
UNE MESURE DE LA RÉPONSE EN CONSOMMATION À DES CHOCS DE REVENUS À PARTIR DES DONNÉES BANCAIRES

Odran Bonnet (), François Le Grand (**), Tom Olivia (*), Xavier Ragot (***), et
Lionel Wilner (*)*

() Insee, Département des Études Économiques*

*(**) EM-Lyon*

*(***) CNRS / OFCE / Sciences-Po*

odran.bonnet@insee.fr

legrand@em-lyon.com

tom.olivia@insee.fr

xavier.ragot@gmail.com

lionel.wilner@insee.fr

Mots-clés. Données bancaires, données hautes-fréquences, économétrie appliquée, évaluation des politiques publiques, micro-économie

Domaines concernés 7.2 ; 7.3

Résumé

La littérature empirique sur le lien entre revenus, consommation et épargne a bénéficié de l'accès récent aux données bancaires dans des pays comme les USA (Baker, 2018, Gelman, 2021b) ou le Danemark (Andersen et al., 2020). En France, de telles données à haute fréquence peuvent désormais être mobilisées afin d'analyser, par exemple, l'impact de la crise sanitaire et des dispositifs de lutte contre la pandémie de Covid-19 sur la consommation et l'épargne (Bounie et al., 2020, Insee, 2021).

L'objectif est de mieux comprendre l'impact de variations de revenu sur la consommation des ménages, en portant une attention particulière aux spécificités françaises mises en regard des éléments empiriques documentés dans d'autres pays. Des analyses exploitant les données bancaires du Crédit Mutuel Alliance Fédérale comme de La Banque Postale permettent en effet d'identifier les réponses en dépenses de consommation, suite à un choc de revenus, sur un panel de 300 000 clients environ, suivis de 2019 à aujourd'hui. Différents types de chocs seront considérés : (i) des chocs transitoires (versement de l'Allocation de Rentrée Scolaire (ARS), de différentes aides exceptionnelles, de la prime de Noël ou encore perception d'un héritage), (ii) des chocs de moyen terme (entrée ou sortie du chômage) ; (iii) des chocs permanents (hausse de salaire, augmentation du Smic, de l'Allocation Adulte Handicapé (AAH)). La méthode retenue consiste à exploiter la dimension longitudinale de ces données pour identifier la réaction, à un niveau fin, des dépenses de consommation à ces chocs de revenus. Elle permet notamment d'estimer des propensions marginales à consommer (MPC) ou des élasticités-revenus de la consommation qui peuvent différer selon les catégories de ménages considérées. L'accent sera ainsi mis sur l'hétérogénéité de ces réactions selon différentes dimensions pertinentes comme le revenu lui-même, mais aussi le patrimoine liquide, le patrimoine illiquide, etc. Cette approche est susceptible de contribuer à déterminer, de manière plus normative, des mécanismes de redistribution appropriés pouvant être opérés par des transferts sociaux ou fiscaux.

Les données bancaires sont une source nouvelle riche en information pour la statistique publique. Elles comportent néanmoins un certain nombre de limites. Notamment, les variables disponibles ne correspondent pas aux concepts que souhaiterait mobiliser le statisticien. Par exemple, seuls les flux entrants sur les comptes sont observés et non le revenu disponible du ménage. De plus, seuls les avoirs sur les comptes d'une unique banque sont observés et non le patrimoine total du ménage. Cette étude s'attachera à décrire les enjeux méthodologiques que posent l'utilisation de telles données et décrira des solutions pour répondre à certains d'entre eux.

Abstract

This paper exploits high-frequency financial data to measure how spending responds to a variety of income shocks: direct transfers, predetermined payments, revaluations of social benefits, job losses, wage increases. Based on individual transaction-level data from France, we estimate marginal propensities to consume out of these shocks. Spending response to transitory shocks displays substantial heterogeneity with respect to liquidity, consistently with incomplete market models. Yet behavior-based explanations are not ruled out: asymmetrical responses arise with respect to the sign of persistent shocks. Surprisingly, there is no notable difference in responses to transitory or permanent shocks.

Keywords: Marginal propensity to consume; excess sensitivity of consumption; bank account data.

JEL Classification: C55; D12; D15; H50.

1 Introduction

How do consumers react to income shocks? This classical question in economics has received a lot of attention in both theoretical and empirical literatures, notably because its answer shapes the design of both fiscal and monetary policy. From a normative point of view, characterizing spending response to income is key to determine the optimal targeting of social policies, i.e. which households should receive direct transfers when necessary, including the amount of stimulus payments during major crises such as the Great Recession or the Covid-19 pandemic.

Despite a now substantial evidence on this topic, estimates of the marginal propensity to consume (MPC) out of income turn out to vary considerably depending on the institutional setting, the time period considered, the types of individuals concerned, among others (Johnson, Parker, and Souleles, 2006; Agarwal, Liu, and Souleles, 2007; Gelman et al., 2014; Gelman, 2021b). Hence the need for up-to-date information on how sensitive spending is with respect to income is perennial, both from a policy viewpoint and from an academic perspective. Researchers have recently benefited from the availability of financial account data in countries like the U.S., Denmark, France, or Spain, which permits to dispose of accurate and contemporaneous (almost real-time) estimates.¹ Taking advantage of the observation of spending at the transaction level, these new sources allow to track each individual's reaction to income shocks over time at a fine granular level, which facilitates the identification of the elasticity of spending to income, in comparison with cross-sectional or low-frequency longitudinal data.² By definition, when that information is available -at most- annually, researchers are not able to estimate short-run, infra-annual responses. In contrast, high-frequency bank account data render unnecessary a modelling of the earnings process and corresponding statistical decompositions usually required in order to disentangle transitory from permanent components of income (as in, e.g., Blundell, Pistaferri, and Preston, 2008) so as to assess the reaction of spending with respect to each and any of these shocks. Last, on top of making the applied researcher's life easier, it provides with a comprehensive picture of the heterogeneity in spending response along various dimensions (income, assets, liquidity, debt, on top of usual socio-demographics) since these datasets contain detailed information not only on transactions, but also on balance sheets, which allows to characterize savings (to some extent), and the corresponding reactions along this margin. However, as emphasized by Baker, 2018, resorting to such data raises at least two concerns that deserve careful attention: representativeness (by

¹The research in this domain has been pioneered by Gelman et al., 2014, and now includes Baker, 2018, Olafsson and Pagel, 2018, Andersen et al., 2020, Carvalho et al., 2020.

²Examples of such databases include the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX) in the U.S., the British Household Panel Survey (BHPS, now Understanding Society) in the UK, the German Socio-Economic Panel (GSOEP) in Germany, and *Budget des Familles* (BDF) and *Enquête sur les Revenus Fiscaux et Sociaux* (ERFS) surveys in France, not claiming here to be exhaustive.

construction, surveys do not suffer from this problem) and completeness (assets may be disseminated into different financial institutions).

The contribution of this paper is to estimate spending response to income shocks based on a large spectrum of such shocks: transitory shocks (direct transfers and predetermined payments) and permanent shocks (revaluations of social benefits, job losses, and other persistent changes in earnings like wage increases). This approach contrasts with most of the literature focusing either (i) on unanticipated, temporary shocks,³ (ii) on unpredictable, persistent shocks like unemployment or disability,⁴ (iii) on anticipated, one-shot payments,⁵ or (iv) on predictable, recurrent payments.⁶ To the best of our knowledge, very few papers study various types of shocks based on the same dataset. We believe that such an approach enables us to test convincingly whether spending reacts more to permanent than to transitory shocks, and whether it responds to anticipated shocks -or not, as the theory would suggest. These findings have direct implications as regards the empirical relevance of the life-cycle model and the permanent income hypothesis (PIH), as well as of deviations from this standard framework (liquidity constraints and incomplete market models, behavioral explanations based on individual failures).

Our results are the following. First, we find average marginal propensities to consume out of income that fall typically within the range of most recent estimates documented in the U.S. or in Denmark, and based on similar transaction data (Baker, 2018; Andersen et al., 2021). Second, excess sensitivity of consumption exhibits substantial heterogeneity with respect to individual characteristics, especially liquid assets (resp. income) as far as transitory (resp. permanent) shocks are concerned; again, this empirical evidence is consistent with previous findings obtained in other institutional settings (Gelman, 2021a). On top of being immediately useful to the policy maker in charge of designing fiscal and social policies, these results are consistent with deviations from the standard model (LC-PIH) due to liquidity constraints, and with incomplete market models making a distinction between liquid and illiquid assets (Kaplan and Violante, 2014). Third, we find some asymmetry of spending response toward negative or positive permanent shocks: consumers smooth more when they are exposed to some persistent loss than when they experience some permanent income gain. Hence behavior-based explanations, possibly related to loss aversion, are also at stake, on top of liquidity constraints. Last, we find rather homogeneous responses

³usually based on lotteries as in Fagereng, Holm, and Natvik, 2021, on experimental designs like in Jappelli and Pistaferri, 2014 and Bunn et al., 2018, exploiting surprise, one-time fiscal cash payouts like in Agarwal and Qian, 2014, or exogenous variation in the price of specific consumption products: gasoline (Gelman et al., 2016), housing (Disney, Gathergood, and Henley, 2010) or equity assets (Andersen, Johannesen, and Sheridan, 2021).

⁴see Browning and Crossley, 2001, Ganong and Noel, 2019, Andersen et al., 2020, among others.

⁵such as tax rebates: Souleles, 1999, Johnson, Parker, and Souleles, 2006, Agarwal, Liu, and Souleles, 2007, Parker et al., 2013 or Misra and Surico, 2014, not claiming to be exhaustive either.

⁶including Olafsson and Pagel, 2018 and Gelman, 2021a devoted to the payday effect, as well as Kueng, 2018a who exploits a large, regular (but specific) payment from the Alaska Permanent Fund.

regardless of the nature of the shock considered (transitory as opposed to permanent). In particular, the reaction of spending does not look stronger consecutive to permanent, rather than to temporary shocks. This empirical evidence supports a rejection of the PIH, at least of the most parsimonious version of the life-cycle model. As a possible explanation rationalizing the data, it also points to the distinction between various types of shocks being neither sharp nor salient to consumers, from an empirical viewpoint.

The rest of the paper is organized as follows. Section 2 presents our data, including comparisons with external sources in order to answer both representativeness and completeness concerns. Section 3 presents our estimates of spending response out of transitory income shocks, namely direct transfers (the expansion of social benefits consecutive to the Covid-19 pandemic) and predetermined payments (the payday effect and back-to-school allowances). Section 4 is devoted to studying spending response out of permanent income shocks: negative shocks (job losses) and positive shocks (wage increases and revaluations of social benefits). Section 5 discusses the implications of our findings from both policy and academic prospects.

2 Data

In this section, we present our de-identified bank account datasets. The first database stems from the oldest French bank, the *Crédit industriel et commercial* (CIC hereafter), created in 1859 and now belonging to the *Crédit Mutuel Alliance Fédérale* group; this bank has about 4.2 million customers. The second dataset we use is issued from *La Banque Postale* (LBP hereafter), a public bank created in 2006 within the postal group *La Poste*, the historical monopoly in charge of mail delivery; this bank has nearly 11 million customers. The construction of key variables follows Baker, 2018, Ganong and Noel, 2019 and Andersen et al., 2021. We dispose of transaction-level data on debit card payments,⁷ paper checks, cash withdrawals, cash deposits, bank transfers, and direct debits; each transaction is characterized by its amount in euros. We have also access to balance sheets data, and thus to end-of-month balances on deposit and various savings accounts (*Livret A*, *Livret Jeune*, *Livret de Développement Durable et Solidaire*, *Compte Épargne Logement*, and time deposits) on top of life insurance and stocks. Deposit accounts include joint accounts. The unit of observation is the household (individual) in the CIC (LBP) dataset. Both datasets are high-frequency in that they contain transaction-level information, hence timestamped, though the aggregation at the daily level only is available to us; again, balances are available at a monthly frequency. Finally, we observe various socio-demographics: age, sex, *département*,⁸ marital status, occupation, dummies for urban/rural/peri-urban areas.

⁷In France, the use of credit cards is scarce: it accounts for less than 10% of bank cards.

⁸an administrative division like, e.g., the county in the U.S. Mainland France, i.e. France at the exclusion of both Corsica and overseas, is divided into 94 *départements*.

Spending We define total spending as the sum of outgoing transactions paid either by debit card or by cash withdrawals. Paper checks may be included as a robustness check.

Income We do not have access to transfers' labels: as a result, we do not know whether given incoming transfers correspond to salaries, pensions, benefits or private transfers. Hence we measure disposable income as the sum of monthly incoming transfers (up to €40,000), and at the exclusion of integers which likely correspond to private transfers. It is of course possible not to cap at the previous threshold, or/and to include integers: corresponding robustness checks are systematically performed in subsequent analyses.

Financial wealth Liquid assets are proxied by the sum of balances on different bank accounts (deposit account and savings accounts).

Illiquid assets are proxied the sum of balances on life insurance, stocks, bonds, mutual funds and certificates of deposits. In France, banks are not in charge of retirement savings plans.

Non-mortgage loans are also available in our databases, and we compute net repayments from these loans by considering the change in end-of-month balances on these accounts: positive (negative) values correspond to net repayments (borrowing).

Working sample Our observation period runs from February 2019 to November 2021. Our main initial raw data is a sample of about 300,000 households who bank primarily at CIC, this sample being stratified by *départements* of mainland France and by 5-year age dummies. In what follows, we restrict our attention to households with the same number of adults aged at least 18 over the period. We focus on customers who spend -either by debit card or in cash- and earn at least €150 on three rolling months. Moreover, we impose that customers be present and meet previous criteria all over the period, which leaves us with about 200,000 households who are active customers banking primarily at CIC.

Two concerns have been raised by the literature as regards external validity of bank account data (Baker, 2018): representativeness and completeness. We resort therefore to several external sources to assess both representativeness and completeness of our databases.

Representativeness To alleviate concerns about representativeness, and to build upon previous works previously mentioned, we proceed to calibration weighting using the method proposed by Deville and Särndal, 1992. We thus compute weights that reproduce exactly known population totals for auxiliary variables, while ensuring that these calibrated weights are as close as possible to original sampling weights. By construction, the

weighted sample has the same distribution as the population on the corresponding variables. As regards such auxiliary information, we consider the following dimensions (called margins): age \times sex and *département*. In what follows, all our estimating equations are weighted using calibration weights.

The distribution of household expenditures with respect to their position in the standard of living distribution obtained in transaction data matches closely the one issued from the representative consumption survey *Budget des Familles* (Figure 1); in particular, and both ends of the income distribution put aside, spending-to-income ratios look remarkably similar, decreasing from 1 to 0.75, which mitigates previous concerns on measurement error as regards income. If any, our data overestimate spending, probably because CIC customers tend to be wealthier. This is confirmed by Table 1 which suggests that CIC customers are wealthier: they dispose of higher income and spend more than the average (the reverse holds true for LBP customers). The pregnancy of liquidity constraints can be assessed by looking at the liquid wealth-to-income ratio, about 10, meaning that, on average, households dispose of liquidity equivalent to 10 months of income. It decomposes into a 3.5 ratio of liquid assets over end-of-month balances on deposit accounts (this number compares well with the one documented in the U.S. by Baker, 2018), and another 3.5 ratio of end-of-month balances on deposit accounts over monthly income. Finally, these customers are on average younger and tend to live in more peripheral areas.

Completeness First, the evolution of our measure of spending follows closely the one issued from the *Groupement des cartes bancaires CB*, the French national interbank network. Breaking down these variations by categories of spending (food at home, food away from home, clothing, drugstore, tobacco, household goods, accomodation expenditures) yields systematically to similar diagnoses.

Second, the dynamics of our measure of income is more volatile (Figure 2) than the one measured by Insee.⁹ This higher dispersion is rather expected: it is intrinsically related to the fact that we do not observe income directly, but rather all incoming transfers, and that there may be some measurement error here. Yet it is reassuring to see that the magnitude of such measurement error is limited, the differential amounting to about -10pp at the beginning of 2019 and to +5pp at the end of 2021 when taking the beginning of 2020 as the common reference.

Third, our measure of liquid assets is slightly more dynamic than the one reported by *Banque de France* that gathers information at the French level from all other bank networks (Figure 3). Still adopting the same normalization as above, the differential now amounts to less than +2pp at the end of 2021, which may reflect a composition effect: CIC customers likely benefit from higher capital gains (Fagereng et al., 2019).

⁹namely, gross disposable income per consumption unit, i.e. gross standard of living.

On the whole, previous comparisons with external sources suggest that representativeness is not too much of a concern, and that the calibration weighting contributes to alleviate this problem; and also that the remaining differences on earnings and assets are mostly due to differences in concepts, rather than to incompleteness.

3 Transitory income shocks

The first part of our empirical analysis is devoted to the spending response to transitory income shocks. On the one hand, we consider direct social transfers such as those consecutive to the Covid-19 pandemic targeted at the poor. On the other hand, we examine regular, predetermined payments: monthly paycheck arrivals and back-to-school allowances, i.e. family benefits, which are perceived every year by eligible individuals. In each of these cases, we determine whether excess sensitivity of spending to income is at stake, and we compare our results with the literature. Besides, we take a stand at these findings in order to draw global conclusions made possible by us exploiting various shocks of different nature.

3.1 Direct transfers

Faced to major economic downturns or to disaster events which are textbook cases for public intervention, states have been used to compensate for such market failures by supporting directly firms or/and households. Recent examples include the Great Recession in 2008 and the on-going Covid-19 pandemic since 2020. Consecutive to the latter, the French government decided to expand social benefits on May and November 2020 in an exceptional, unexpected fashion. This quasi-natural experiment provides with a transparent example of unanticipated, transitory income shock and we rely on these direct transfers to infer the MPC of concerned households.

3.1.1 The expansion of social benefits due to the Covid-19 pandemic

The strict lockdown arising in France from March 17th, 2020 to May 11th, 2020 made it clear that the poor were the most suffering from restrictions: first, through poor living conditions, among others, but also from the economic crisis, due for instance to the scarcity of job opportunities on the labor market, which prevented many unemployed from getting back to work. As a consequence, President Macron announced on April 15th, 2020 the provision of an exceptional €150 bonus at the destination of the beneficiaries of welfare benefits (mostly the *Revenu de Solidarité Active* or RSA), or other specific social benefits (the *Allocation de solidarité spécifique* or ASS, which is granted to unemployed people at

the exhaustion of their UI benefits, the *Revenu de solidarité* or RSO, some kind of RSA at the destination of overseas, etc.) including housing benefits (the so-called *Aide Personnalisée au Logement* or APL). The payment was made effective on May 15th, 2020. Before the second lockdown, which was less severe than the first one, the same decision was taken once again, with an extra €150 being announced on October 14th, 2020, and attributed on November 27th, 2020 (week 48) to individuals meeting the same criteria.

We base the subsequent analysis on the second payment in order to avoid possible confounding factors related to the strict lockdown, namely severe restrictions on spending. Many so called “non-essential” stores were closed (including clothing, haircuts, etc.), which led to a huge drop in spending. By contrast, during the second lockdown, restrictions were far less severe. We are able to identify the bonus in the data based on the day when it was available on deposit accounts. Our identification strategy consists in matching beneficiaries of this unexpected aid in 2020 (treated individuals) to similar individuals in 2019. Our working sample here is therefore composed of 5,441 treated and their 5,441 siblings. In practice, we keep individuals for whom more than 2/3 of incoming transfers correspond to integer amounts perceived on that day. We then regress weekly spending on week fixed effects as well as on their interactions with the treatment, the latter coefficients being provided by Figure 4. Standard errors are two-way clustered at individual and group levels in order to take the matching process (and the possibly resulting autocorrelation of residuals) into account, namely the fact that the same individuals may belong to both groups at different moments. The marginal propensity to consume (MPC) out of that aid is computed from the cumulated sum of estimated spending differences between the two groups from week 42, i.e. 7 weeks before treatment, to week 51, i.e. 4 weeks after treatment. This conservative choice is guided by our willingness of taking possible anticipations and corresponding changes in spending behavior into account; doing so enables us not to forget small pre-trend differences (visible on Figure 4), instead of spuriously attributing post-treatment differences in spending to the policy shock at stake. The MPC is finally obtained by regressing previous cumulated difference in weekly spending on the pecuniary amount of extra social benefits (€150 for the treated group as opposed to 0 for the control group). Our point estimate is 0.89 and has a standard error of 0.04. Since this population has a spending-to-income ratio that is quite close to 1,¹⁰ the estimated elasticity of spending to income is nearly the same, about 0.9. The heterogeneity of previous effects with respect to liquid assets is investigated by Figures 16 to 19 in Appendix. Moreover, we estimate separately treatment effects depending on the household’s family status since an extra aid of €100 per child was granted: see Figures 20 and 21 in Appendix.

¹⁰Remember Figure 1 showing that this ratio decreases from 1 at the lower end of households’ disposable income to 0.75 at the upper end of that distribution, first and last deciles put aside: measurement error is likely below the first decile and above the last decile of income.

3.2 Predetermined payments

This subsection studies how spending reacts to different sources of regular anticipated income, either on a monthly basis (paycheck arrival) or at an annual frequency (peculiar family benefits, namely back-to-school allowances, which are received by a sub-sample of eligible households just before school year).

3.2.1 The payday effect

We document first the existence of deviations from consumption smoothing upon the arrival of regular monthly payment on our data, i.e. what the literature has called the “payday effect” (Stephens Jr, 2003; Shapiro, 2005; Stephens Jr, 2006; Mastrobuoni and Weinberg, 2009; Hastings and Washington, 2010; Olafsson and Pagel, 2018; Gelman, 2021a; Kuchler and Pagel, 2021). Paycheck is the most salient, anticipated payment: many individuals receive this regular income, the main source of their labour earnings, on a monthly basis. Due to its timely recurrence, this event should come at no surprise. However, it has been shown that consumption is not perfectly smooth around that point in time, and that at least some individuals tend to spend much more on that day.

The granularity of our data enables us to estimate finely this effect at a daily level on both datasets (CM-CIC and LBP), and thus to quantify by how much spending increases in comparison with the average spending day of the month. The payday is directly inferred from the data, based on the highest incoming transfer received within a month, provided that this transfer is worth at least 80% of the sum of all incoming transfers in a 30-day rolling window centered around this day. In a conservative fashion, we consider customers whose highest two income transfers are separated from between 25 and 35 days, and do not differ from more than 10%; on top of that, the amount received should exceed €400; last, we keep observations that meet all these criteria during at least three consecutive months. On the CIC dataset, this selection leaves us with 258,520 observations, corresponding to 58,531 different customers.

Figure 5 depicts the daily evolution of the ratio of spending by individual i on day t over average daily spending, denoted by x_{it} in what follows -the latter average is computed on the basis of the 29 days around the payday. By construction, this ratio is equal to one, on average, for every individual during a month, but fluctuates from one day to another. A ratio higher (lower) than 1 means that current daily spending is higher (lower) than the average daily spending. A peak of about 45% is observed on payday, but also on the day after.¹¹ The pattern of spending decreases then sharply during the first week following payday. Last,

¹¹This looks like a specificity of our data compared with other studies devoted to the payday effect; the fact that this effect is diluted on both payday and the day immediately after could be due to information frictions.

a weekly cycle can be observed within a month, which is in line with paychecks arriving during weekdays only.¹² Interestingly, this effect concerns almost all categories of spending (Figure 6): durables (clothing, home improvements), strict non-durables like groceries, public transportation, sin goods (alcohol, tobacco), health or personal care expenditures exhibit a spike on payday, but so do expenditures on short-run or even instantaneous consumables like food away from home, fuel, or sports and activities, i.e. on categories that look closer to immediate consumption. As expected, strongest effects (higher than +40%) are found for food, clothing, alcohol and tobacco, personal care, and cash withdrawals, while lowest deviations from consumption smoothing arise on discretionary spending like restaurants and leisure; yet, even for the latter categories, the observed spike is still about +20% at least, which constitutes a substantial deviation from LC-PIH.

Turning now to an econometric specification which aims at controlling for potential observed and unobserved confounding factors, we follow Olafsson and Pagel, 2018 and specify:

$$x_{it} = \sum_{k=-7}^{+7} \beta_i^k \text{Paid}_{i,t+k} + \delta_{dow} + \phi_{wom} + \psi_m + \zeta_y + \eta_i + \varepsilon_{it} \quad (1)$$

The dummy $\text{Paid}_{i,t+k}$ indicates whether individual i gets her paycheck on day $t+k$; the coefficient of interest β_i^k measures the relative payday-specific deviation from average spending. A bunch of time fixed-effects (FE) are included: δ_{dow} accounts for a day-of-the-week FE and takes 7 possible values, ϕ_{wom} is a week-of-the-month FE and takes 5 values, ψ_m is a month FE (12 values), and ζ_y year FE (3 possible values). On top of that, individual FE η_i are included in the regression.

Consistently with previous descriptive evidence, the estimates point out to an average payday effect of +0.43 (Table 3), which falls to +0.37 (Table 5) after controlling for time FE (column 2) and further for individual FE (column 3), meaning that consumers spend on average 37% more the day when they receive their paycheck.¹³ From a statistical viewpoint, the impact is very precisely estimated and highly significant at usual levels. The economic significance is also meaningful: everything happens as if consumption expenditures during the payday were equivalent to about 1.5 usual day, which constitutes a substantial deviation from perfect smoothing, i.e. from the null hypothesis $H_0 : \beta_i^k = 0, \forall i, k$ which is strongly rejected. The magnitude of our effect lies in the range obtained by Olafsson and Pagel, 2018, while seemingly above estimations obtained in the U.S. or in the UK: Stephens Jr, 2003 observed that total expenditures increased by about 80% on payday, but that the

¹²The corresponding frequency is as follows: 6% on Mondays, 17% on Tuesdays, 34% on Wednesdays, 22% on Thursdays, and 21% on Fridays. The distribution of paydays within a month is displayed by Figure 15; a mode can be observed every 9th of the month, which corresponds to pensions being paid on that day.

¹³Replicating the same specification as Olafsson and Pagel, 2018 yields to a +0.46, see Table 2.

magnitude of the increase as regards food away from home, closer to immediate consumption, relative to average daily spending was 6 percent only; Shapiro, 2005 finds a 10-15% decline in caloric intake all over the month; in the UK, according to Stephens Jr, 2006, the increase in total expenditures is roughly 16% of the average total weekly spending, but restricting attention to items closer to immediate consumption yields lower estimates (about 5%); Gelman, 2021a concludes to similar effects on food expenditures, about 5.5% weekly, but he focuses on individuals being paid bi-weekly, while Hastings and Washington, 2010 find evidence of a sharp decrease in food expenditures of more than 20 percent from the first to the second week of the month among households that are eligible to food stamps. When aggregating our data at the weekly level though, we estimate an effect of about +0.23 (Table 5), i.e. +23% on a payweek, which compares well to previous findings -the corresponding decline all over the month being then of about -19%.

Many factors may rationalize this empirical evidence, among which binding borrowing constraints and behavioral explanations (myopia, hyperbolic discounting, self-control problems, etc.). To disentangle these explanations, we investigate how heterogeneous the effect is, and along which dimensions: age, marital status, liabilities, income, liquid and illiquid wealth (see Figure 7 and Table 4). The effect is much higher among single parents, but much lower in the presence of mortgage loans; it is also lower for older individuals and for couples. Again, these results are obtained with a high statistical significance. Gradients with respect to both income and financial assets are strong; the latter is even stronger, though. Moreover, among financial assets, illiquid wealth plays a minor role in these deviations from consumption smoothing compared with liquid wealth. To illustrate, we estimate $\beta_i^k = 0.64$ for low-income individuals, i.e. individuals in the lowest quartile of the income distribution: for them, a payday is equivalent to less than two usual spending days. The corresponding figures for individuals belonging to the lowest quartile of the financial wealth distribution is 1.03 -hence a payday is worth slightly more than two usual spending days. Compared with income, liquid assets segment more the population in this respect. Also, even at the top of these distributions, statistically significant deviations from smoothing are found: 0.04 (liquid wealth) and 0.14 (income), which is consistent with the “wealthy hand-to-mouth” phenomenon described, e.g., by Aguiar, Bils, and Boar, 2020; however, despite the statistical significance due to the precision of the estimation, the economic significance of the effect at the top of the distribution of liquid wealth suggests that these individuals hardly deviate at all.¹⁴

To sum up, the current empirical analysis leads us to conclude to excess sensitivity of spending to regular anticipated income, in particular for individuals who detain little liquid wealth.

¹⁴The dispersion of these effects is higher than in Olafsson and Pagel, 2018; explanations could stem from less unequal distributions in Iceland, for instance, due themselves to various reasons including institutional setting, generosity of social insurance and redistribution schemes, etc.

3.2.2 Back-to-school allowances

We turn to another form of regular anticipated income, perceived on an annual basis, though. We replicate the analysis made by Bounie et al., 2020. Back-to-school allowances (*Allocations de rentrée scolaire* or ARS in French) designate a one-shot, annual payment intended for poorest families with children above 6 years old, in order to help them dealing with schooling expenditures. This welfare transfer is made before the start of the school year, in August, more precisely on week 34. In what follows, we focus on year 2020. Note that the take-up rate is expected to be high since the payment is automatic once families have registered to family insurance, which is almost mandatory at childbirth. The transfer is means-tested: eligible households with one child (two children) must earn no more than €25,000 (€31,000) per year. Its amount depends on the household's structure, it is for instance equal to €369.97 for families with a single child aged between 6 and 10. Importantly, due to the Covid-19 pandemic, these benefits were expanded by a supplementary €100 bonus. Though the main part of the benefit (namely, the €369.97) ought to be anticipated by households due to the annual nature of the payment, it should not be the case of the extra €100; disentangling MPCs out of the anticipated part from the one out of the unanticipated part of the benefit sounds nevertheless a difficult task. In the following analysis, our working sample is restricted to active customers between weeks 28 and 40.

Causal inference follows from a difference-in-difference strategy where 2,856 eligible households with two children, a child aged 6 to 10 and a child aged 3 to 5, are the treatment group, and compared with their 2,237 siblings having either one or two children aged 3 to 5, and earning no more than €45,000. This identification strategy follows closely the one used by Bounie et al., 2020. It relies on a common trend assumption, according to which the evolution of spending would have been the same in both comparison and treatment groups in the absence of the ARS payment. The plausibility of the latter can be assessed by searching for any differences in pre-trends; as in Bounie et al., 2020, an eyeball analysis (Figure 8) suggests that the corresponding patterns moved in parallel.

Our results point out to an estimated 4-week MPC of about 0.29, which is very comparable to the 5-week MPC of 0.38 obtained by Bounie et al., 2020. An event study analysis (Figure 9) confirms that most of the effect is immediate and stems from the very first week the payment is received. To go further this average treatment effect, it is possible to recover the heterogeneity of MPC along various dimensions (Figure 10); it turns out that neither income, illiquid wealth, nor liabilities like mortgage loans affect these MPCs significantly. In contrast, liquid wealth matters a lot: we estimate that MPCs follow a U-shape with respect to the distribution of liquid wealth-to-income. In the first quintile, MPCs are as high as 0.58 (comprised between 0.45 and 0.71), in the second quintile they amount to 0.2 only, while one cannot reject their nullity in the third and fourth quintiles of that distribution; finally, the estimated MPC in the upper part of this distribution reaches 0.35. Computing

the elasticity of spending to income leads to a value of 0.8 in the bottom of the distribution, given the average spending-to-income ratio (about 1, remember Figure 1) within that quintile.

All-in-all, previous results are consistent with those obtained by Bounie et al., 2020, but lie above those found by Gelman, 2021b (about 0.1 for the average MPC) estimated from tax rebates, which stems likely from the fact that family benefits concern poorer individuals. As far as heterogeneity is concerned, they contrast with Kueng, 2018b who finds increasing MPC over the distribution of liquid wealth, and who interprets his results as a rejection of liquidity constraints being the main source of deviation from LC-PIH, concluding rather to empirical evidence of near rationality. However, they are in line with the presence of wealthy hand-to-mouth described in Aguiar, Bils, and Boar, 2020 and already present in Kaplan and Violante, 2014.

3.3 Discussion

Based on the former analysis devoted to various transitory shocks, we conclude first that irregular and imperfectly anticipated shocks like one-shot transfers are accompanied by large responses, the estimated MPC (as well as the corresponding pass-through) being about 0.9, not far from 1 as theory would predict. In contrast with the most basic version of the LC-PIH model, though, we find excess sensitivity of spending to regular anticipated income, regardless of its frequency (monthly or annual). Typical estimates of MPC out of transitory income are 0.3 on average. Yet highest MPCs (say, 0.58) are found at the bottom (and, to a smaller extent, at the top) of the distribution of liquid wealth, which is consistent with a role played by liquidity constraints, hence with incomplete market models à la Kaplan and Violante, 2014 or/and with wealthy hand-to-mouth individuals (Aguiar, Bils, and Boar, 2020).

4 Permanent income shocks

We now examine how spending responds to permanent income shocks. Among them, we distinguish positive income shocks that could fix temporarily liquidity issues, and hence be a channel through which spending increases, from negative income shocks which neutralize that channel. In that sense, excess sensitivity of consumption consecutive to such shocks is often interpreted as empirical evidence of behavior-based explanations to deviations from PIH. We focus first on job losses as examples of negative income shocks, and then on wage increases as cases for positive income shocks, on top of investigating spending's reaction to a revaluation of disability benefits. Most of these shocks are unanticipated, at

least several months before their realization. The standard LC-PIH model predicts large responses of spending, up to full pass-through.

4.1 Negative shocks: job losses

Institutional setting In France, the notice period before job loss varies with seniority, but it tends to be 2 months in general. It may be accompanied with severance payments, depending on collective bargaining agreements and seniority. The generosity of unemployment insurance (UI) is described in details by Boutchenik and Lardeux, 2020. UI benefits are characterized by two parameters: duration and replacement rate. Their duration is some function of past employment history: each working day entitles workers with one day of benefit, provided that total employment history exceeds 4 months within the last 28 months (before 53) or 36 months (above 53). The maximal duration of UI benefits is 2 or 3 years, depending on age (before or above 53). The replacement rate ranges from 57% to 75% of a reference daily wage w computed as the average daily wage on a 4 to 12 months period before the end of the last job in employment history. Daily benefits b are computed as follows: $b = \min\{0.75w, \max(29.56\alpha, 0.404w + 12.12\alpha, 0.57w)\}$ where $\alpha \in [0, 1]$ designates the ratio of hours worked divided by legal employment duration ($\alpha = 1$ for full-time jobs). Claimants to UI benefits should not have deliberately quit their jobs.

Definition of job loss (unemployment shock) Two kinds of income shocks may be examined: the one that arises at onset of unemployment, and the other arising at exit (exhaustion of UI benefits). Due to statistical reasons (namely a small number of observations at exit), we choose to focus on onset only, contrary to what Ganong and Noel, 2019 do. In the subsequent analysis, the unit of observation is individual-by-month. The observation window for the event analysis is 10 months before to 4 months after. In the data we observe neither individual employment history nor inflows from unemployment insurance.¹⁵ Yet our data contain monthly dummies equal to one when households have received some UI benefits. We therefore rely on the latter to define an unemployment spell as follows: such an episode should include at least two consecutive months with some UI benefits preceded by at least ten months without any UI benefit. We assume that job loss occurs the month before first UI allowance; we have performed robustness checks with respect to this assumption, after which our results remain mostly unchanged. We select out of our working sample individuals whose monthly labor earnings do not exceed €1,000 during at least 5 months over the 10 months before first UI allowance, as well as those whose average monthly earnings from 10 to 4 months before job loss are less than 95% of monthly earnings during the unemployment spell from one month before to one month after job loss. This selection leaves us with 2,923 individuals over the period.

¹⁵Remember that each incoming transfer has a label that is unknown to us.

Econometric specification We follow the approach used by Andersen et al., 2021, namely an event study around job loss. We consider several outcomes on top of income (spending, liquid savings, borrowing, private transfers), which enable us to provide a decomposition of the income shock along these margins. Each outcome y_{it} is normalized with respect to pre-event income, and is detrended from individual (δ_i) and time (γ_t) fixed-effects:

$$y_{it} = \sum_{h=-10}^{+4} \beta_h \mathbb{1}[e_{it} = h] + \delta_i + \gamma_t + \epsilon_{it} \quad (2)$$

where e_{it} is event time defined as distance in months to job loss. Standard errors are clustered at the individual level and computed from a bootstrapping procedure that is carried out with resampling of individuals, rather than individual observations, to account for heteroskedasticity and autocorrelation within observations for the same individual.

Identification Identifying assumptions are (i) parallel trends and (ii) no anticipation of the treatment at $t - a$, where a is the anticipation horizon.¹⁶ Contrary to difference-in-differences designs comparing a treatment group and a control group, in our setting every individual is treated, i.e. loses her job, at some point. Using the terminology found, e.g., in Callaway and Sant’Anna, 2021, the first assumption requires outcomes to evolve similarly for both “treated” (unemployed from time t onward) and “not-yet-treated” (not yet unemployed at time t); violations of this hypothesis include distinct pre-trends before treatment, i.e. before time t . The second assumption means that individuals do not anticipate the treatment by time $t + a$, hence restricting their foresight. In practice, we choose to fix $a = 10$ months before job loss, so that we assume away anticipation effects 10 months before this event. In that sense, we consider that job loss is both persistent and unanticipated. As a result, we normalize $\beta_{-10} = 0$; since another normalization is required in order to achieve full identification of remaining parameters of the model, and instead of the usual normalization $\beta_{-1} = 0$ in event study designs, we impose rather $\beta_{-4} = 0$ due to possible measurement error in the definition of job loss. Indeed, despite possible measurement error as regards the exact moment of job loss, individuals are most likely employed 4 months prior to the beginning of our unemployment spells.

Results To empirically assess the plausibility of the first identifying assumption, we estimate an alternative specification to equation (2) in which individual FE δ_i have been set to zero, but in which we allow for a linear trend common to both groups, on top of imposing $\beta_h = 0, \forall h = -10, \dots, -4$, hence neutralizing any group-specific pre-trend. For the sake of comparability, we adopt the same normalization as before, namely $\beta_{-10} = 0$.

¹⁶For more details as regards identification of two-way fixed-effect models used in event study designs, see Borusyak and Jaravel, 2017.

Reinsuringly, both specifications provide us with very similar estimates, which confirms the plausibility of our first identifying assumption, and gives some credit to the validity of the current identification strategy. Removing then the common pre-trend permits to observe that both spending and liquid savings followed some increasing trend before event, but that job loss put an end to this trend, as is the case in Andersen et al., 2021. Last, our preferred estimating equation (2) confirms that no significant change in outcomes occurs between 10 and 4 months before event, contrary to what happens afterwards: we interpret this empirical evidence as supporting our common trend assumption.

Figure 11 displays the point estimates obtained from the event study. From an event study viewpoint, we compute the cumulative impact (Table 6) over the observation window, hence summing our $\hat{\beta}$ estimates from month -10 to 4. Over that period, the income loss amounts to 1.58 months of pre-event household disposable income. Meanwhile, spending drops by 0.25 month of pre-event household disposable income, i.e. 16% of the income loss. This finding lies below the empirical evidence in Ganong and Noel, 2019 and Andersen et al., 2021: the corresponding numbers are 28% in the U.S. and 30% in Denmark. Liquid assets are, by far, the most important margin by which households respond to job loss (about 2/3): households stop saving, while they would have kept on accumulating liquid wealth in the absence of job loss. Households resort slightly more to credit, a specific form of self-insurance; though statistically significant, the effect is nevertheless not economically significant since that margin is almost negligible when compared to previous ones. Finally, other inflows including private transfers account for about 14% of the income loss. Hence, as in Andersen et al., 2021, we are almost able to achieve a perfect decomposition of the unemployment shock along previous different dimensions on our data: in practice, previous margins account for 96% of the initial income shock. But contrary to both Danish and U.S. cases, spending seems to adjust less than liquid wealth.

In the long run, we estimate that the persistent income loss consecutive to a job loss is 26% of pre-event disposable income. The reduction of spending is durable, about 6% of pre-event disposable income, hence a corresponding long-term MPC of 0.23.

We now seek to estimate how that response varies with income and liquid assets. To that aim, we estimate a first-difference version of previous estimating equation:

$$\Delta C_{it} = \beta(X_i)\Delta Y_{it} + \Gamma_t + \Delta \epsilon_{it} \quad (3)$$

where the Δ operator corresponds to the difference between the average last three months (post-event) and the average first three months (pre-event). C and Y designate spending and income respectively, and both are still normalized by pre-event income. In the case of an homogeneous coefficient $\beta(X_i) = \beta_0$, we obtain a 0.196 point estimate (see Table 7), an effect that compares well with previous numbers. Interestingly, we now allow this coefficient to depend on the location in the distribution of income or liquid assets. The

coefficient decreases from 0.33 in the lowest quartile to 0 in the highest quartile of the liquid wealth distribution. The dispersion found is more substantial than the one related to income, where the coefficient varies from 0.25 to 0.08, and exhibits non-monotonic behavior. When we allow the coefficient to depend both on income and liquid assets, it is confirmed that liquidity matters more than income in this regard: conditional on liquid wealth, the estimated coefficient does not vary much with respect to income.

4.2 Positive shocks

4.2.1 Wage increases

Definition of a wage increase In the subsequent analysis, the unit of observation is individual-by-month. The observation window is 9 months before wage increase (the event hereafter) to 6 months after. We pay attention here to “treated” individuals who experience some durable increase in their earnings, which we call a wage increase, and which we measure as follows: (i) each quartile of the post-event 6-month distribution of monthly earnings must be greater or equal than each quartile of the post-event 9-month distribution of monthly earnings; (ii) the event occurs during the month with the highest relative rise in earnings. We compare these individuals with their siblings having each quartile of their last 6 month distribution of monthly earnings comprised between 0.95 and 1.05 times each quartile of their first 9-month distribution of monthly earnings. We select out of our working sample individuals whose monthly labor earnings do not exceed €1,000 during at least 5 months over the nine months before event. This selection leaves us with 3,338 treated individuals and with about 6,375 non-treated over the period.

We resort to an event study design issued from a difference-in-differences method, i.e. on the existence of the treatment group and the comparison group defined above. The identification strategy is based on a common trend assumption between these two groups. We adopt the normalization $\beta_{-6} = 0$ to avoid measurement error on the beginning of wage increase as much as possible. Figure 12 confirms that no significant pre-trend in outcomes can be observed between 9 and 6 months before event, which constitutes empirical support for our identifying assumption. Our estimating equation is:

$$y_{it} = \sum_{h=-9}^{+6} \beta_h T_i \mathbb{1}[e_{it} = h] + T_i + \gamma_t + \epsilon_{it} \quad (4)$$

where T_i is the treatment dummy.

Results The magnitude of income shocks considered here is substantial: about 8.75% of pre-event income. We are therefore confident in the shock being salient to their benefi-

ciaries. A few months after, concerned individuals have increased their spending by about 3% of pre-event income, hence an estimated MPC of 0.31. The cumulative impact over the observation window¹⁷ yields to an income gain that represents nearly 54% of the pre-event income. We estimate on Table 8 that the spending response is worth 16% of that pre-event income, which points out to a MPC out of that event of 0.3. Liquid savings are still the most important margin of adjustment: the estimated marginal propensity to save liquid assets out of a wage increase is 0.51. Credit and private transfers only explain a marginal part of households' response to the shock. The decomposition along different margins is slightly more imperfect (86%) than the one obtained from job losses. We suspect that outflows, in particular towards illiquid assets located in different financial institutions (especially life insurance and retirement savings), are missing from this picture.¹⁸

Investigating further the heterogeneity in MPCs, we consider a first-difference version of the difference-in-differences equation based on treated individuals only:

$$\Delta C_{it} = \beta(X_i)\Delta Y_{it} + \Gamma_t + \Delta\epsilon_{it} \quad (5)$$

The estimated MPC is higher, amounting now to 0.41 (Table 9), which can be explained by slightly lower levels of spending before our normalized month (6 months before wage increase). Allowing this coefficient to depend on income or/and liquid wealth, we find that it decreases from 0.58 in the lowest quartile to 0.18 in the highest quartile of the liquid wealth distribution. The dispersion found is more substantial than the one related to income, where the coefficient varies from 0.17 to 0.49 in a non-monotonic fashion. Moreover, when we allow the coefficient to depend both on income and liquid assets, it is confirmed that liquidity matters more than income in this regard: conditional on liquid wealth, the estimated coefficient does not vary much with respect to income. On top of that, it seems to be fairly homogeneous all along the distribution of liquidity, since one cannot reject that the MPC is in fact even higher, up to 0.62, except in the highest part where it falls to 0.15.

4.2.2 Revaluation of disability benefits

Last, we exploit a quasi-natural experiment based on the revaluation of disability benefits, the *Allocation aux Adultes Handicapés* (AAH), intervened on November 1st, 2019, which we call the event hereafter. The government decided to increase the monthly amount of this allowance from €860 to €900, i.e. by €40 or +4.7%.¹⁹ Though the post-revaluation

¹⁷including 4 months post-event, hence 5 months of wage increase.

¹⁸Interestingly, the negative shock consecutive to unemployment was immune to such a “missing variable” issue.

¹⁹These integer amounts are included in our sample here, contrary to what prevails in the rest of our analysis.

era has been much affected by lockdowns and by various measures to limit the incidence of the Covid (hence by harsh restrictions on consumption from March to May 2020, from November 2020 to January 2021 as well as in April 2021), Figure 13 suggests that the pre-event median level of spending for AAH beneficiaries amounted roughly to €835. By contrast, the post-event median level of spending oscillates around €875=€835+€40, with strong fluctuations due to the restrictions mentioned, which indicates full pass-through of that revaluation to spending. Since the average spending-to-income ratio is almost equal to one for this sub-population that belongs to the bottom of the income distribution (remember Figure 1),²⁰ it also means that the corresponding long-term elasticity is close to one.

4.3 Discussion

From this analysis devoted to permanent shocks, we conclude first that imperfectly anticipated shocks like wage increases, job losses or revaluation of social benefits are accompanied by large spending responses, the corresponding MPCs reaching 0.6 in the case of positive shocks (and even 1 for poorest individuals depending on disability benefits), as theory would predict. Yet some asymmetry is found since MPCs out of negative shocks (e.g. job losses) are much lower (0.2) than MPCs out of positive shocks (e.g. wage increases), comprised between 0.3 and 0.4. This asymmetry is not consistent with liquidity constraints fully explaining observed deviations from PIH. Rather, it points out to behavior-based explanations. This empirical finding contrasts notably with the one unraveled, e.g., by Baker, 2018. Nevertheless, a growing literature including at least Baugh et al., 2021, has started to take a close look at asymmetric responses to income shocks: for instance, tax payments and tax refunds seem to induce different consumption smoothing behaviors, with consumers smoothing in the case of a loss but exhibiting excess sensitivity in the presence of a gain. More generally, this is reminding of Kahneman and Tversky, 1979 emphasizing the role of loss aversion as a driving mechanism for the observed asymmetry at stake.

Second, we find substantial heterogeneity behind those average effects: much higher MPCs are found at the bottom of the liquid wealth distribution (about 0.6 consecutive to positive shocks, 0.33 following negative shocks) whereas lower MPCs prevail at the top of this distribution (about 0.15 consecutive to positive shocks and 0 after negative shocks). The latter empirical evidence is consistent with a role played by liquidity constraints, on top of behavioral explanations, as is the case in Gelman, 2021a who concludes to an equal role of both explanations.

²⁰To be precise, we can estimate this ratio to $835/860 \approx 0.97$ ($875/900 \approx 0.97$) before (after) November 1st, 2019.

5 Concluding remarks

Our main contribution has been to explore how spending responds to a collection of distinct income shocks based on a single (or similar) high-frequency bank account database(s), on the same period of time (2019-2021) and facing identical institutional rules, those prevailing in France. We stress that our approach permits to enhance the comparability of various estimated MPCs that can be found in the literature, based on different income shocks. As a result, our empirical approach provides with convincing tests of theoretical predictions issued from basic life-cycle or standard incomplete market models.

On the whole, our main finding is that MPCs do not differ much when looking either at transitory or at permanent income shocks. In particular, (i) income shocks are fully passed through to spending for welfare benefits beneficiaries, regardless of their nature (temporary -Covid-19 bonuses- or persistent -the revaluation of disability benefits), and (ii) MPCs out of positive transitory shocks (like the payment of back-to-school allowances) and MPCs out of positive permanent shocks (like wage increases) look remarkably similar -not only on average (0.3) but also when we focus on comparable sub-populations, e.g. at the bottom of the liquid wealth distribution (0.6 in both cases). If the message conveyed here does not differ much from what a meta-analysis of various MPCs found in the literature would permit to learn, we believe nevertheless that the main advantage of our approach resides in us neutralizing confounding factors (sample composition effects, institutional setting, time period, etc.) that would threaten the comparability of these estimates. This is made possible by the availability of high-frequency bank account data and the variety of income shocks we exploit in order to infer corresponding MPCs. Overall, our findings challenge standard predictions derived from the PIH according to which responses should be higher consecutive to permanent shocks.

On top of that, our results are consistent with some empirical evidence on MPCs: (i) they are non-zero, consistently with a rejection of the most parsimonious version of the LC-PIH model; (ii) they exhibit much heterogeneity with respect to liquid wealth, which suggests a role for liquidity constraints; (iii) lower (higher) MPCs are found consecutive to (positive) negative shocks, which indicates, by contrast, some role for behavior-based explanations, on top of liquidity-based explanations, in observed deviations from PIH.

We believe that these up-to-date estimates are interesting for policy-makers since they are derived from almost real-time data, and thus immediately useful to them in order to design optimal fiscal and social policies. They are also precious to researchers, in relation with the theoretical implications mentioned above. A natural extension would consist in estimating a structural incomplete-market model that allows for individual failures such as, e.g., present bias, hyperbolic discounting, etc. A test of such model should help quantifying the relative importance of the two likely channels at stake accounting for deviations from PIH, liquidity constraints and behavior-based explanations.

References

- Agarwal, Sumit, Chunlin Liu, and Nicholas S Souleles (2007). “The reaction of consumer spending and debt to tax rebates—evidence from consumer credit data”. In: *Journal of political Economy* 115.6, pp. 986–1019.
- Agarwal, Sumit and Wenlan Qian (2014). “Consumption and debt response to unanticipated income shocks: Evidence from a natural experiment in Singapore”. In: *The American Economic Review* 104.12, pp. 4205–30.
- Aguiar, Mark A, Mark Bilz, and Corina Boar (2020). *Who are the Hand-to-Mouth?* Tech. rep. National Bureau of Economic Research.
- Andersen, Asger Lau, Niels Johannesen, and Adam Sheridan (2021). *Dynamic Spending Responses to Wealth Shocks: Evidence from Quasi-lotteries on the Stock Market*. CEPR Discussion Paper No. DP16338.
- Andersen, Asger Lau et al. (2020). *Consumer responses to the COVID-19 crisis: Evidence from bank account transaction data*. Available at SSRN 3609814.
- Andersen, Asger Lau et al. (2021). *How Do Households Respond to Job Loss? Lessons from Multiple High-Frequency Data Sets*. CEPR Discussion Paper No. DP16131.
- Baker, Scott R (2018). “Debt and the response to household income shocks: Validation and application of linked financial account data”. In: *Journal of Political Economy* 126.4, pp. 1504–1557.
- Baugh, Brian et al. (2021). “Asymmetric consumption smoothing”. In: *American Economic Review* 111.1, pp. 192–230.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008). “Consumption inequality and partial insurance”. In: *The American Economic Review* 98.5, pp. 1887–1921.
- Borusyak, Kirill and Xavier Jaravel (2017). *Revisiting event study designs*. Working paper.
- Bounie, David et al. (2020). *Dynamiques de consommation dans la crise: les enseignements en temps réel des données bancaires*. Working paper.
- Boutchenik, Béatrice and Raphaël Lardeux (2020). *The Take-Up of Unemployment Benefit Extensions*. Document de travail Insee G2020/02.
- Browning, Martin and Thomas F Crossley (2001). “Unemployment insurance benefit levels and consumption changes”. In: *Journal of Public Economics* 80.1, pp. 1–23.
- Bunn, Philip et al. (2018). “The consumption response to positive and negative income shocks”. In: *Journal of Monetary Economics* 96, pp. 1–15. ISSN: 0304-3932. DOI: <https://doi.org/10.1016/j.jmoneco.2017.11.007>. URL: <https://www.sciencedirect.com/science/article/pii/S0304393217301411>.
- Callaway, Brantly and Pedro HC Sant’Anna (2021). “Difference-in-differences with multiple time periods”. In: *Journal of Econometrics* 225.2, pp. 200–230.
- Carvalho, Vasco M et al. (2020). *Tracking the COVID-19 crisis with high-resolution transaction data*. CEPR Discussion Paper No. DP14642.
- Deville, Jean-Claude and Carl-Erik Särndal (1992). “Calibration estimators in survey sampling”. In: *Journal of the American statistical Association* 87.418, pp. 376–382.

- Disney, Richard, John Gathergood, and Andrew Henley (2010). “House price shocks, negative equity, and household consumption in the United Kingdom”. In: *Journal of the European Economic Association* 8.6, pp. 1179–1207.
- Fagereng, Andreas, Martin B Holm, and Gisle J Natvik (2021). “MPC heterogeneity and household balance sheets”. In: *American Economic Journal: Macroeconomics* 13.4, pp. 1–54.
- Fagereng, Andreas et al. (2019). *Saving behavior across the wealth distribution: The importance of capital gains*. NBER Working paper 26588.
- Ganong, Peter and Pascal Noel (2019). “Consumer spending during unemployment: Positive and normative implications”. In: *The American Economic Review* 109.7, pp. 2383–2424.
- Gelman, Michael (2021a). *The Self-Constrained Hand-to-Mouth*. *The Review of Economics and Statistics*, forthcoming.
- (2021b). “What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics”. In: *Journal of Monetary Economics* 117, pp. 521–542.
- Gelman, Michael et al. (2014). “Harnessing naturally occurring data to measure the response of spending to income”. In: *Science* 345.6193, pp. 212–215.
- Gelman, Michael et al. (2016). *The response of consumer spending to changes in gasoline prices*. Tech. rep. National Bureau of Economic Research.
- Hastings, Justine and Ebonya Washington (2010). “The first of the month effect: consumer behavior and store responses”. In: *American economic Journal: economic policy* 2.2, pp. 142–62.
- Insee (2021). “En 2020, la chute de la consommation a alimenté l'épargne, faisant progresser notamment les hauts patrimoines financiers.” In: *Note de conjoncture*.
- Jappelli, Tullio and Luigi Pistaferri (2014). “Fiscal policy and MPC heterogeneity”. In: *American Economic Journal: Macroeconomics* 6.4, pp. 107–36.
- Johnson, David S, Jonathan A Parker, and Nicholas S Souleles (2006). “Household expenditure and the income tax rebates of 2001”. In: *The American Economic Review* 96.5, pp. 1589–1610.
- Kahneman, Daniel and Amos Tversky (1979). “Prospect Theory: An Analysis of Decision under Risk”. In: *Econometrica* 47.2, pp. 263–292.
- Kaplan, Greg and Giovanni L Violante (2014). “A model of the consumption response to fiscal stimulus payments”. In: *Econometrica* 82.4, pp. 1199–1239.
- Kuchler, Theresa and Michaela Pagel (2021). “Sticking to your plan: The role of present bias for credit card paydown”. In: *Journal of Financial Economics* 139.2, pp. 359–388.
- Kueng, Lorenz (2018a). “Excess sensitivity of high-income consumers”. In: *The Quarterly Journal of Economics* 133.4, pp. 1693–1751.
- (June 2018b). “Excess Sensitivity of High-Income Consumers”. In: *The Quarterly Journal of Economics* 133.4, pp. 1693–1751. ISSN: 0033-5533. DOI: [10.1093/qje/qjy014](https://doi.org/10.1093/qje/qjy014). eprint: <https://academic.oup.com/qje/article-pdf/133/4/1693/>

25985656/qjy014_supplemental_file.pdf. URL: <https://doi.org/10.1093/qje/qjy014>.

- Mastrobuoni, Giovanni and Matthew Weinberg (2009). “Heterogeneity in intra-monthly consumption patterns, self-control, and savings at retirement”. In: *American Economic Journal: Economic Policy* 1.2, pp. 163–89.
- Misra, Kanishka and Paolo Surico (2014). “Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs”. In: *American Economic Journal: Macroeconomics* 6.4, pp. 84–106.
- Olafsson, Arna and Michaela Pagel (2018). “The liquid hand-to-mouth: Evidence from personal finance management software”. In: *The Review of Financial Studies* 31.11, pp. 4398–4446.
- Parker, Jonathan A et al. (2013). “Consumer spending and the economic stimulus payments of 2008”. In: *The American Economic Review* 103.6, pp. 2530–53.
- Shapiro, Jesse M (2005). “Is there a daily discount rate? Evidence from the food stamp nutrition cycle”. In: *Journal of public Economics* 89.2-3, pp. 303–325.
- Souleles, Nicholas S (1999). “The response of household consumption to income tax refunds”. In: *The American Economic Review* 89.4, pp. 947–958.
- Stephens Jr, Melvin (2003). ““ 3rd of tha Month”: Do social security recipients smooth consumption between checks?” In: *American Economic Review* 93.1, pp. 406–422.
- (2006). “Paycheque receipt and the timing of consumption”. In: *The Economic Journal* 116.513, pp. 680–701.

A Figures

A.1 Data: External validity




Figure 1 – Household expenditures and spending-to-income ratio, by income (transaction data vs. survey data from *Budget des Familles*, Insee)




Figure 2 – Income (transaction data vs. aggregate data from national accounts, Insee)




Figure 3 – Liquid Assets (transaction data vs. aggregate data from Banque de France)

A.2 The expansion of social benefits due to the Covid-19 pandemic




Figure 4 – Spending response to the expansion of social benefits

A.3 The payday effect




Figure 5 – The payday effect

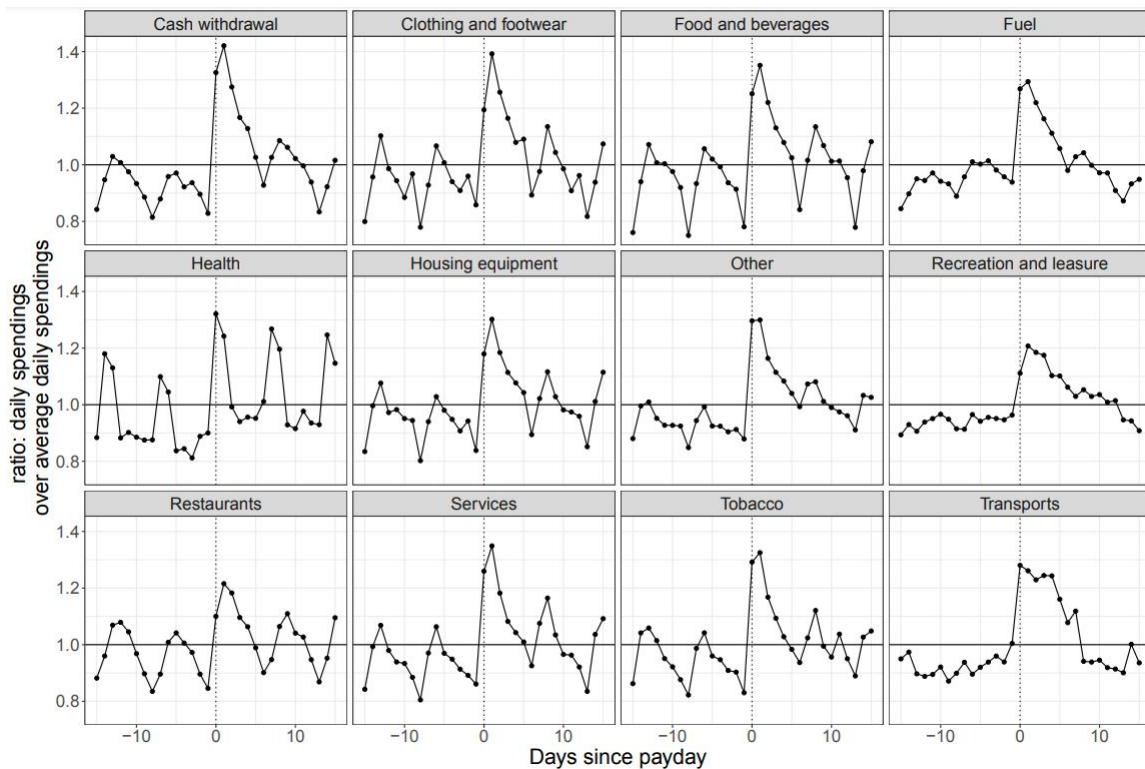

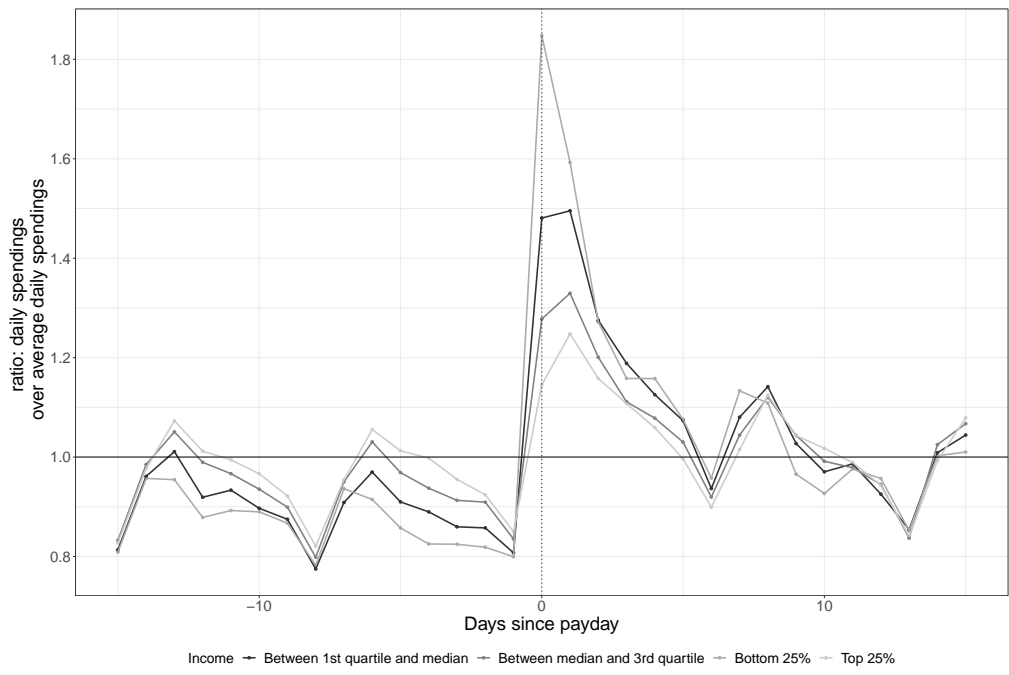


Figure 6 – The payday effect (by category of spending)



(a) By assets



(b) By income

Figure 7 – Heterogeneity of the payday effect

A.4 Back-to-school allowances

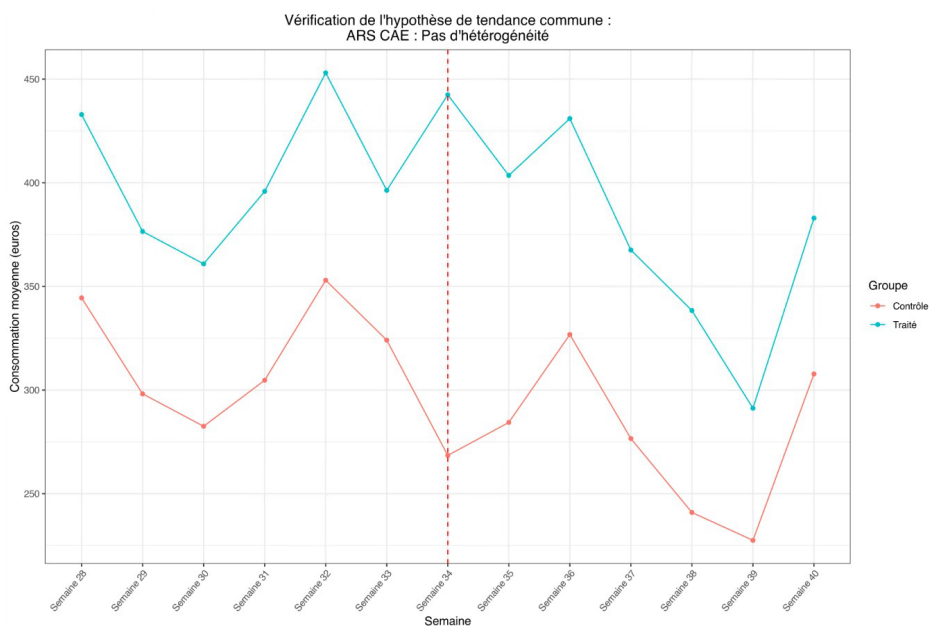


Figure 8 – Evolution of spending in comparison and treatment groups

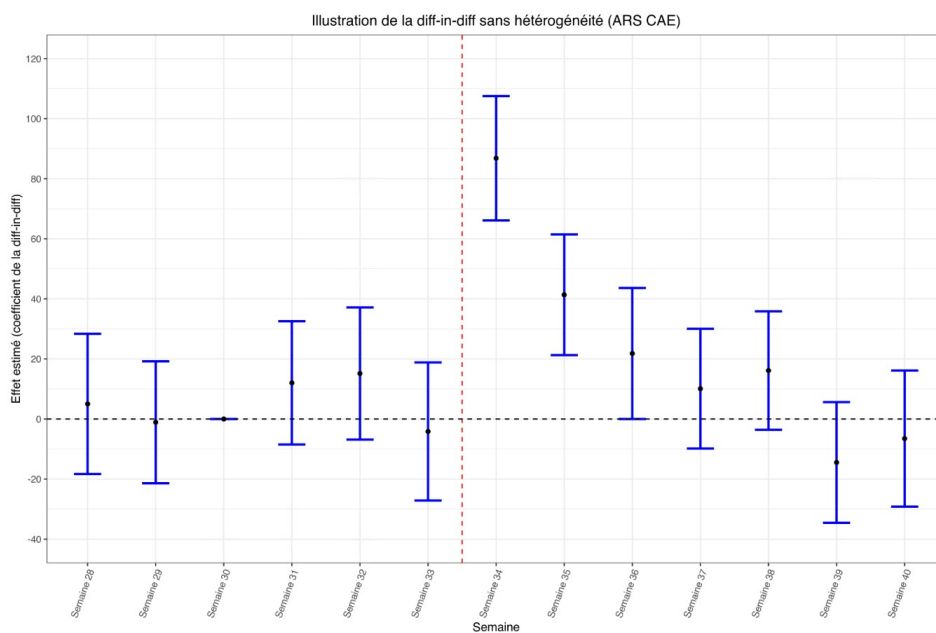


Figure 9 – Spending response to the ARS payment

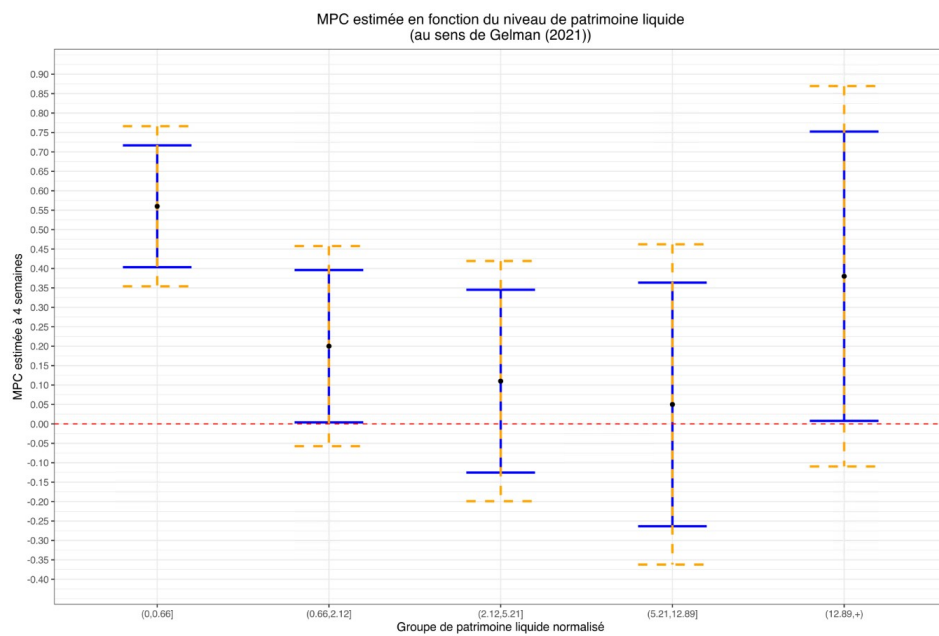


Figure 10 – Heterogeneity of MPCs out of the ARS payment

A.5 Job losses

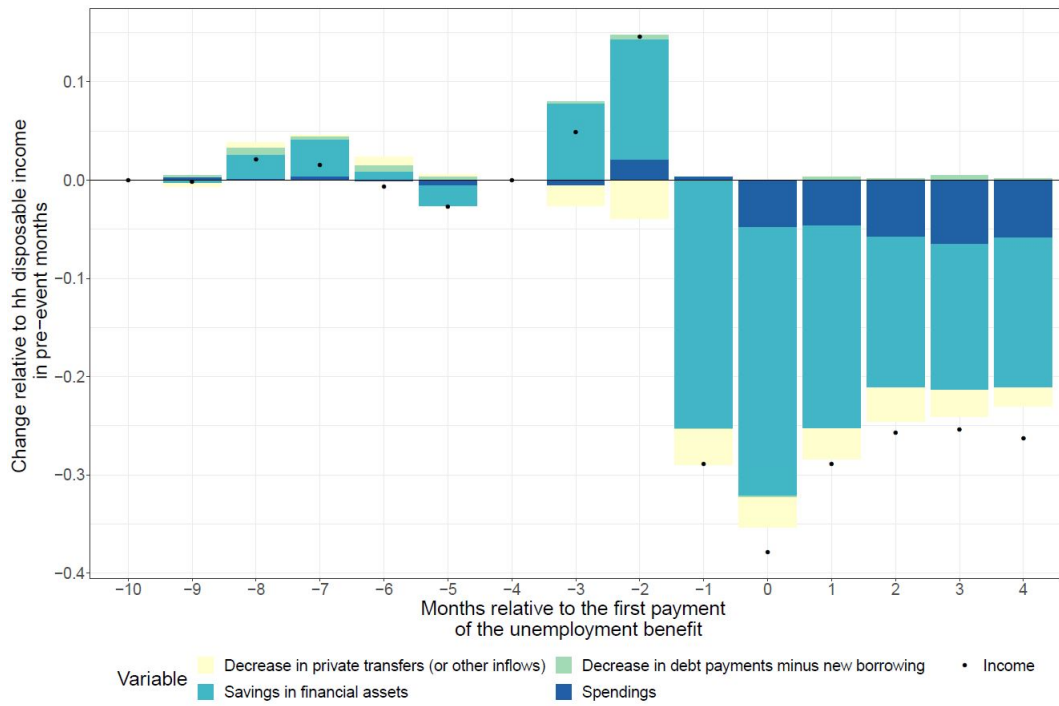


Figure 11 – Income, spending and other self-insurance responses to job losses

A.6 Wage increases

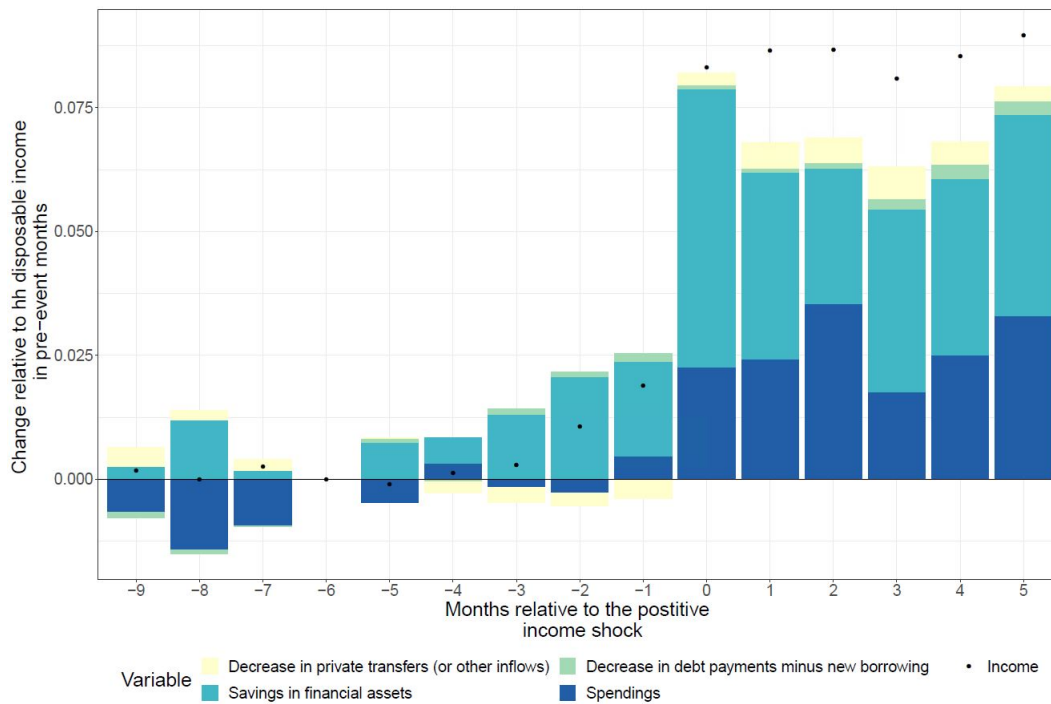


Figure 12 – Income, spending and other self-insurance responses to wage increases

A.7 Revaluation of disability benefits




Figure 13 – Income and spending responses to a revaluation of disability benefits

B Tables

Table 1 – Summary statistics

	(1)	(2)	(3)
	Crédit Mutuel		
	Unweighted sample	Weighted sample	National surveys
# of observations	169,163	169,163	
# of months	35	35	
	<i>Sample means</i>		
Spending	2,371	2,461	2,284
<i>Credit cards</i>	1,650	1,698	
<i>Bills</i>	713	756	
<i>Checks</i>	8	7	
<i>Utilities (bills and cards)</i>	140	148	113
<i>Groceries (cards)</i>	232	248	368
<i>Restaurants (cards)</i>	98	97	136
<i>Fuel (cards)</i>	78	79	92
Income	3,497	3,492	2,924
Financial Assets	50,657	55,615	50,882
<i>Liquid financial Assets</i>	32,858	35,241	24,270
Deposit account	9,514	10,525	4,046
Savings account	23,345	24,716	20,224
<i>Illiquid financial Assets</i>	17,799	20,374	26,612
Life insurance	13,597	15,748	18,947
Securities account	4,202	4,626	7,664
Monthly savings	95	83	
Loan net repayments	-390	-389	
Non-mortgage debt	-3,024	-3,086	-5,377
Mortgage debt	-34,793	-35,203	-38,605
Private transfers (or other inflows)	1,457	1,542	
Ratio liquid assets/deposit account	3.45	3.35	5.99
Age	49	52	52
Female	0.52	0.51	0.51
Craftsmen, merchants and business owners	0.06	0.06	0.04
Managerial and professional occupations	0.13	0.12	0.10
Technicians and associate professionals	0.14	0.13	0.13
Employees	0.24	0.22	0.14
Workers	0.13	0.12	0.11
Periphery areas	0.42	0.41	0.18
Rural areas	0.20	0.20	0.21
Urban areas	0.32	0.33	0.61

Pecuniary amounts: in €.

Table 2 – The payday effect: main estimates

	All sample	Quartiles of income			
		1st	2nd	3rd	4th
Our estimates	0.43*** (0.01)	0.79*** (0.04)	0.46*** (0.02)	0.26*** (0.01)	0.14*** (0.01)
Olafsson-Pagel		0.88*** (0.01)	0.59*** (0.01)	0.44*** (0.01)	0.34*** (0.01)
Time FE	✓	✓	✓	✓	✓

Table 3 – The payday effect: robustness

Variable	Model 1	Model 2	Model 3
Pay Day	0.46*** (0.01)	0.37*** (0.01)	0.37*** (0.01)
Time FE		✓	✓
Individual FE			✓

Table 4 – The payday effect: heterogeneity

	(1)	(2)	(3)	(4)
	0.46 (0.01)			1.14 (0.04)
Financial assets				
<i>Bottom 25%</i>		1.03 (0.02)		
<i>Between 1st quartile and median</i>		0.33 (0.01)		
<i>Between median and 3rd quartile</i>		0.10 (0.01)		
<i>Top 25%</i>		0.04 (0.01)		
Illiquid financial assets				
<i>No</i>				Ref.
<i>Yes</i>				-0.16 (0.01)
Liquid financial assets				
<i>Bottom25%</i>				Ref.
<i>Between 1st quartile and median</i>				-0.58 (0.02)
<i>Between median and 3rd quartile</i>				-0.78 (0.02)
<i>Top 25%</i>				-0.80 (0.03)
Income				
<i>Bottom 25%</i>			0.64 (0.03)	Ref.
<i>Between 1st quartile and median</i>			0.42 (0.01)	-0.10 (0.02)
<i>Between median and 3rd quartile</i>			0.24 (0.01)	-0.14 (0.02)
<i>Top 25%</i>			0.14 (0.01)	-0.14 (0.02)
Consumption credit undertaken				
<i>No</i>				Ref.
<i>Yes</i>				-0.02 (0.01)
Home credit undertaken				
<i>No</i>				Ref.
<i>Yes</i>				-0.20 (0.01)
Age				
<i>Below 30</i>				Ref.
<i>30-60</i>				0.04 (0.02)
<i>Above 60</i>				-0.06 (0.02)
Household structure				
<i>Single man</i>				Ref.
<i>Couple</i>				-0.09 (0.02)
<i>Family</i>				-0.00 (0.03)
<i>Single parent family</i>				0.38 (0.05)
<i>Single woman</i>				-0.01 (0.02)
Time FE	✓	✓	✓	✓
Individual FE	✓	✓	✓	✓

Table 5 – The payday effect: aggregation at the weekly level

Variable	Model 1	Model 2	Model 3
Pay Week	0.24*** (0.00)	0.22*** (0.00)	0.23*** (0.00)
Time FE		✓	✓
Individual FE			✓

Table 6 – Quantifying income, spending and other self-insurance responses to job losses

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Month -1	Month 0	Month 1	Month 2	Month 3	Month 4	Cumulative -1 to 4	Frac. of income loss
Income	-0.29 (0.02)	-0.38 (0.02)	-0.29 (0.02)	-0.26 (0.02)	-0.25 (0.02)	-0.26 (0.03)	-1.58 (0.14)	100.00 (0.00)
Spending	0.00 (0.01)	-0.05 (0.01)	-0.05 (0.01)	-0.06 (0.01)	-0.06 (0.01)	-0.06 (0.01)	-0.25 (0.06)	16.04 (4.02)
Spending (in % of income loss)	-1.20 (2.56)	12.74 (2.20)	16.31 (3.16)	22.73 (4.36)	25.82 (5.33)	22.52 (5.48)	16.04 (4.02)	16.04 (4.02)
Savings	-0.25 (0.03)	-0.27 (0.03)	-0.21 (0.03)	-0.15 (0.03)	-0.15 (0.04)	-0.15 (0.04)	-1.07 (0.20)	66.98 (10.35)
Inflows	-0.04 (0.01)	-0.03 (0.01)	-0.03 (0.01)	-0.03 (0.01)	-0.03 (0.01)	-0.02 (0.01)	-0.22 (0.07)	13.95 (4.79)
Credit	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)	0.01 (0.03)	-0.95 (1.67)
Total (<i>Spend.</i> + <i>Savs.</i> + <i>Infl.</i> + <i>Cred.</i>)	-0.29 (0.03)	-0.35 (0.03)	-0.28 (0.03)	-0.24 (0.03)	-0.24 (0.03)	-0.23 (0.04)	-1.52 (0.18)	96.04 (8.12)
# of observations	2,923	2,923	2,923	2,923	2,923	2,923	2,923	2,923

Table 7 – Heterogeneity of the MPC for job losses

	(1)	(2)	(3)	(4)
	0.20*** (0.02)			0.31*** (0.04)
Financial assets				
<i>Bottom 25%</i>		0.33*** (0.03)		
<i>Between 1st quartile and median</i>		0.20*** (0.03)		
<i>Between median and 3rd quartile</i>		0.11*** (0.03)		
<i>Top 25%</i>		0.04 (0.04)		
Income				
<i>Bottom 25%</i>			0.25*** (0.04)	Ref.
<i>Between 1st quartile and median</i>			0.30*** (0.03)	0.04 (0.04)
<i>Between median and 3rd quartile</i>			0.20*** (0.03)	-0.02 (0.05)
<i>Top 25%</i>			0.08*** (0.03)	-0.09* (0.05)
Liquid financial assets				
<i>Bottom 25%</i>				Ref.
<i>Between 1st quartile and median</i>				-0.09** (0.04)
<i>Between median and 3rd quartile</i>				-0.14*** (0.04)
<i>Top 25%</i>				-0.17*** (0.05)
Illiquid financial assets				
<i>No</i>				Ref.
<i>Yes</i>				-0.05 (0.04)
Consumption credit undertaken				
<i>No</i>				Ref.
<i>Yes</i>				0.02 (0.03)
Time FE		✓	✓	✓
Socio-demographic controls		✓	✓	✓
Financial pre-treatment vars		✓	✓	✓
R^2		0.19	0.21	0.20
# of obs.		2,923	2,923	2,923

Table 8 – Quantifying income, spending and other self-insurance responses to positive permanent shocks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Month -1	Month 0	Month 1	Month 2	Month 3	Month 4	Cumulative -1 to 5	Frac. of income loss
Income	0.02 (0.00)	0.08 (0.00)	0.09 (0.00)	0.09 (0.00)	0.08 (0.00)	0.09 (0.00)	0.54 (0.02)	100.00 (0.00)
Spending	0.00 (0.01)	0.02 (0.01)	0.02 (0.00)	0.04 (0.01)	0.02 (0.01)	0.03 (0.01)	0.16 (0.03)	29.50 (6.01)
Spending (in % of income loss)	24.77 (32.96)	27.15 (6.30)	28.10 (5.60)	40.90 (6.50)	21.74 (7.19)	29.38 (6.66)	29.50 (6.01)	29.50 (6.01)
Savings	0.02 (0.01)	0.06 (0.01)	0.04 (0.01)	0.03 (0.01)	0.04 (0.01)	0.04 (0.01)	0.27 (0.06)	50.58 (10.09)
Inflows	0.00 (0.00)	0.00 (0.00)	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)	0.00 (0.00)	0.02 (0.01)	3.84 (2.09)
Credit	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.01 (0.01)	2.27 (1.26)
Total (<i>Spend. + Savs.</i> <i>+Infl. + Cred.</i>)	0.02 (0.01)	0.08 (0.01)	0.07 (0.01)	0.07 (0.01)	0.06 (0.01)	0.07 (0.01)	0.47 (0.05)	86.20 (7.71)
# of ind. (treatment)	3,338	3,338	3,338	3,338	3,338	3,338	3,338	3,338
# of ind. (comparison)	6,375	6,375	6,375	6,375	6,375	6,375	6,375	6,375

Table 9 – Heterogeneity of the MPC for positive permanent shocks

	(1)	(2)	(3)	(4)
	0.41*** (0.04)			0.62*** (0.10)
Financial assets				
<i>Bottom 25%</i>		0.58*** (0.73)		
<i>Between 1st quartile and median</i>		0.47*** (0.07)		
<i>Between median and 3rd quartile</i>		0.36*** (0.07)		
<i>Top 25%</i>		0.18 (0.08)		
Income				
<i>Bottom 25%</i>			0.48*** (0.07)	Ref.
<i>Between 1st quartile and median</i>			0.49*** (0.08)	0.08 (0.11)
<i>Between median and 3rd quartile</i>			0.17* (0.09)	-0.20 (0.12)
<i>Top 25%</i>			0.43*** (0.11)	0.1* (0.14)
Liquid financial assets				
<i>Bottom25%</i>				Ref.
<i>Between 1st quartile and median</i>				-0.21** (0.11)
<i>Between median and 3rd quartile</i>				-0.18*** (0.13)
<i>Top 25%</i>				-0.47*** (0.13)
Illiquid financial assets				
<i>No</i>				Ref.
<i>Yes</i>				-0.06 (0.10)
Consumption credit undertaken				
<i>No</i>				Ref.
<i>Yes</i>				0.01 (0.10)
Time FE		✓	✓	✓
Socio-demographic controls		✓	✓	✓
Financial pre-treatment vars		✓	✓	✓
R^2		0.28	0.28	0.28
# of obs.		3,338	3,338	3,338

Supplementary material




Figure 14 – Card and cash spending, by category and income (transaction data vs. survey data from *Budget des familles*, Insee)




Figure 15 – The distribution of regular payment arrival over the month




Figure 16 – Spending response to the expansion of social benefits (1st quartile of assets)




Figure 17 – Spending response to the expansion of social benefits (2nd quartile of assets)




Figure 18 – Spending response to the expansion of social benefits (3rd quartile of assets)




Figure 19 – Spending response to the expansion of social benefits (4th quartile of assets)




Figure 20 – Spending response to the €150 expansion of social benefits




Figure 21 – Spending response to the €250 expansion of social benefits




Figure 22 – Spending response to unemployment shocks (bottom 50% of financial assets)

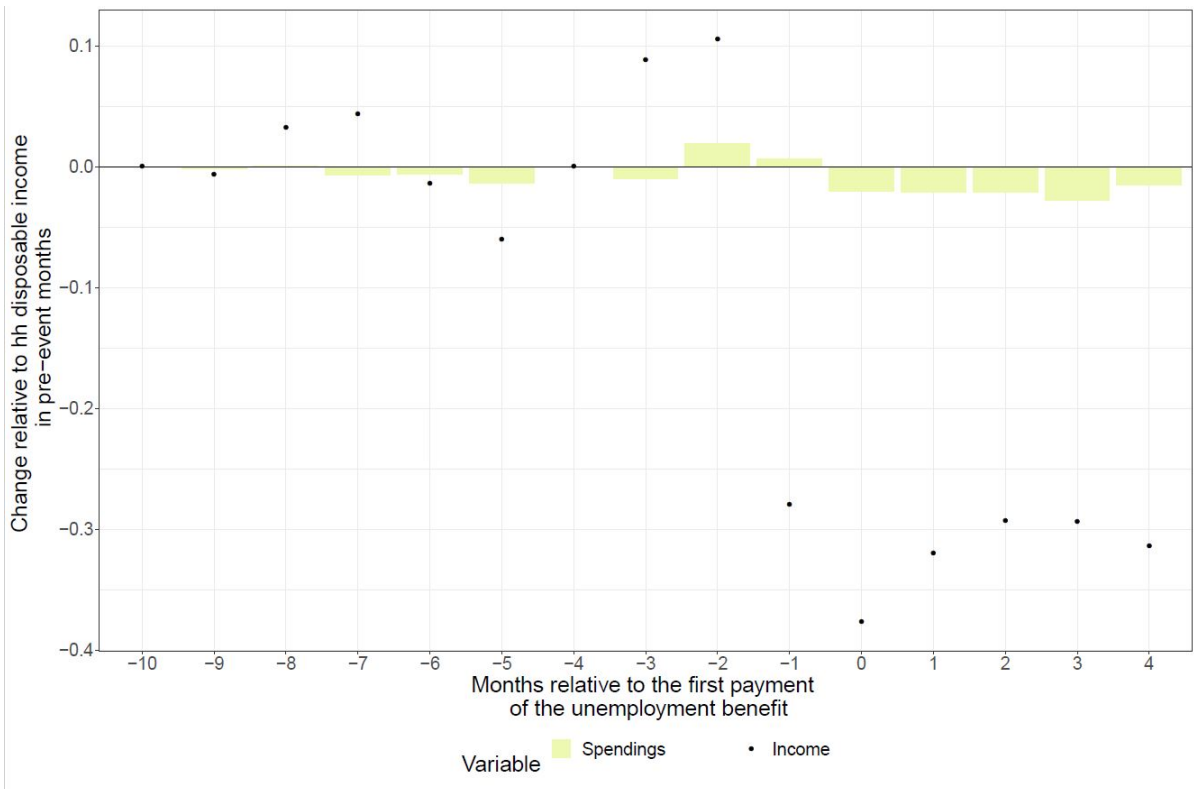


Figure 23 – Spending response to unemployment shocks (bottom 50% of financial assets)