of savings buildup are clearest for households in higher quintiles.

The decline in income after job loss is smaller (or not a fall) for the bottom 20% (pre-crisis) income group because of a higher take-up rate of CERB and a higher replacement rate. The right panel of figure 7 shows that while high-income households accounted for a larger fraction of the excess savings buildup, the size of savings buildup relative to income is larger among lower-income households. This implies that while low-income households do not account for a large fraction of aggregate savings, consumption restrictions and CERB transfers still resulted in a sizeable rise in their balances relative to income.

Heterogeneous buildup of debt

The earnings losses arising from unemployment and reduced hours drive an increase in debt for some households. To assess the implications of this rise in debt for the number of financially vulnerable households, we examine the ratio of debt payments to income, the debt service ratio (DSR).

Before discussing the distributional impacts, we examine the impact of deferrals and income changes on the aggregate DSR. With mortgage deferrals (the solid line in figure 8), the DSR declines because of the (temporarily) lower payments by deferrers. The recovery in income (with employment) after April

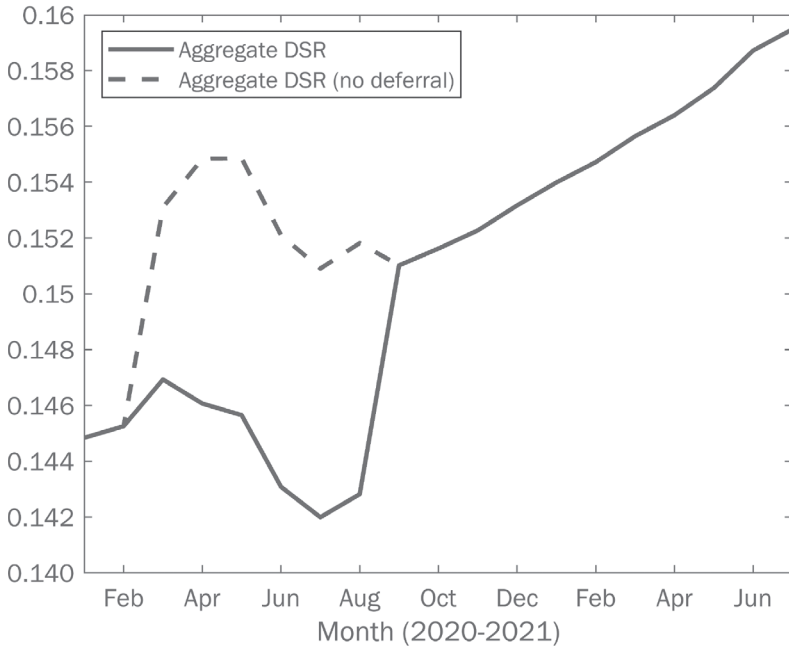


FIGURE 8 Aggregate debt service ratio

NOTES: This figure plots the aggregate debt service ratio (DSR), which is calculated as total monthly debt payments divided by total monthly income. The dashed lines represent a counterfactual path of DSR under a scenario without a mortgage deferral option.

pushes down the DSR until mortgage deferrals end in August. After deferrals end, the aggregate DSR begins to rise as households resume mortgage payments and face additional debt accumulated over the crisis. By the end of the simulation period, the aggregate DSR is 1.5 pp higher than pre-crisis.²³

To decompose the debt buildup, we group households by their pre-COVID earnings quintiles. We divide the change in debt by mean income in each quintile to highlight the group's shift in debt relative to income. Similar to our treatment of savings, we define increased debt as the difference between the debt in our baseline simulation and debt in the no-shock counterfactual. As can be seen from figure 9, middle income households account for the largest increase in debt. This reflects two forces. The first is that middle-income

²³ Our simulation abstracts from both an increase in debt for new housing purchases as well as the pass through of lower interest rates. We do not explicitly model the reasons behind mortgage deferrals because, even if a large fraction of deferrals arise from precautionary motives (not actual income loss), this would result in the same short-lived decline in aggregate DSR shown in figure 8. However, once deferrals end, the rise in DSR may not be as steep as in our simulation because precautionary deferrers are potentially in better capacity to resume payments than those who defer because of hardship.

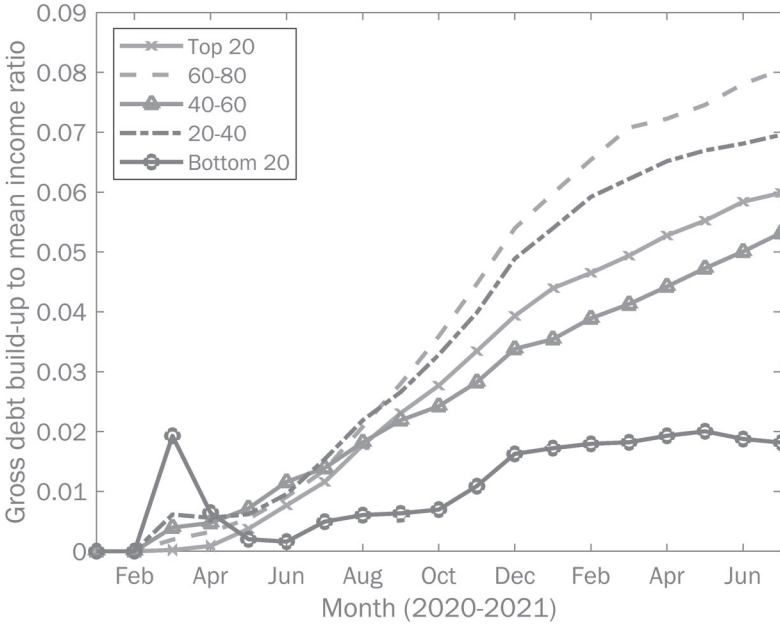


FIGURE 9 Gross increase in debt by pre-COVID earnings quintiles (employed households who only borrow)
NOTES: Plotted are the simulated rise in debt when employed households in the Survey of Financial Security (SFS) are grouped into quintiles of initial income. Debt buildup is the ratio of total debt to total pre-crisis monthly income of each group. To isolate the effect of the crisis, any debt buildup that would have occurred in a counterfactual no-crisis scenario is subtracted.

earners see only partial replacement of lost income from CERB. Second, some middle-income households have relatively large (compared with income) mortgage or rent payments and modest expenditure shares on hard-to-distance consumption. This results in smaller declines in expenditures than income, which drives a rise in debt for these households.

Households in the bottom income quintile experience an initial rise in debt during the first month of the crisis. This short-lived increase subsides as CERB payments are introduced and the labour market begins to recover. However, the expiration of CERB benefits and the second round of employment losses lead to a resurgence of debt for lower-quintile households. Households in the top quintiles experience a relatively slower buildup of debt during the first phase of the crisis before July 2020 because of the large drop in hard-to-distance spending and relatively lower exposure to job loss. However, the adjustments in consumption spending towards luxury goods and away from hard-to-distance goods is particularly pronounced for higher-income households. As a result, the build of up debt rises more rapidly afterwards.

Whilst informative, the average DSR does not provide information on the demographics of borrowers from financially vulnerable households. We follow

Faruqui (2008) and define a borrower as financially vulnerable if their debt service exceeds 40% of income. In figure 10, we plot the DSR distribution pre-COVID, in April and in October. In our simulations, there is a significant and sustained rise in the fraction of households with high debt service ratios. The fraction of households with DSR above 40% rises by roughly 1.2 pp by April. This rise in financially vulnerable households persists through late 2020 despite the bounceback in employment in our simulations.

This rise in financially vulnerable households is attributable largely to renters and homeowners with mortgages. Further decomposing these households by income in table 4 reveals an interesting pattern. The fraction of mortgagors

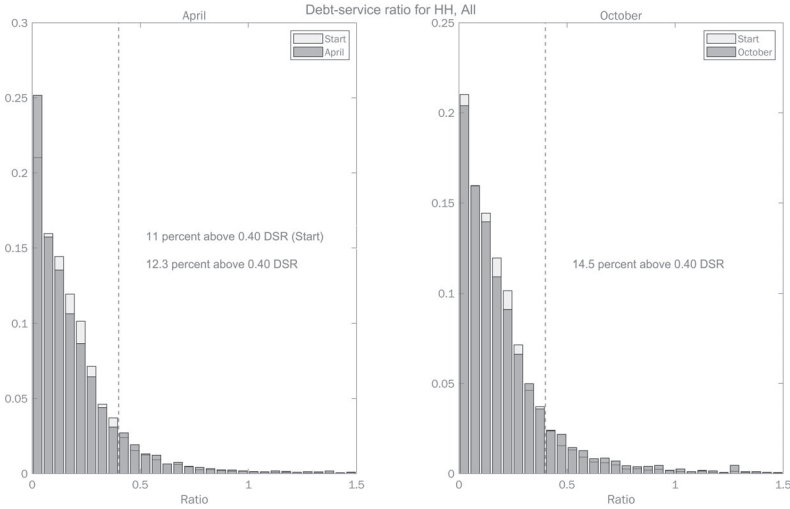


FIGURE 10 Distribution of DSR, all households

NOTES: This figure compares the distribution of debt service ratio (DSR) across households prior to the crisis with the simulated distribution in April and October. We consider a household to be financially vulnerable if DSR is above 0.40.

TABLE 4

Fraction of households with debt service ratio > 0.40, in percent

Household type	Initial (Feb.)	Trough (April)	Recovery (Oct.)
All	11.0	12.3	14.5
Renter (low Y)	4.9	5.2	8.8
Renter (high Y)	1.6	4.7	2.9
Owner, with mortgage (low Y)	32.1	30.6	36.6
Owner, with mortgage (high Y)	7.2	12.4	12.4
Owner, without mortgage	4.5	5.7	5.6

NOTES: This table reports the fraction of households with a debt service-to-income ratio above 40%. Renters and owners with mortgages are subdivided into those with incomes above and below the group-specific median. Values are in percent.

with low income and high DSR declines at the onset of the crisis (1.5 pp, from 32.1% to 30.6%) following the introduction of CERB. In contrast, the fraction of high-income mortgagors with high DSR rises from 7.2% to 12.4% as high earners who lose their jobs have a small fraction of lost income replaced by CERB. However, after CERB ends in September, low-income mortgagors experience a sustained rise in financial vulnerability, while high-income mortgagors remain close to levels seen in April 2020. Renters also exhibit heterogeneity in financial vulnerability across income groups because both low- and high-income renters see an increase in the fraction with high DSR, although this effect is larger and more persistent for low-income renters.

Mortgage deferrals: A modest slope, not a cliff

The early stages of the COVID-19 crisis saw substantial debate over the likelihood of a rise in mortgage defaults after the “mortgage deferral cliff” of expiring deferrals (e.g., Siddall, E. 2016). As we discuss below, our approach points to a modest rise in mortgage defaults as deferrals end, rather than a “cliff.”

The likelihood of a “mortgage deferral cliff” after the mortgage deferral window closes depends on the incidence of the so-called double trigger. The first trigger involves the persistence of elevated levels of unemployment, rendering some mortgagors unable to resume payments on their mortgages. The second involves low (or negative) levels of home equity among those who defer mortgages, which may be exacerbated if home prices decline. Using our simulation and LTV and wealth information from the SFS, we ask whether the incidence of these triggers were a cause for concern.

Figure 11 plots the unemployment rate of mortgagors and mortgagors who deferred payments.²⁴ By early 2021, the COVID-adjusted unemployment rate of mortgagors who defer is projected to decline to 15% after peaking at close to 50% in April 2020. Moreover, the fraction of unemployed households with high LTVs (above 80%) remains low (see table A2). This holds even in a scenario where house prices decline by 10%. The low incidence of double triggers predicted by our approach thus suggests a modest increase in mortgage defaults after the end of deferrals.

7. Impact of unplanned savings and higher debt on consumption

The unprecedented magnitude of the COVID-19 shock and its heterogeneous impact on household balance sheets has led to debate over the potential for “pent-up demand” to drive consumption higher after the relaxation of social

24 Our calibration implies that mortgage deferrers have low liquid wealth, high DSR and high LTV. See Allen et al. (2021) for an examination of which households selected deferrals.

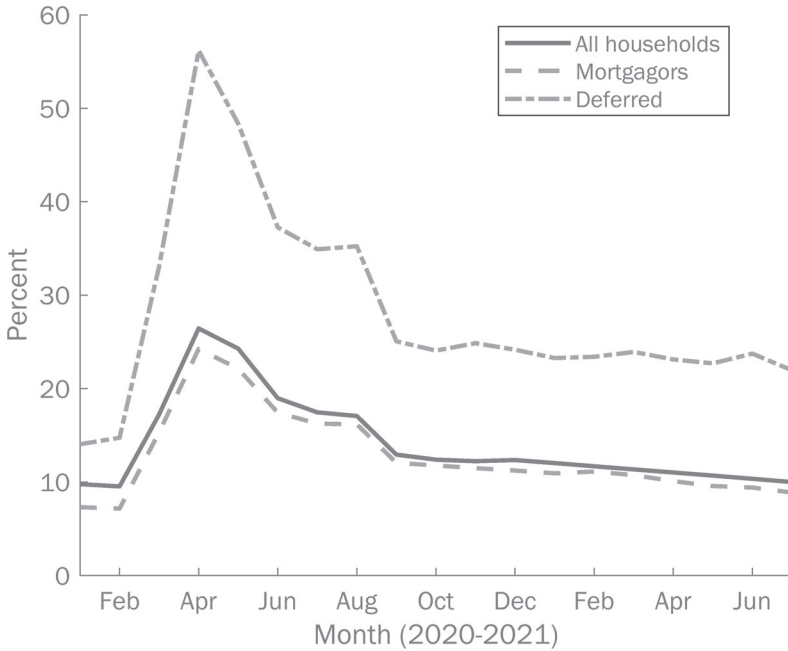


FIGURE 11 Unemployment rate of mortgagors
NOTES: Plotted are the simulated unemployment rates for all households, all mortgagors and mortgagors who defer. Mortgage deferrals are triggered by low liquid wealth and high loan-to-value ratio (LTV). Because the incidence of unemployment affects both metrics, mortgagors who defer have a higher incidence of unemployment.

distancing restrictions (e.g., Deloitte 2020). In large part, this debate reflects a range of views on how the accumulation of unplanned or precautionary savings by some households will impact spending as well as the magnitude of higher debt payments by households whose debt rose as a result of the crisis.

The potential size of these effects is large. Based on our benchmark exercise in section 3, by February 2021, total accumulated “lost” hard-to-distance spending (relative to a non-COVID counterfactual) had reached 70% of monthly consumption expenditures. Using our baseline model’s scenario, if households were to recoup this lost spending over six months following the (approximate) general population vaccination programme start date of September 2021, it would imply a 15% rise in monthly consumption.²⁵

²⁵ If one subtracts the observed rise in excess luxury spending from lost hard-to-distance spending, then the monthly consumption rise would be 9%. We choose September for illustrative purposes. Because the majority of the lost spending occurred during the 2020 lockdowns, varying the starting date for the rise in spending does not have a large impact on the amount of missing spending to be made up.

TABLE 5

Norwegian lottery spending proportion by liquid deposits and prize size

Magnitude of prize quartiles Liquid deposit group	0–3,299	3,300–8,242	8,243–13,798	13,799+
Bottom 25	1.047	0.745	0.720	0.490
25–50	0.762	0.640	0.559	0.437
50–75	0.663	0.546	0.390	0.386
Top 25	0.354	0.325	0.242	0.216

NOTES: Obtained from Fagereng et al. (2021); prize values converted to 2016 CAD. Deposit groups non-inclusive at lower bound.

Given the limited direct empirical evidence on the impact of a buildup of excess savings or debt as a result of consumption restrictions, we develop a novel counterfactual to tackle the question of post-vaccine consumption dynamics. Key to our approach are two assumptions. Firstly, for households whose debt rose that spending adjusts dollar for dollar with the rise in debt payments implied by higher debt. We construct these debt payments by imposing an amortization period of 15 years at the household's current estimated line of credit interest rate.²⁶

Second, we assume that households treat excess savings as akin to lottery winnings. This allows us to draw on Fagereng et al. (2021), who estimate how the marginal propensity to consume varies with the size of the lottery winnings, income and liquid asset positions. For example, table 5 shows that individuals at the bottom quartile of liquid assets spend close to the entirety of a lottery winning below \$3,300 within the first year after receipt but those in the top quartile spend only around 35% of this. Furthermore, as the lottery size grows, this MPC schedule shifts lower.

The intuition behind our approach is grounded in the idea that household post-pandemic income risk reverts to pre-pandemic levels and consumption behaviour normalizes. In this case, the source for the excess wealth accumulated, be it precautionary saving, consumption restrictions or a lottery win, does not affect forward-looking risks, so should not influence a rational agent's spending decisions.²⁷ Furthermore, our model's cash flow-based exercise is motivated by the large fraction of excess savings in the form of liquid assets, as evidenced by the large increase in M1+ and M1++ monetary aggregates measures in 2020.²⁸

26 Reducing the amortization period increases the level of drag on consumption but decreases the horizon which it affects (and vice versa).

27 This reasoning is consistent with the dynamics of buffer-stock saving in incomplete market models such as Bewley (1972), Huggett (1993) and Aiyagari (1994) or in models with default such as Livshits et al. (2007).

28 See www.bankofcanada.ca/rates/indicators/key-variables/monetary-aggregates/ for data on the evolution of various measures of monetary aggregates over the crisis.

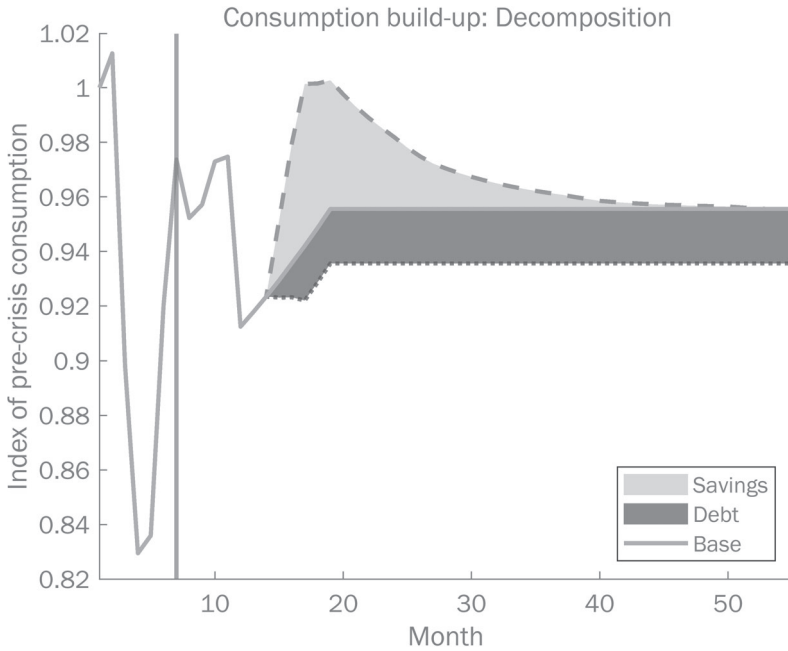


FIGURE 12 Impact of higher savings and debt on consumption

NOTES: This figure plots the simulated dynamics of consumption when accounting for: (i) the potential spending of unanticipated savings accumulated by households during the crisis (light shaded area) and (ii) the additional debt service obligations resulting from debt buildup (dark shaded area).

Our thought experiment assumes households stochastically switch from our baseline scenario to a “post-pandemic” state between March to May of 2021.²⁹ The implied drag on consumption from households with a higher debt burden is approximately 2.1% of pre-crisis monthly consumption. This is captured in the dark shaded area in figure 12, where the solid line is our baseline consumption path. The light shaded area shows the larger initial but less-persistent rise in consumption due to the spending of accumulated savings.

Given the large rise in savings in our baseline model, it is not surprising that the pent-up demand channel is initially much larger than the debt drag channel. The peak effect of the rise in savings is 5.9% of aggregate monthly consumption versus the 2.1% drag from households whose debt rose. However, the increase in expenditures due to unplanned savings dissipates rapidly. This follows from Fagereng et al. (2021) finding that small lottery winnings are mostly spent while large prizes are not, and households with low liquid assets

²⁹ Households draw from a uniform distribution of switching dates in this window. The selected dates are illustrative and not intended to provide precise statements about when consumption restrictions are lifted.

spend a larger fraction of winnings. The excess savings are large in aggregate but are largely held by higher-income and wealthy households. Thus, the size of these individual accumulations (figure 7) are small enough that marginal propensities to consume are relatively large (e.g., average excess savings for the highest income quintile is \$3,000 by December 2020, suggesting significant numbers of high-income households are in the 0–3,300 category in table 5).

8. Conclusion

The impact of COVID-19 on the Canadian economy, much like in other countries, is heterogeneous and severe. The combination of government-mandated lock-downs and voluntary social distancing by consumers has shifted consumption expenditures across goods and time. To capture this shock, we model a labour market with COVID-related hours reductions, non-participation and absences; separate goods and services into categories with different social/physical distancing characteristics and discretionary/essential status and include major policy interventions (e.g., CERB, mortgage deferrals) in a scenario analysis that quantifies the changes in debt and savings.

We find that the lowest quintile of earners are cushioned by the widespread CERB payments. Combined with their relatively large expenditures on easy-to-distance essentials, their consumption expenditures decline modestly. However, low- to middle-earning households see a faster and larger increase in debt than the bottom quintile; because their spending and rent/mortgage payments are not fully covered by CERB, they experience reductions in income.

High earners and older households with relatively large expenditures on hard-to-distance and luxury goods saw a larger decline in consumption expenditures during the spring and winter lockdowns. Combined with a modest rise in unemployment among high earners, this resulted in higher savings. Our analysis indicates this “excess” saving (“pent-up demand”) implies a substantial upside to future consumption. This effect is countered by the higher debt service payments for some households. While overall debt does not rise much, the number of high-debt households increases modestly.

Our work points to several directions for future research. First, our modelling does not include household expectations, risk aversion, or smoothing desires, except implicitly from the use of the survey data, nor developments in housing market during the pandemic. Secondly, the substitution between hard-to-distance, luxury and easy-to-distance goods is important for the path of future consumption. Finally, fiscal programs have limited debt rises and, in a loose sense, vulnerabilities among lower-earning households. However, we identify that low- to middle-income earners may be a future concern for policy-makers.

Appendix: Additional figures and tables

In sections 2 and 3, we discussed the procedure we develop to construct paths for different consumption spending classes (essential, hard-to-distance, luxury/easy-to-distance and shelter). A part of this procedure involves assigning spending class weights to the SHS and National Accounts consumption categories as well as mapping consumption categories between the National Accounts and high-frequency spending data. Because of the large number of categories, we have published them online in the following repository: <https://drive.google.com/drive/folders/10B8LSMrAD3LHTAj30aojvItjE74bubFC?usp=sharing>.

Here, we show three figures not held in the main text. Figure A1 displays the transition probabilities underlying the calibration of the model to unemployment and reduced hours statistics. Table A1 depicts the terms and

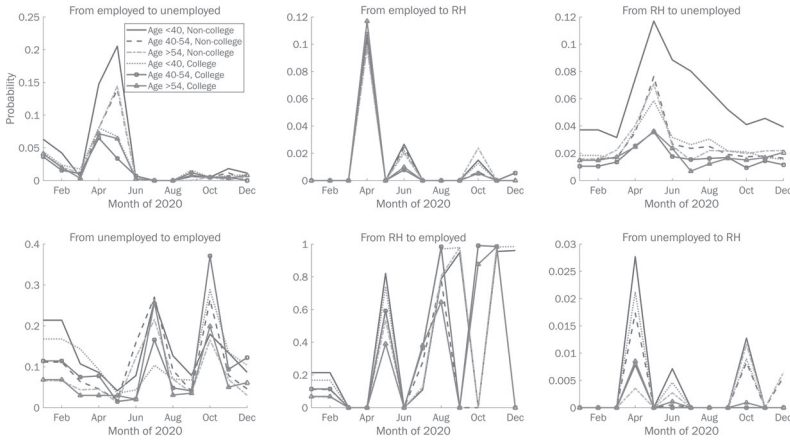


FIGURE A1 Employment transition probabilities

NOTE: This figure plots the probability of a move between different employment statuses in different periods, as calibrated in our model.

TABLE A1

Interest rates and terms of debt categories

Debt category	Interest rate (annualized)	Term (years)
Line of credit	Mortgage rate + 3%	15
Credit card	15%	13
Instalment	2%	2
Vehicle	2%	3
Student	1%	10
Other	8%	7

NOTE: This table presents the assumed interest rates and term lengths of various debt categories in order to derive an amortization schedule for each agent in the simulation procedure.

TABLE A2

Loan-to-value ratios (LTVs) of mortgagors

Percent of mortgagors with > 0.80 LTV by August	Unemployed	Unemployed + deferred mortgage
Baseline	1.9	0.5
House price decline (10%)	3.8	1.1

NOTE: LTV is calculated as the ratio of debt on principal residence and the value of the principal residence.

interest rate assumptions used for the model simulation, whilst table A2 shows the proportion of households with a mortgage who become unemployed, or unemployed and defer mortgages payments, and reach a loan-to-value ratio of over 80% by August.

Supporting information

Supplementary material accompanies the online version of this article.

References

- Achou, B., D. Boisclair, P. d’Astous, R. Fonseca, F. Glenze, and P.-C. Michaud (2020) “The early impact of the COVID-19 pandemic on household finances in Quebec,” *Canadian Public Policy* 46, S217–S235
- Aiyagari, S.R. (1994) “Uninsured idiosyncratic risk and aggregate saving,” *Quarterly Journal of Economics* 109, 659–84
- Allen, J., R. Clark, S. Li, and N. Vincent (2021) “Debt-relief programs and money left on the table: Evidence from Canada’s response to COVID-19,” mimeo
- Baker, S.R., R. A. Farrokhnia, S. Meyer, M. Pagel, and C. Yannelis (2020) “Income, liquidity, and the consumption response to the 2020 economic stimulus payments,” NBER working paper no. 27097
- Bewley, T.F. (1972) “Existence of equilibria in economies with infinitely many commodities,” *Journal of Economic Theory* 4(3), 514–40
- Bilyk, O., A. T. Y. Ho, M. Khan, and G. Vallée (2020) “Household indebtedness risks in the wake of COVID-19,” Bank of Canada Staff Analytical Note, no. 2020-8
- Carroll, C.D., E. Crawley, J. Slacalek, and M. N. White (2020) “Modeling the consumption response to the Cares Act,” NBER working paper no. 27876
- Chetty, R., J. N. Friedman, N. Hendren, M. Stepner, and O. I. Team (2020) “The economic impacts of COVID-19: Evidence from a new public database built using private sector data,” NBER working paper no. 27431
- Christelis, D., D. Georgarakos, and T. Jappelli (2015) “Wealth shocks, unemployment shocks and consumption in the wake of the Great Recession,” *Journal of Monetary Economics* 72, 21–41
- Coibion, O., Y. Gorodnichenko, and M. Weber (2020) “The cost of the COVID-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending,” NBER working paper no. 27141

- Cox, N., P. Ganong, P. Noel, J. Vavra, A. Wong, D. Farrell, and F. Greig (2020) "Initial impacts of the pandemic on consumer behavior: Evidence from linked income, spending, and savings data," *Brookings Papers on Economic Activity*, forthcoming
- Deloitte (2020) "COVID-19: Voice of Canadians and impact to retailers." Available at www2.deloitte.com/ca/en/pages/finance/articles/covid-19-voice-of-canadians-and-impact-to-retailers.html
- Fagereng, A., M. B. Holm, and G. J. Natvik (2021) "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, forthcoming
- Faruqui, U. (2008) "Indebtedness and the household financial health: An examination of the Canadian debt service ratio distribution," Bank of Canada staff working paper no. 2008-46
- Gellatly, G., and E. Richards (2019) "Indebtedness and wealth among Canadian households," *Economic Insights* (Statistics Canada), 11-626-X, no. 89. Available at www150.statcan.gc.ca/n1/pub/11-626-x/11-626-x2019003-eng.htm
- Hacıoğlu-Hoke, S., D. R. Känzig, and P. Surico (2021) "The distributional impact of the pandemic," *European Economic Review* 134, 103680. <https://doi.org/10.1016/j.euroecorev.2021.103680>
- Hackethal, A., and A. Weber (2020) "Fiscal policies and household consumption during the COVID-19 pandemic: A review of early evidence," SAFE white paper, no. 76
- Huggett, M. (1993) "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of economic Dynamics and Control* 17(5–6), 953–69
- Jappelli, T., and L. Pistaferri (2010) "The consumption response to income changes," *Annual Review of Economics* 2, 479–506
- Kaplan, G., B. Moll, and G. L. Violante (2020) "The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the U.S.," NBER working paper no. 27794
- Kaplan, G., G. L. Violante, and J. Weidner (2014) "The wealthy hand-to-mouth," *Brookings Papers on Economic Activity* 45(1 Spring), 77–153
- Lemieux, T., K. Milligan, T. Schirle, and M. Skuterud (2020) "Initial impacts of the COVID-19 pandemic on the Canadian labour market," *Canadian Public Policy* 46, S55–S65
- Livshits, I., J. MacGee, and M. Tertilt (2007) "Consumer bankruptcy: A fresh start," *American Economic Review* 97(1), 402–18
- MacGee, J. (2012) "The rise in consumer bankruptcy and consumer credit: Cause for concern?," Commentary No. 356, C.D. Howe Institute
- Norris, S., and K. Pendakur (2015) "Consumption inequality in Canada, 1997 to 2009," *Canadian Journal of Economics* 48, 773–92
- Siddall, E. (2016) "Supporting financial stability during the COVID-19 pandemic – Speaking notes for the standing committee on finance (05/19/2021)." Available at www.cmhc-schl.gc.ca/en/media-newsroom/speeches/2020/supporting-financial-stability-during-covid19-pandemic
- Statistics Canada (2016) "Survey of financial security: Public-use microdata file, 2016," (13m0006x2018001). Available at www150.statcan.gc.ca/n1/en/catalogue/13M0006X2018001

- (2017) “Survey of household spending: Public-use microdata file” (62m0004x). <https://doi.org/10.25318/62m0004x-eng>
- (2020a) “Distribution of household economic accounts.” Available at www150.statcan.gc.ca/n1/pub/13-607-x/2016001/858-eng.htm
- (2020b) “Labour Force Survey: Public-use microdata file” (71m0001x). <https://doi.org/10.25318/71m0001x-eng>
- (2020c) “National balance sheet and financial flow accounts, first quarter 2020.” *The Daily*. Available at www150.statcan.gc.ca/n1/daily-quotidien/200612/dq200612a-eng.htm
- (2021) “Household economic well-being during the COVID-19 pandemic, experimental estimates, first quarter to third quarter of 2020.” *The Daily*. Available at www150.statcan.gc.ca/n1/daily-quotidien/210301/dq210301b-eng.htm
- van Buuren, S., and K. Groothuis-Oudshoorn (2011) “mice: Multivariate imputation by chained equations in R,” *Journal of Statistical Software* 45(3), 1–67