

Pozzi, Lorenzo; Sadaba, Barbara

Working Paper

Macroeconomic disasters and consumption smoothing: International evidence from historical data

Bank of Canada Staff Working Paper, No. 2023-4

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Pozzi, Lorenzo; Sadaba, Barbara (2023) : Macroeconomic disasters and consumption smoothing: International evidence from historical data, Bank of Canada Staff Working Paper, No. 2023-4, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2023-4>

This Version is available at:

<https://hdl.handle.net/10419/297389>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Macroeconomic Disasters and Consumption Smoothing: International Evidence from Historical Data

by Lorenzo Pozzi¹ and Barbara Sadaba²

¹Erasmus University Rotterdam and Tinbergen Institute
pozzi@ese.eur.nl

²International Economic Analysis Department
Bank of Canada
bsadaba@bankofcanada.ca



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

DOI: <https://doi.org/10.34989/swp-2023-4> | ISSN 1701-9397

©2023 Bank of Canada

Acknowledgements

For constructive comments, we thank Tino Berger, Narayan Bulusu, Christian Friedrich, Ana Maria Herrera, Felix Ward and participants of the Economics Department seminar at Erasmus School of Economics, the International Economic Analysis Department seminar at the Bank of Canada and the Midwest Econometrics Group 2022 Conference in Michigan. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

Abstract

This paper uses a large historical dataset (1870–2016) for 16 industrial economies to show that during macroeconomic disasters (e.g., wars, pandemics, depressions) aggregate consumption and income are significantly less decoupled than during normal times. That is, during these times of turmoil, the consumer intertemporal budget constraint holds more strictly, implying a structural reduction in consumption smoothing. While we also observe this for the ongoing COVID-19 pandemic, this is not the case for more conventional post-war recessions. Our results are obtained using a predictive regression approach that follows directly from the forward-looking nature of consumption theory. Using a savers-spenders type of model, we show that our findings can be interpreted as stemming from an increase in rule-of-thumb consumer behavior during disasters as well as from a stronger precautionary savings motive of optimizing consumers.

Topics: Business fluctuations and cycles; Coronavirus disease (COVID-19); Econometric and statistical methods

JEL codes: E21, C23

Résumé

Dans cette étude, nous utilisons un vaste ensemble de données historiques (de 1870 à 2016) provenant de seize économies industrielles afin de montrer que, lors de catastrophes macroéconomiques (comme une guerre, une pandémie ou une dépression), la consommation et le revenu agrégés sont beaucoup moins dissociés qu'en temps normal. Autrement dit, en période de crise, la contrainte budgétaire intertemporelle des consommateurs est plus forte, ce qui se traduit par une diminution structurelle du lissage de la consommation. Bien que nous observions cette tendance pour la pandémie actuelle de COVID-19, ce n'est habituellement pas le cas durant les récessions d'après-guerre plus classiques. Nous obtenons nos résultats grâce à une méthode de régression prévisionnelle qui découle directement du caractère prospectif de la théorie de la consommation. En employant un modèle qui met en relation les épargnants et les personnes qui ne sont pas en mesure d'épargner, nous montrons que nos conclusions peuvent être interprétées comme provenant de l'adoption plus répandue de comportements non prospectifs de la part des consommateurs lors de catastrophes ainsi que d'un motif d'épargne de précaution plus sérieux chez les consommateurs ayant un comportement d'optimisation.

Sujets : Cycles et fluctuations économiques; Maladie à coronavirus (COVID-19); Méthodes économétriques et statistiques

Codes JEL : E21, C23

1 Introduction

The COVID-19 pandemic and the lockdown measures implemented to contain it in countries around the world have triggered significant changes in consumption and saving. Large downward shifts have been reported in the propensities to consume of US and European households during 2020 (see e.g., Dossche and Zlatanos, 2020; Vandenbroucke, 2021). Currently, the Russian invasion of Ukraine and its potential geopolitical and economic repercussions—for example, escalation of the conflict, political instability, ever-increasing prices, supply chain disruptions, energy crises and even famines—also have the potential to strongly affect consumption around the world. Indeed, past macroeconomic disaster episodes, such as pandemics, wars and depressions, have been characterized by drastic declines in private consumption far above and beyond what we typically observe during ordinary recessions (see e.g., Barro and Ursúa, 2008; Nakamura et al., 2013).

This paper investigates whether rare macroeconomic disaster periods are characterized by reductions in welfare-optimizing consumption smoothing opportunities. The drastic falls in consumption that occur during disasters which, as noted by Barro and Ursúa (2008), often coincide with large reductions in income, suggest that the relationship between consumption and income may be substantially different during these times.¹ In particular, it is conceivable that consumption and income are more in tandem during disasters, implying less decoupling of consumption from income and, as a consequence, less consumption smoothing. Consumption smoothing may falter during extreme crises that confront consumers with large income shocks due to market incompleteness. During wars and depressions, it is much more likely that banks and financial markets malfunction or fail. For example, the causes and consequences of the numerous US bank failures during the Great Depression have been discussed extensively in the literature (see e.g., Richardson, 2013, and references therein). As for financial markets, Silber (2005) discusses the liquidity repercussions of the temporary New York Stock Exchange (NYSE) suspension during World War I, while Frey and Kucher (2000) document drastic government bond price reductions in Austria, Belgium, France and Germany at the onset of World War II.² And while banks appear to have been more resilient during the COVID-19 pandemic, the literature does document severe COVID-related financial market disruptions, in particular in corporate bond markets (see Goldstein et al., 2021, and references therein). With respect to the consequences of market incompleteness for consumers, Parker and Preston

¹Rare historical macroeconomic disaster periods are defined by Barro and Ursúa (2008) as peak-to-trough cumulative declines in GDP and/or private consumption of at least 10%.

²An impression of the breaking down of financial markets in continental Europe during World War I is given by a comment in a contemporaneous newspaper article, i.e., “...in Belgium, where people with securities in their pockets, and fleeing from war and starvation, sold them for cash at thirty and forty percent discount to some itinerant peddler” (Wall Street Journal, January 7th 1915).

(2005) note that missing markets to transfer consumption across time imply credit constraints, while missing markets to transfer consumption across states of the world imply precautionary saving motives. Hence, a priori, the channels through which a reduction in consumption smoothing could occur during disasters are consumers facing tighter credit constraints and consumers having a stronger precautionary saving motive.

This paper's contribution is both theoretical and methodological. Theoretically, we investigate the impact of disasters on consumption smoothing using a standard consumer intertemporal budget constraint (IBC) that relates consumption and income in the long run. More specifically, the IBC can be expressed as a relationship between the current log consumption-income ratio and future income and consumption growth rates. The question is then how disasters affect this relationship. Does the current log consumption-income ratio have a stronger or a weaker link with future income and consumption growth rates during disasters? If the link is stronger (weaker), then the IBC holds more (less) strictly and consumption and income are less (more) decoupled during disasters. When consumption and income are less (more) decoupled, this is indicative of a lower (higher) degree of consumption smoothing. Subsequently, we investigate the channels through which consumption smoothing is affected during disasters by imposing structure on consumer behavior. To this end, we consider a savers-spenders type of model in the spirit of Mankiw (2000), where the spenders consume according to their current income—for instance, because they face binding credit constraints—and where the savers optimize and have a precautionary saving motive. Changes in the structural parameters of this model—in particular, in the parameters governing rule-of-thumb behavior and the strength of the precautionary motive—affect the link between the current log consumption-income ratio and future income and consumption growth rates and the degree of consumption smoothing.

Methodologically, we estimate the IBC-implied relationship between the consumption-income ratio and future income and consumption growth rates using predictive regressions. In effect, according to the IBC, the log consumption-income ratio should have predictive power for future income and/or consumption growth rates. To find out whether disaster episodes affect the IBC and the degree of consumption smoothing, we then check whether the predictive power of this ratio differs between normal times and disaster periods. While less common in macroeconomics, our focus on predictive regressions borrows from the finance literature (see e.g., Cochrane, 2005).³ Its advantage is that the predictive regression equations follow immediately from the forward-looking nature of consumption theory; that is, both the IBC and the savers-spenders set-up that we consider imply that the log consumption-income ratio depends on (expected) future variables and therefore should have predictive power for these variables. This approach

³For an application in international macro, that is, in exchange rate modeling, we refer to Sarno and Schmeling (2014).

is also advantageous because it avoids complications related to causal inference as the error term in a predictive regression is expected to be orthogonal to the regressor.⁴ Violations of this orthogonality are possible but are, as we show in the paper, easily dealt with. With respect to the data, we use historical annual data over the period 1870–2016 for sixteen industrial economies to estimate cross-country panel predictive regressions. These data are taken from Jordà et al. (2016)’s macrohistory database. Furthermore, Barro and Ursúa (2008), and more recently Nakamura et al. (2013), provide start and end dates of rare disasters over the full period for all countries in the sample. The use of historical data, while not without complications, is motivated by the absence of major macro disasters in most postwar industrial economies.⁵ While the magnitude of the shifts in consumption and income during the COVID-19 pandemic is comparable to that of the shifts that occurred during historical disaster episodes, the COVID-19 pandemic constitutes only one (ongoing) crisis and, as such, provides insufficient information from which to draw general conclusions about consumption smoothing during disasters. We do, however, check whether our historical findings also hold for the ongoing COVID-19 pandemic by supplementing our results based on historical data with results obtained from recent quarterly data for twenty industrial countries over the period 1995Q1–2021Q4. With respect to the estimation method, we use a variety of mean-group (MG) estimators that, facilitated by the lengthy series at our disposal, allow for full parameter heterogeneity to obtain estimates for the average predictive effects across countries (see e.g., Pesaran and Smith, 1995; Pesaran, 2006; Chudik and Pesaran, 2015).

Our findings suggest that the predictive ability of the log consumption-income ratio for future income and consumption growth rates is significantly higher during macroeconomic disaster episodes. This new result survives a battery of robustness checks. While it holds both for historical disaster episodes and for the ongoing COVID-19 pandemic, we do not find evidence of a structural change in predictability when looking at conventional postwar recessions. Interpreted through the lens of the theory, it implies that the IBC holds more strictly and that consumption and income are significantly less decoupled during disaster episodes. This, in turn, points to a structural decrease in consumption smoothing during disasters, the like of which cannot be observed during ordinary recessions. From our savers-spenders framework, we argue that the increased predictive power during disasters of the log consumption-income ratio for future income growth rates can be attributed to an increase in rule-of-thumb consumption behavior. In addition, the increased predictive power of the log consumption-income ratio for future consumption growth rates can be attributed to a stronger precautionary saving motive of the optimizing consumers. Importantly, the precaution result does not (merely) pertain to additional saving resulting from higher

⁴As argued by Cochrane (2005, page 392), predictive regressions do not have causes on the right and effects on the left. Rather, regressions are run with the variable that is orthogonal to the (prediction) error on the right.

⁵An important complication when using historical data is measurement error, for which we control in our estimations.

uncertainty during disasters, a finding that is commonly reported for more conventional recessions (see e.g., Mody et al., 2012). Rather, it refers to a structural shift, namely that *a given amount* of uncertainty has a larger impact on saving during these rare crises. As our evidence suggests that disasters are characterized both by tighter credit constraints—that is, the increased rule-of-thumb behavior—and by stronger precautionary saving motives, it supports an incomplete markets-based interpretation of the decrease in consumption smoothing observed during rare disasters.

A large literature focuses on the asset pricing and, to a lesser extent, business cycle implications of the presence of ex-ante disaster risk in the economy (see e.g., Rietz, 1988; Barro, 2006, 2009; Gourio, 2012; Barro and Ursúa, 2012; Nakamura et al., 2013; Gillman et al., 2015; Farhi and Gabaix, 2016). Our paper, in contrast, contributes to a growing literature that looks at the ex-post macroeconomic effects of rare disasters (i.e., the effects of both historical and current disasters on the real economy, in particular on consumption and saving). Jordà et al. (2020), for example, use aggregate time series data for European countries going back to the 14th century and argue that historical pandemics lead to medium-term reductions in the real (natural) interest rate that are potentially caused by increases in (precautionary) saving. Coibion et al. (2020) use US survey data to investigate how local lockdown measures implemented in reaction to COVID-19 affect consumer spending and the macroeconomic expectations of households. Ludvigson et al. (2020) use US aggregate monthly data to study the dynamic and potentially non-linear impact of (natural) disaster shocks on industrial production, employment and uncertainty. Using US micro data, Levine et al. (2021) observe a large flow of household deposits into bank branches in US counties with high COVID infection rates, and attribute this to precautionary saving motives. We note that our paper is also related to a literature that investigates how and why conventional recessions affect consumption and saving (see e.g., Mody et al., 2012; Alan et al., 2012; Adema and Pozzi, 2015; Carroll et al., 2019, who generally point to the role of higher uncertainty during recessions). Our paper deviates from these studies by focussing on shifts that characterize consumption behavior—that is, consumption smoothing—during macro disasters and that do not seem to be present during ordinary recessions.

The outline of the paper is as follows. Section 2 discusses the IBC, its predictability implications and how it is related to consumption smoothing. Section 3 presents the results of estimating cross-country panel predictive regressions to investigate the IBC and consumption smoothing, both during normal times and during disaster episodes. Section 4 proposes a savers-spenders model to give a theoretical interpretation to our predictability findings; that is, we look at channels that affect consumption smoothing and provide additional empirical evidence to support these channels. Section 5 concludes.

2 The intertemporal budget constraint and consumption smoothing

In this section, we discuss the intertemporal budget constraint (IBC) and show that it implies that the log consumption-income ratio has predictive power for future income and consumption growth rates. We then discuss what different degrees of predictability imply for consumption smoothing.

2.1 The intertemporal budget constraint and the predictive ability of the consumption-income ratio

The intertemporal budget constraint (IBC) implies that we can write the period t log consumption to income ratio $c_t - y_t$ (up to a constant and an approximation error) as,

$$c_t - y_t = \sum_{j=1}^{\infty} \rho^j [E_t(\Delta y_{t+j}) - E_t(\Delta c_{t+j})] \quad (1)$$

(see Campbell and Mankiw, 1989), where ρ is the discount factor (with $0 < \rho < 1$), E_t is the expectations operator conditional on period t information, c_t is the log of real consumption C_t , y_t is the log of real total income Y_t , which equals the sum of labor and capital income. Refer to Appendix A for the derivation.⁶

The intuition behind eq.(1) is straightforward. In ex-post form (i.e., without the expectations operator E_t), the budget constraint tells us that a high current consumption-income ratio (or, conversely, a low current saving ratio) coincides with high future income growth rates and/or low future consumption growth rates, while a low current consumption-income ratio (or, conversely, a high current saving ratio) coincides with low future income growth rates and/or high future consumption growth rates. In ex ante form, the budget constraint tells us that expected future income decreases and expected future consumption increases lower the current consumption-income ratio (or, conversely, augment the current saving ratio) while, expected future income increases and expected future consumption decreases augment the current consumption-income ratio (or, conversely, lower the current saving ratio).

Importantly, eq.(1) implies that the log consumption-income ratio $c_t - y_t$ may have predictive ability for future income and consumption growth rates. To see this, we first write eq.(1) in ex-post form (i.e., without the expectations operator E_t) and then write the variance of $c_t - y_t$ as,

$$V(c_t - y_t) = \sum_{j=1}^{\infty} \rho^j [cov(c_t - y_t, \Delta y_{t+j}) - cov(c_t - y_t, \Delta c_{t+j})] \quad (2)$$

where $V(\cdot)$ denotes the variance and $cov(\cdot)$ denotes the covariance. This equation shows that, unless $c_t - y_t$ is constant so that $V(c_t - y_t) = 0$, $c_t - y_t$ has predictive power for either future income growth

⁶The derivation includes a more general expression for $c_t - y_t$ that includes expected real rates of return on wealth. We note that the IBC-based link between the current log consumption-income ratio and expected future returns is ambiguous and not substantial if the discount factor for future income growth rates is close to that of future consumption growth rates.

rates, future consumption growth rates or both. We refer to Cochrane (2005, pages 398-399) for a similar argument in the context of asset pricing.⁷ We can therefore write the following predictive equations for Δy_{t+j} and Δc_{t+j} ,

$$\Delta y_{t+j} = \phi_j^y (c_t - y_t) + \eta_{t+j}^y \quad (3)$$

$$\Delta c_{t+j} = \phi_j^c (c_t - y_t) + \eta_{t+j}^c \quad (4)$$

with error terms η_{t+j}^y and η_{t+j}^c . These prediction errors should, in principle, be orthogonal to the predictor variable $c_t - y_t$ but, as we discuss in Section 3, there are reasons why this may not be the case. Furthermore, we note that the IBC itself does not impose restrictions on the coefficients ϕ_j^y and ϕ_j^c for particular horizons j . In general, however, the predictive ability is expected to be positive for future income growth rates and/or negative for future consumption growth rates (i.e., we generally expect $\phi_j^y > 0$ and/or $\phi_j^c < 0$). Moreover, we expect that, in absolute value, the coefficients ϕ_j^y and ϕ_j^c are decreasing with the horizon j . These expectations are confirmed by our empirical evidence reported below.

We note that by subtracting eq.(4) from eq.(3), we obtain a predictive equation for the income-consumption growth differential $\Delta y_{t+j} - \Delta c_{t+j}$,

$$\Delta y_{t+j} - \Delta c_{t+j} = \phi_j (c_t - y_t) + \eta_{t+j} \quad (5)$$

where $\phi_j = \phi_j^y - \phi_j^c$ and $\eta_{t+j} = \eta_{t+j}^y - \eta_{t+j}^c$. For $\phi_j^y > 0$ and/or $\phi_j^c < 0$, we generally expect $\phi_j > 0$.

2.2 Predictability and consumption smoothing

The magnitude of the coefficients ϕ_j^y and ϕ_j^c is informative about the horizon over which the IBC holds. When ϕ_j^y and ϕ_j^c are close to zero, the current consumption-income ratio coincides with relatively small future adjustments in income and consumption (i.e., the IBC holds more loosely over a longer horizon). Hence, the decoupling episodes between c_t and y_t , that is, the deviations from the long-run equilibrium implied by the IBC, are more prolonged. More prolonged saving and dissaving episodes, in turn, suggest more consumption smoothing. On the other hand, when the coefficients ϕ_j^y are more positive or when the coefficients ϕ_j^c are more negative, the current consumption-income ratio coincides with relatively large future adjustments in income and consumption (i.e., the IBC holds more strictly over a shorter horizon). Hence, the decoupling episodes between c_t and y_t , that is, the deviations from the long-run equilibrium implied by the IBC, are less prolonged. Less prolonged saving and dissaving episodes, in turn, suggest less consumption smoothing.⁸

⁷In asset pricing, this concerns the predictive ability of the equity price-dividend ratio for future returns and/or future dividend growth rates.

⁸An alternative way to look at our set-up is by noting that if Δy_{t+1} and Δc_{t+1} are stationary, then, given eq.(1), $c_t - y_t$ should also be stationary and c_t and y_t are cointegrated. By Engle and Granger (1987), there then exists an error

In the next section, we empirically investigate how macroeconomic disasters affect the predictive power of $c_t - y_t$ for future income and consumption growth rates. Our main finding is that, during macro disasters, the log consumption-income ratio has a more positive predictive impact on future income growth rates, while it has a more negative predictive impact on future consumption growth rates. This suggests that the IBC holds more strictly and consumption and income are less decoupled during macroeconomic disaster episodes. This, in turn, points to a structural reduction in consumption smoothing during these crises. In Section 4, we impose additional theoretical structure on our set-up by explicitly specifying consumption behavior and we give a model-based interpretation to the reduction in consumption smoothing observed during disasters.

3 Empirical results

In this section, we investigate whether disaster episodes affect the IBC by looking at the predictive ability of the log consumption-income ratio for future income and consumption growth rates during both disasters and more normal times. To this end, we estimate predictive regressions for a panel of industrial economies.

3.1 Data

For most estimations, we use long-term historical macro data over the period 1870–2016. These are available at the annual frequency. Data availability determines the countries included in the dataset and the periods considered per country.⁹ Our sample consists of sixteen economies (i.e., $N = 16$). These are Australia, Belgium, Denmark, Finland, France, Germany, Italy, Japan, the Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, the UK and the US. For c_t , we use the log of per capita real consumption, while for y_t we use the log of per capita real GDP. Per capita real personal consumer expenditures and per capita real GDP are taken from the Jordà-Schularick-Taylor Macrohistory Database (Jordà et al., 2016).¹⁰

To investigate the impact of macroeconomic disasters on the predictive ability of the log consumption-income ratio, we construct country-specific disaster dummies that take on the value of one during disaster correction model between c_t and y_t where deviations from the long-run equilibrium relationship between c_t and y_t implied by the IBC affect the next period's values of c_t and y_t . Hence, our eqs.(3)-(4) written for $j = 1$ can be considered an error correction model, with the predictability parameters ϕ_1^y and ϕ_1^c reflecting the speed of adjustment towards equilibrium. A more positive predictability parameter ϕ_1^y or a more negative parameter ϕ_1^c implies a faster adjustment towards equilibrium, that is, less prolonged saving and dissaving episodes.

⁹For some countries and variables, a number of data points are missing at the beginning of the sample period that renders the panel unbalanced.

¹⁰The website is <http://www.macrohistory.net/data>. The series have codes “rconpc” and “rgdppc”. We note that the series that we use are both expressed as indices.

episodes. They are constructed from the macroeconomic disaster episodes identified by Barro and Ursúa (2008). The authors define a disaster as a peak-to-trough cumulative decline in real per capita GDP and/or real per capita personal consumer expenditure of at least 10%. We construct a general dummy that contains all identified disaster episodes over the sample period. Additionally, we also consider specific disaster episodes. In particular, we construct dummies for each of the four principal world economic crises identified by Barro and Ursúa (2008): World War I (WW1), the Spanish flu pandemic of the late 1910s/early 1920s (PAN), the Great Depression (GRD) and World War II (WW2). More details on the construction of the disaster dummies are provided in Appendix B.

3.2 Baseline setting and results

3.2.1 Specification and method

Our discussion in the previous sections suggests that the current log consumption-income ratio may have predictive power for future income and consumption growth rates and that this predictive ability may be different during disaster episodes. To check this empirically, we estimate the following baseline specification,

$$x_{i,t+1} = \mu_i + \alpha_i d_{it} + \beta_i (c_{it} - y_{it}) + \gamma_i (c_{it} - y_{it}) d_{it} + \epsilon_{i,t+1} \quad (6)$$

where $x_{i,t+1}$ is the predicted variable in period $t + 1$ in country i (with $i = 1, \dots, N$), μ_i is a country fixed effect, d_{it} is a country-specific dummy variable that is equal to zero in normal times and equal to one during disaster episodes, $c_{it} - y_{it}$ is the log consumption-income ratio, and $\epsilon_{i,t+1}$ is the error term. Given the relatively long time series at our disposal for every country i , we allow for heterogeneity across countries in all slope coefficients.

With respect to the regressors of interest, from IBC logic, we expect that the current log consumption-income ratio $c_{it} - y_{it}$ has a positive impact on next period's income growth rate $\Delta y_{i,t+1}$. If during macroeconomic disaster episodes the IBC holds more strictly, we further expect that this predictive ability is higher—that is, more positive—during such episodes. As such, for $x_{i,t+1} = \Delta y_{i,t+1}$, we expect $\beta_i > 0$ and $\gamma_i > 0$. On the other hand, from IBC logic, we expect that the current log consumption-income ratio $c_{it} - y_{it}$ has a negative impact on next period's consumption growth rate $\Delta c_{i,t+1}$. If during macroeconomic disaster episodes the IBC holds more strictly, we further expect that this predictive ability is higher—that is, more negative—during such episodes. As such, for $x_{i,t+1} = \Delta c_{i,t+1}$, we expect $\beta_i < 0$ and $\gamma_i < 0$. We further add the disaster dummy separately to eq.(6) to control for a potential predictive impact of disasters on the dependent variable that is unrelated to the predictive impact of the consumption-income ratio.

The error term $\epsilon_{i,t+1}$ is a prediction error that should, in principle, be unpredictable based on period t information. It is nonetheless possible that it is autocorrelated, however, where the autocorrelation is of the moving average (MA) type. For example, it could follow an MA(1) process due to measurement error or time aggregation in the data.^{11,12} Further complications include the possibility that the error term is correlated across countries (cross-sectional dependence) and that it is correlated with the included regressors. These complications are dealt with in the robustness checks discussed in Section 3.3.

For the baseline results reported below, we estimate eq.(6) country-by-country using ordinary least squares (OLS). Pesaran and Smith (1995) then show that for a heterogeneous (dynamic) panel with country-specific parameter vector Ψ_i and with a sufficiently large T and N , consistent estimates of the average effects $\bar{\Psi} = N^{-1} \sum_{i=1}^N \Psi_i$ can be obtained by averaging over the country-specific coefficient estimates, that is, $\hat{\bar{\Psi}} = N^{-1} \sum_{i=1}^N \hat{\Psi}_i$. The average over the N country-specific OLS estimates is referred to as the mean-group (MG) estimator. It is consistent provided that the country-specific coefficients are consistently estimated by OLS. Following Pesaran et al. (1996), the asymptotic covariance matrix Σ for the mean-group estimator is consistently estimated nonparametrically by,

$$\hat{\Sigma} = \frac{1}{N-1} \sum_{i=1}^N \left(\hat{\Psi}_i - \hat{\bar{\Psi}} \right) \left(\hat{\Psi}_i - \hat{\bar{\Psi}} \right)' \quad (7)$$

Finally, we note that our estimation method requires regression equations that include stationary variables. This is true for the OLS-based estimations considered in this section, but also for results obtained using alternative estimators like the common correlated effects (CCE) or instrumental variables (IV) estimators that are applied in the robustness checks. The log consumption-income ratio $c_{it} - y_{it}$, while expected to be stationary on theoretical grounds, is the one variable in our estimations for which stationarity is not immediately evident in the data. In Appendix C, we report the results of panel unit root tests applied to this variable. From these tests, we conclude that, over the historical period 1870–2016, the regressor $c_{it} - y_{it}$ is stationary for the overall panel and for a majority of countries in our sample. In the robustness checks discussed in Section 3.3, we check whether the presence of a unit root in $c_{it} - y_{it}$ in some countries has an impact on our findings by conducting estimations that use a stochastically detrended version of the log consumption-income ratio.

3.2.2 Baseline results

Turning to our findings, Table 1 (columns 2-4) presents the baseline results from estimating eq.(6) for the sixteen economies in our sample over the period 1870–2016 with $x_{i,t+1} = \Delta y_{i,t+1}$, $x_{i,t+1} = \Delta c_{i,t+1}$ and

¹¹See Sommer (2007), for example, for measurement error in aggregate consumption data and its implications.

¹²The error term $\epsilon_{i,t+1}$ can also be conditionally heteroskedastic (see e.g., Hamilton, 2008; Nakamura et al., 2017, who document changes over time in the volatilities of macroeconomic variables like GDP growth).

$x_{i,t+1} = (\Delta y_{i,t+1} - \Delta c_{i,t+1})$. The table reports the OLS-based mean-group estimates of the coefficients α_i , β_i and γ_i and their corresponding standard errors calculated from eq.(7). The country-specific coefficient estimates β_i and γ_i that are used in the calculation of the mean-group coefficient estimates for the regressors $c_{it} - y_{it}$ and $(c_{it} - y_{it})d_{it}$ are reported in Appendix D. Table 1 further reports the average Cumby and Huizinga (1992) autocorrelation test and its corresponding p-value which tests the null hypothesis that there is no autocorrelation in the error term.^{13,14}

From the baseline results reported in the table, we note the following. First, a look at Cumby and Huizinga (1992)'s test for autocorrelation shows that the null hypothesis of no autocorrelation is not rejected for any of the conducted regressions. Second, while it can be expected that the disaster dummy d negatively affects income and consumption growth *in the same period*, the reported results show that it also negatively affects the next period's income growth. It has no predictive impact for consumption growth, however. Third, in accordance with the discussion in Section 2 of the IBC and its predictability implications, the log consumption-income ratio $c - y$ has significant positive predictive ability for the next period's income-consumption growth differential. The separate results for Δy and Δc as dependent variables then show that this stems mainly from the significant predictive power that $c - y$ has for the consumption growth rate where the sign of the coefficient on $c - y$ is in accordance with IBC logic (i.e., a high consumption-income ratio today is followed by future decreases in consumption growth).¹⁵ Finally, from the estimated coefficients on the regressor $(c - y)d$, we note that the predictive ability of $c - y$ for *both* Δy and Δc is significantly higher during disasters as opposed to normal times; that is, during disasters $c - y$ has a positive predictive impact on Δy and a more negative predictive impact on Δc . Whereas during normal times a 1% increase in $\frac{C}{Y}$ implies a next period increase in Δy of only one basis point on average (across time and countries) and a next period decrease in Δc of only three basis points, these numbers equal twelve, respectively seventeen basis points during disaster episodes. Interpreted through

¹³More specifically, it tests the null hypothesis that the error term follows a moving average process of known order $q \geq 0$ against the alternative that the autocorrelations of the error term are nonzero at lags greater than q . Most statistics reported in this paper are for $q = 0$. We note that this test is particularly suitable as, besides allowing to test for MA errors, it provides an autocorrelation test that is valid also if the errors are conditionally heteroskedastic. Moreover, it can also be applied when using estimators other than OLS, such as IV (see Cumby and Huizinga, 1992, for details).

¹⁴We calculate the statistic per country and then average it across countries. The Cumby and Huizinga (1992) test statistic follows a χ^2 distribution. Assuming that the country-specific test statistics are independent, the average Cumby and Huizinga (1992) test still follows a χ^2 distribution with the same number of degrees of freedom as its country-specific counterparts.

¹⁵The coefficient on the regressor $c - y$ is a semi-elasticity. For example, for the coefficient of Δy on $c - y$, we have $\frac{\partial \Delta y}{\partial (c - y)} = \frac{\partial \Delta y}{\partial \ln(\frac{C}{Y})}$ (i.e., the coefficient equals the change in Δy divided by the percentage change in $\frac{C}{Y}$). A coefficient equal to 0.1 then implies that if $\frac{C}{Y}$ increases with 1% (e.g., from 100% to 101%), then Δy increases with 0.1 percentage points (e.g., from 1% to 1.1%). A coefficient equal to 1 then implies that if $\frac{C}{Y}$ increases with 1% (e.g., from 100% to 101%), then Δy increases with 1 percentage point (e.g., from 0.5% to 1.5%).

the lens of the intertemporal budget constraint, these findings suggest that the IBC holds more strictly and that there is substantially less decoupling between consumption and income during disaster episodes. This, in turn, points to a reduction in consumption smoothing during disasters.

Table 1: Baseline results: OLS-based mean-group estimates

	Baseline results			With lagged dependent variable		
	Dependent variable $x_{i,t+1}$			Dependent variable $x_{i,t+1}$		
	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$
d_{it}	-0.043*** (0.014)	-0.014 (0.016)	-0.029 (0.022)	-0.036** (0.016)	-0.010 (0.016)	-0.023 (0.023)
$(c_{it} - y_{it})$	0.014 (0.026)	-0.035* (0.020)	0.049*** (0.014)	0.013 (0.027)	-0.038* (0.020)	0.052*** (0.013)
$(c_{it} - y_{it})d_{it}$	0.111* (0.058)	-0.136* (0.078)	0.247*** (0.062)	0.115* (0.066)	-0.158** (0.081)	0.270*** (0.060)
x_{it}				0.048 (0.045)	0.030 (0.049)	0.096** (0.047)
Cumby-Huizinga AC	2.503 [0.286]	3.902 [0.142]	2.243 [0.326]	2.320 [0.314]	2.502 [0.286]	2.244 [0.326]

Notes: Reported are the mean-group results based on OLS estimation of eq.(6) (baseline results) and eq.(8) (results with lagged dependent variable). Estimation is based on panel data for sixteen countries over the period 1870–2016. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation tests, testing the null hypothesis of no autocorrelation against the alternative that the autocorrelations of the error term are nonzero at lags greater than zero (with maximum lag equal to two).

To conclude, our baseline results suggest that the predictive ability of the consumption-income ratio for the next period's income and consumption growth rates is significantly higher during disaster episodes. In the following section, we first conduct a number of robustness checks to more firmly establish our empirical finding. In Section 3.4, we then check to what extent our results hold up at longer horizons, while in Section 3.5, we investigate whether we can draw the same conclusions when looking at ordinary recessions. Finally, we look at the predictive impact of the log consumption-income ratio during specific disaster episodes; that is, we look at major historical disaster periods in Section 3.6 and at the current COVID-19 pandemic in Section 3.7.

3.3 Robustness checks

This section checks the robustness of our baseline results with respect to the regression equation specification, estimation methodology and variables included in the regression equation.

Lagged dependent variable

Our first robustness check consists of looking at a dynamic panel setting where the regression equation includes a lag of the dependent variable under consideration. Controlling for the lagged dependent variable is useful to make sure that, when detecting a relationship between $c_{it} - y_{it}$ and the dependent variable $x_{i,t+1}$, this relationship is not driven solely by the combination of an autocorrelated $x_{i,t+1}$ variable and the possible covariance between $c_{it} - y_{it}$ and x_{it} (i.e., $c_{it} - y_{it}$ only affects $x_{i,t+1}$ because it is correlated with x_{it} and x_{it} has predictive power for $x_{i,t+1}$). To deal with this, we estimate an extended version of eq.(6) where one lag of the dependent variable is added as a control variable. That is, we have,

$$x_{i,t+1} = \mu_i + \alpha_i d_{it} + \beta_i (c_{it} - y_{it}) + \gamma_i (c_{it} - y_{it}) d_{it} + \delta_i x_{it} + \epsilon_{i,t+1} \quad (8)$$

where $x_{i,t+1} = \Delta y_{i,t+1}, \Delta c_{i,t+1}, (\Delta y_{i,t+1} - \Delta c_{i,t+1})$. We add only one lag of the dependent variable because, when conducting estimations with more lags, we find that the coefficient estimates on additional lags are not significant.

Table 1 (columns 5-7) presents the OLS-based mean-group estimates obtained from estimating eq.(8) using our historical sample.¹⁶ While the significance of the impact of the regressors $c - y$ and $(c - y)d$ is somewhat higher compared to our baseline results, our findings are generally not affected much when including a lagged dependent variable to the regression equation (which itself enters the regression equation significantly only in column 7).

Detrended log consumption-income ratio

While panel unit root tests suggest that the regressor $c - y$ is stationary for the overall panel and for a majority of countries in our sample, for some countries the hypothesis of a unit root in $c - y$ cannot be rejected. Moreover, the number of countries for which a unit root can/cannot be rejected varies depending on the considered panel unit root test. Refer to Appendix C for details. In this section, we therefore check the robustness of our results to stochastically detrending the variable $c - y$. We detrend $c - y$ by calculating the deviation of $c - y$ from its stochastic trend $\overline{c - y}$. The latter is approximated by a ten-year moving average as $\overline{c - y} = \frac{1}{10} \sum_{j=0}^9 (c_{-j} - y_{-j})$.¹⁷

¹⁶Since T is large, the time series bias in OLS—and, therefore, in MG—that results from including the lagged dependent variable to the specification can be considered negligible.

¹⁷This detrending approach takes out the low frequency movements (i.e., long swings) in the data as well as high frequency noise as opposed to a first-differencing approach that takes out low- and medium-frequency movements (see e.g., Sarno and Schmeling, 2014). Note that if we proxy the stochastic trend using a moving average calculated over five or twenty years instead of ten years, we obtain similar results.

Table 2: Results using detrended $c - y$ variable: OLS-based mean-group estimates

	Stoch. detrended $c - y$ for all countries			Stoch. detrended $c - y$ for a subset of countries		
	Dependent variable $x_{i,t+1}$			Dependent variable $x_{i,t+1}$		
	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$
d_{it}	-0.024 (0.022)	-0.053*** (0.011)	0.029 (0.031)	-0.023 (0.023)	-0.032** (0.015)	0.010 (0.033)
$(c_{it} - y_{it})$	0.060 (0.058)	-0.042 (0.044)	0.102** (0.050)	0.004 (0.029)	-0.033 (0.024)	0.036** (0.015)
$(c_{it} - y_{it})d_{it}$	0.082 (0.083)	-0.171** (0.068)	0.254** (0.105)	0.131* (0.071)	-0.187** (0.078)	0.318*** (0.095)
Cumby-Huizinga AC	3.261 [0.196]	3.838 [0.147]	2.734 [0.255]	2.643 [0.267]	3.846 [0.146]	2.214 [0.330]

Notes: Reported are the mean-group results based on OLS estimation of eq.(6). Columns 2-4 report the results based on using the stochastically detrended log consumption-income ratio for all sixteen countries in the sample. Columns 5-7 report the results based on using the stochastically detrended log consumption-income ratio only for those countries for which the hypothesis of a unit root in $c - y$ cannot be rejected. For the remaining countries, the actual variable $c - y$ is used after removal of a deterministic linear time trend (if present). Estimation is based on panel data for sixteen countries over the period 1870–2016. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation tests, testing the null hypothesis of no autocorrelation against the alternative that the autocorrelations of the error term are nonzero at lags greater than zero (with maximum lag equal to two).

The results obtained with this stochastically detrended version of the log consumption-income ratio are presented in Table 2. In columns 2-4, we report the results of estimating eq.(6) where the detrended version of $c - y$ is used for all sixteen countries in the sample. The results obtained for the dependent variables Δc and $(\Delta y - \Delta c)$ are generally in accordance with the baseline results from Table 1. For the dependent variable Δy however, the positive impact of $c - y$ during disasters is no longer significant. This may be the result of unnecessarily throwing away potentially relevant information when detrending $c - y$ variables that, according to our unit root tests, are not stochastically trended. We therefore also estimate eq.(6) with the stochastically detrended version of $c - y$ used only for those countries for which a unit root in $c - y$ cannot be rejected.¹⁸ For the other countries, we use the actual $c - y$ variable (albeit cleansed from a deterministic linear time trend, if one is present). The results are presented in columns 5-7 of Table 2 and now unequivocally support our baseline results.

From the findings reported in this section, we generally conclude that our predictability results are not driven by the presence of stochastic or deterministic trends in the regressor $c - y$.

¹⁸Based on the CIPS and CIPS* panel unit root tests that control for the presence of cross-sectional dependence and a deterministic linear time trend in $c - y$, the six countries (out of sixteen) for which the null hypothesis of a unit root cannot be rejected are Denmark, Finland, Germany, Japan, Norway and Portugal. See Appendix C for details.

Cross-sectional dependence

Our baseline estimations do not control for cross-sectional dependence in the regression error term.¹⁹ The latter may be caused by unobserved factors that are common across countries. Examples of common factors are international business or financial cycles or changes in trade or financial integration that occur simultaneously in most or all countries of the sample. Ignoring these common factors may imply less efficient estimation and, more seriously, may lead to biased and inconsistent OLS estimates if the unobserved common factors are correlated with the regressors. To control for unobserved common factors, we consider the following specification,

$$x_{i,t+1} = \mu_i + \alpha_i d_{it} + \beta_i (c_{it} - y_{it}) + \gamma_i (c_{it} - y_{it}) d_{it} + \kappa_i f_{t+1} + \epsilon_{i,t+1} \quad (9)$$

where $x_{i,t+1} = \Delta y_{i,t+1}, \Delta c_{i,t+1}, (\Delta y_{i,t+1} - \Delta c_{i,t+1})$ and where the regression equation now includes a vector of unobserved common factors f_{t+1} with a corresponding vector of country-specific factor loadings κ_i . To estimate eq.(9), we follow Pesaran (2006) and use cross-sectional averages of the dependent variable and all regressors as proxies for f_{t+1} . After replacing f_{t+1} by these cross-sectional averages, we estimate eq.(9) country-by-country using OLS. This is the common correlated effects (CCE) estimator. The average over the N country-specific CCE estimates is referred to as the common correlated effects mean group (CCEMG) estimator. For a dynamic setting such as ours, Chudik and Pesaran (2015) propose to additionally include lagged cross-sectional averages of the dependent variable and the regressors. In this case, we obtain N country-specific dynamic CCE estimates from which we calculate the dynamic CCEMG estimator. Standard errors of both mean-group estimators are calculated from eq.(7).

The results of estimating eq.(9) using the standard and the dynamic CCEMG estimator for the sixteen economies in our sample over the period 1870–2016 are presented in Table 3. The estimation results are in accordance with our baseline results as we find that the predictive ability of the consumption-income ratio for future income and consumption growth rates is significantly higher during disasters (i.e., more positive for income growth and more negative for consumption growth).

¹⁹When testing explicitly for cross-sectional dependence in the error terms of our baseline specification, we reject cross-sectional independence. These results are not reported but are available upon request.

Table 3: Results controlling for cross-sectional dependence: CCE-based mean-group estimates

	CCEMG estimator			dynamic CCEMG estimator		
	Dependent variable $x_{i,t+1}$			Dependent variable $x_{i,t+1}$		
	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$
d_{it}	-0.066*** (0.018)	-0.003 (0.016)	-0.052** (0.022)	-0.065*** (0.018)	-0.020 (0.023)	-0.035 (0.028)
$(c_{it} - y_{it})$	0.076* (0.039)	-0.149*** (0.030)	0.216*** (0.057)	0.047 (0.029)	-0.181*** (0.047)	0.213*** (0.062)
$(c_{it} - y_{it})d_{it}$	0.150** (0.065)	-0.134** (0.060)	0.290*** (0.066)	0.157** (0.067)	-0.117* (0.071)	0.270*** (0.081)
Cumby-Huizinga AC	2.465 [0.292]	2.518 [0.284]	3.510 [0.173]	2.639 [0.267]	3.062 [0.216]	4.975 [0.083]

Notes: Reported are the mean-group results based on static CCE estimation (see Pesaran, 2006) and dynamic CCE estimation (see Chudik and Pesaran, 2015) of eq.(9). In the former case, we proxy the unobserved common factors f_{t+1} by adding the cross-sectional averages of the dependent variable and all regressors into the regression equation. In the latter case, we proxy the unobserved common factors f_{t+1} by adding contemporaneous values as well as lags of the cross-sectional averages of the dependent variable and all regressors into the regression equation. Given the sample size, we add five lags of each cross-sectional average. We refer to Chudik and Pesaran (2015) for details. Estimation is based on panel data for sixteen countries over the period 1870–2016. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation tests, testing the null hypothesis of no autocorrelation against the alternative that the autocorrelations of the error term are nonzero at lags greater than zero (with maximum lag equal to two).

Measurement error

The estimations so far have been conducted under the assumption that the regressors are uncorrelated with the error term. The lack of autocorrelation implied by the results of the Cumby-Huizinga tests reported in Table 1, for instance, suggests that the error term $\epsilon_{i,t+1}$ in eqs.(6) or (8) is indeed *iid*, so uncorrelated with the period t regressors. However, if a potential correlation between regressors and error term renders OLS estimation inconsistent, then the results of autocorrelation tests based on these OLS results may also be flawed. Hence, more scrutiny is needed here. We focus in particular on the case of measurement error. Measurement error most likely is present in our historical dataset and may be more important during macroeconomic disasters as it may be harder to construct GDP and its components during wars and swift economic declines. It is easy to show that if the variables y_{it} and c_{it} are measured with noise, this leads to correlation between the regressors and the error term in our regression specifications.²⁰ In this case, an instrumental variables (IV) approach is necessary. Using our

²⁰To see this, assume that the observed log income and log consumption variables are given by $y_t = \bar{y}_t + \nu_t^y$ and $c_t = \bar{c}_t + \nu_t^c$ with \bar{y}_t and \bar{c}_t denoting true log income and true log consumption and with ν_t^y and ν_t^c denoting noise terms. If for the true data we have $\Delta \bar{y}_{t+j} = \psi_j^y(\bar{c}_t - \bar{y}_t) + e_{t+j}^y$ and $\Delta \bar{c}_{t+j} = \psi_j^c(\bar{c}_t - \bar{y}_t) + e_{t+j}^c$ with $E_t(e_{t+j}^y) = 0$ and $E_t(e_{t+j}^c) = 0$, then the corresponding empirical specifications based on observed data are given by $\Delta y_{t+j} = \psi_j^y(c_t - y_t) + \varepsilon_{t+j}^y$ and $\Delta c_{t+j} = \psi_j^c(c_t - y_t) + \varepsilon_{t+j}^c$ where $\varepsilon_{t+j}^y = e_{t+j}^y + \Delta \nu_{t+j}^y + \psi_j^y \nu_t^y - \psi_j^y \nu_t^c$ and $\varepsilon_{t+j}^c = e_{t+j}^c + \Delta \nu_{t+j}^c + \psi_j^c \nu_t^y - \psi_j^c \nu_t^c$. As such,

historical sample, we therefore estimate eq.(6) country-by-country using IV and calculate the mean-group results (i.e., the average of the country-specific IV estimates across countries). Standard errors of the mean-group estimates are calculated from eq.(7). With respect to the choice of instruments, we note that log consumption-income ratio is highly persistent.²¹ It therefore makes sense to use lags of the regressors as instruments which, along with being sufficiently correlated with the regressors, can be expected to be uncorrelated with the error term. To make sure our findings are robust across instrument sets, we consider two instrument sets, one with four lags of each regressor (instrument set 1) and one with two lags of each regressor (instrument set 2). We calculate the Sargan-Hansen overidentifying restrictions statistic that tests the null hypothesis that the instruments are orthogonal to the error term. We also calculate the Cragg-Donald statistic of instrument strength, which tests the null hypothesis that the instruments are weak (i.e., that the instruments used are not sufficiently correlated with the potentially endogenous regressors). The latter test is a multivariate extension of the first-stage F statistic used to evaluate instrument strength in the case of one endogenous regressor.

The results presented in Table 4 confirm our baseline findings that macroeconomic disasters magnify the predictive impact of the log consumption-income for both future income and consumption growth rates. We further note that the magnitude and significance of the estimates is generally higher compared to the baseline results and that our findings are robust across both instrument sets. The reported statistics support the validity and quality of the instruments used. First, based on the Sargan-Hansen OR test, we cannot reject the orthogonality of instruments and the error term.²² Second, based on the Cragg-Donald WI test, we do reject the null hypothesis that the used instruments are weak.²³

there is correlation between the regressor $c_t - y_t$ and the error terms ε_{t+j}^y and ε_{t+j}^c (irrespective of the horizon $j > 0$).

²¹The OLS-based mean-group AR parameter of an AR(1) process estimated for $c_{it} - y_{it}$ equals 0.918 (with standard error 0.020). While persistent, the variable $c_{it} - y_{it}$ does not contain a unit root for a majority of countries and for the overall panel, however, as can be concluded from the panel unit root tests reported in Appendix C.

²²Establishing the validity of the instrument sets through this test is important as this validity is not necessarily guaranteed a priori. For example, if measurement error in c_{it} or y_{it} takes the form of an autocorrelated *MA* process instead of an *iid* process, then some lagged instruments (e.g., for period $t - 1$ in case of an *MA*(1) process) may be invalid and it may be necessary to start with deeper lags (e.g., starting from $t - 2$ in case of an *MA*(1) process). This typically is detrimental to instrument quality.

²³Stock and Yogo (2004), in Table 1, provide the 5% critical values for the null hypothesis that the bias of the IV estimator relative to the bias of the OLS estimator exceeds the threshold of $x\%$ (see the notes to Table 4 for the critical values).

Table 4: Results controlling for measurement error: IV-based mean-group estimates

	Instrument set 1			Instrument set 2		
	Dependent variable $x_{i,t+1}$			Dependent variable $x_{i,t+1}$		
	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$
d_{it}	-0.039** (0.016)	0.004 (0.021)	-0.043 (0.030)	-0.034* (0.018)	0.024 (0.031)	-0.058 (0.037)
$(c_{it} - y_{it})$	0.000 (0.021)	-0.026 (0.021)	0.026*** (0.010)	-0.002 (0.023)	-0.025 (0.021)	0.023** (0.011)
$(c_{it} - y_{it})d_{it}$	0.195*** (0.072)	-0.185** (0.092)	0.380*** (0.082)	0.270** (0.113)	-0.210** (0.107)	0.480*** (0.095)
Cumby-Huizinga AC	2.735 [0.255]	3.105 [0.212]	1.839 [0.399]	3.004 [0.223]	2.443 [0.295]	2.076 [0.354]
Sargan-Hansen OR	10.335 [0.324]	9.573 [0.386]	9.065 [0.431]	5.310 [0.150]	4.470 [0.215]	3.895 [0.273]
Cragg-Donald WI	9.577			15.537		

Notes: Reported are the mean-group results based on IV estimation of eq.(6). Estimation is based on panel data for sixteen countries over the period 1870–2016. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. Instrument set 1 consists of a constant and lags one to four of the regressors d_{it} , $(c_{it} - y_{it})$ and $(c_{it} - y_{it})d_{it}$. Instrument set 2 consists of a constant and lags one to two of the regressors d_{it} , $(c_{it} - y_{it})$ and $(c_{it} - y_{it})d_{it}$. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation tests, testing the null hypothesis of no autocorrelation against the alternative that the autocorrelations of the error term are nonzero at lags greater than zero (with maximum lag equal to two). The Sargan-Hansen OR test reported is the average of the country-specific Sargan-Hansen overidentifying restrictions statistics that test the null hypothesis of the joint validity of the instruments used (see Sargan, 1958; Hansen, 1982). The Cragg-Donald WI test is the average of the country-specific Cragg-Donald weak instrument test statistics (see Cragg and Donald, 1993). Stock and Yogo (2004) in their Table 1 provide the 5% critical values for the null hypothesis that the bias of the IV estimator relative to the bias of the OLS estimator exceeds the threshold of $x\%$. Assuming all three regressors in eq.(6) are measured with noise and are therefore potentially endogenous, these critical values are 10.01 (for $x = 10\%$), 5.90 (for $x = 20\%$) and 4.42 (for $x = 30\%$) for instrument set 1 (which contains twelve instruments excluding the constant) and 7.77 (for $x = 10\%$), 5.35 (for $x = 20\%$) and 4.40 (for $x = 30\%$) for instrument set 2 (which contains six instruments excluding the constant).

Alternative disaster dummy

Our results so far have been based on disaster dummies constructed from the consumption and GDP disaster episodes identified by Barro and Ursúa (2008). More recently, Nakamura et al. (2013) estimate a model of consumption disasters that generates endogenous estimates of the timing and length of disasters. We use the start and end dates of their identified disaster episodes (see Table 2 in Nakamura et al., 2013) to construct an alternative disaster dummy.

Table 5 then presents our predictability results when estimating eqs.(6) and (8) with this alternative dummy variable for d . We report results both without and with a lagged dependent variable included in the equation as, in contrast to the results obtained with our standard disaster dummy that are reported in Table 1, the lagged dependent variable is now significant in all regressions. The reported results—in

particular, those obtained from the equation that includes the lagged dependent variable—confirm our main finding that the predictive power of $c - y$ is higher for both future income and consumption growth rates during macro disasters.

Table 5: Results using an alternative disaster dummy: OLS-based mean-group estimates

	Without lagged dependent variable			With lagged dependent variable		
	Dependent variable $x_{i,t+1}$			Dependent variable $x_{i,t+1}$		
	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$
d_{it}	-0.022** (0.009)	0.017 (0.014)	-0.038** (0.015)	-0.016 (0.010)	0.025* (0.014)	-0.039** (0.018)
$(c_{it} - y_{it})$	-0.026*** (0.009)	-0.054*** (0.018)	0.028 (0.018)	-0.022** (0.011)	-0.048*** (0.017)	0.025** (0.012)
$(c_{it} - y_{it})d_{it}$	0.196*** (0.071)	-0.080 (0.069)	0.276*** (0.037)	0.196*** (0.071)	-0.131** (0.067)	0.311*** (0.041)
x_{it}				0.149*** (0.038)	0.120*** (0.043)	0.110** (0.046)
Cumby-Huizinga AC	3.340 [0.188]	3.329 [0.189]	2.066 [0.356]	2.270 [0.321]	3.207 [0.201]	2.225 [0.329]

Notes: Reported are the mean-group results based on OLS estimation of eq.(6) (results without lagged dependent variable) and eq.(8) (results with lagged dependent variable) using a disaster dummy d_{it} based on disasters identified by Nakamura et al. (2013). Estimation is based on panel data for sixteen countries over the period 1870–2016. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation tests, testing the null hypothesis of no autocorrelation against the alternative that the autocorrelations of the error term are nonzero at lags greater than zero (with maximum lag equal to two).

Disposable income

The baseline results are based on estimations that use real GDP as a proxy for income. Theoretically, using an after-tax measure of income is more appropriate, but historical data on disposable income are not widely available. Piketty and Zucman (2014) provide historical data on national income after taxes that are available for only four countries out of the sixteen considered when using GDP data.²⁴ These countries are France, Germany, the UK and the US.

²⁴The website is <http://piketty.pse.ens.fr/fr/capitalisback>. The data used are in the country excel files, Table 1, columns 9 and 14. From the reported per capita real national income series and the reported series for the ratio of national income after taxes to national income, a series is constructed for per capita real disposable national income (=national income minus taxes plus transfers). Note that, in line with our consumption data (see Section 3.1), we express this series as an index. The data used are available uninterruptedly from 1870 onward. One exception is the UK where the ratio of after-tax national income to national income is only available from 1948 onward. Here, we extrapolate the 1948 value of this ratio to the period 1870–1947. Note further that we update the calculated historical per capita real disposable income series from 2011 to 2016 using data from OECD Economic Outlook.

Table 6: Results using disposable income: OLS-based mean-group estimates

	Disposable income			GDP (for comparison)		
	Dependent variable $x_{i,t+1}$			Dependent variable $x_{i,t+1}$		
	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$	$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$
d_{it}	-0.015 (0.021)	-0.026*** (0.003)	0.010 (0.020)	-0.040* (0.021)	-0.029*** (0.003)	-0.012 (0.023)
$(c_{it} - y_{it})$	-0.013 (0.016)	-0.085*** (0.031)	0.072** (0.036)	-0.016* (0.009)	-0.050*** (0.018)	0.034** (0.015)
$(c_{it} - y_{it})d_{it}$	0.173** (0.080)	0.047 (0.069)	0.125** (0.064)	0.173*** (0.046)	-0.006 (0.065)	0.179*** (0.021)
Cumby-Huizinga AC	4.274 [0.118]	3.501 [0.174]	3.905 [0.142]	3.939 [0.140]	4.308 [0.116]	3.160 [0.206]

Notes: Reported are the mean-group results based on OLS estimation of eq.(6) using log per capita real disposable national income for y_{it} . Estimation is based on panel data for four countries over the period 1870–2016. The results for this sample when using per capita real GDP for y_{it} are added for comparison. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation tests, testing the null hypothesis of no autocorrelation against the alternative that the autocorrelations of the error term are nonzero at lags greater than zero (with maximum lag equal to two).

In Table 6 (columns 2-4), we therefore report the OLS-based mean-group estimates obtained from estimating eq.(6) with y now calculated as the log of per capita real national disposable (after-tax) income. For the purpose of comparison, the table also reports the mean-group estimates obtained from this reduced sample of four countries when using our standard variable for y , namely the log of per capita real GDP (columns 5-7). We note that the results obtained for both measures of y are quite similar. The results for $x_{i,t+1} = \Delta y_{i,t+1}$ and $x_{i,t+1} = (\Delta y_{i,t+1} - \Delta c_{i,t+1})$ confirm our baseline findings, that during macro disasters, the predictive power of $c - y$ for Δy and $\Delta y - \Delta c$ is higher. Contrary to our baseline results, however, we do not find a significantly higher predictive impact of $c - y$ on Δc . Importantly, this result is obtained for *both* measures of income and therefore cannot be attributed to our use of an alternative income measure. Rather, it stems from the low N dimension of the panel used here (i.e., $N = 4$) which can make the mean-group results less stable and driven by outliers.²⁵

3.4 Longer horizons

The intertemporal budget constraint discussed in Section 2 implies that the current log consumption-income ratio may have predictive power, not only for next period's income and consumption growth rates,

²⁵Specifically, the insignificant mean-group estimate for the coefficient of Δc on $(c - y)d$ is driven by the outlier result for Germany for which the estimate for γ_i in this regression is positive rather than negative. To see this for the GDP-based result, refer to Appendix D, which reports the country-specific estimates of β_i and γ_i that underlie our baseline mean-group results.

but also for income and consumption growth rates further into the future. In this section, we therefore investigate how macro disasters affect the predictive power of $c_{it} - y_{it}$ at longer horizons. To this end, we consider our baseline specification at longer horizons; that is, we estimate,

$$x_{i,t+j} = \mu_i^j + \alpha_i^j d_{it} + \beta_i^j (c_{it} - y_{it}) + \gamma_i^j (c_{it} - y_{it}) d_{it} + \epsilon_{i,t+j} \quad (10)$$

with horizon j and where $x_{i,t+j} = \Delta y_{i,t+j}, \Delta c_{i,t+j}, (\Delta y_{i,t+j} - \Delta c_{i,t+j})$.

Table 7: Results at longer horizons: OLS-based mean-group estimates

	Horizon $j = 2$			Horizon $j = 3$		
	Dependent variable $x_{i,t+j}$			Dependent variable $x_{i,t+j}$		
	$\Delta y_{i,t+j}$	$\Delta c_{i,t+j}$	$(\Delta y_{i,t+j} - \Delta c_{i,t+j})$	$\Delta y_{i,t+j}$	$\Delta c_{i,t+j}$	$(\Delta y_{i,t+j} - \Delta c_{i,t+j})$
d_{it}	-0.024*	0.001	-0.025	-0.010	0.002	-0.012
	(0.012)	(0.018)	(0.022)	(0.011)	(0.017)	(0.011)
$(c_{it} - y_{it})$	-0.013	-0.038**	0.025***	-0.019*	-0.032**	0.013*
	(0.012)	(0.015)	(0.009)	(0.010)	(0.015)	(0.008)
$(c_{it} - y_{it})d_{it}$	0.143**	-0.094*	0.237***	0.022	-0.116**	0.138***
	(0.057)	(0.054)	(0.048)	(0.048)	(0.054)	(0.031)
Cumby-Huizinga AC	3.438	3.893	2.483	3.297	4.071	3.044
	[0.329]	[0.273]	[0.478]	[0.509]	[0.396]	[0.550]

Notes: Reported are the mean-group results based on OLS estimation of eq.(10) for horizons $j = 2$ and $j = 3$. Estimation is based on panel data for sixteen countries over the period 1870–2016. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation tests. For $j = 2$, we test the null hypothesis of autocorrelation of either order zero or order one against the alternative that the autocorrelations of the error term are nonzero at lags greater than one (with maximum lag equal to three). For $j = 3$, we test the null hypothesis of autocorrelation of either order zero, order one or order two against the alternative that the autocorrelations of the error term are nonzero at lags greater than two (with maximum lag equal to four).

The OLS-based mean-group estimates obtained from estimating eq.(10) for $j = 2$ and $j = 3$ are reported in Table 7. Compared to the baseline results reported in Table 1, the coefficients on our regressor of interest $(c - y)d$ are generally somewhat smaller (in absolute value). With the exception of the impact of $(c - y)d$ on Δy at horizon $j = 3$, they are all significant. We note that the impact of $(c - y)d$ is significant until $j = 4$ for Δc and until $j = 5$ for $\Delta y - \Delta c$ (results are unreported but available upon request). Hence, while also significant for $j > 1$, the predictive ability of $c - y$ during disasters clearly decreases with the horizon j .

In sum, in line with the validity of the IBC, we also find evidence of the predictive power of the log consumption-income ratio—and of its different impact during disasters—at horizons larger than one.

3.5 What about ordinary recessions?

We now investigate whether our results hold, not only for disasters, but also for more conventional recessions. To this end, we conduct estimations using recession dummies instead of the disaster dummies considered previously. To focus on ordinary recessions, we restrict our sample to the period 1960–2016 with the same $N = 16$ countries considered in the analysis of historical disasters. Over this period, almost no disasters of the type defined by Barro and Ursúa (2008) have occurred, while a large number of ordinary recessions have taken place. We calculate an annual recession dummy d^{rec} from the OECD Composite Leading Indicator (CLI) of activity, which provides monthly data on recession dates—that is, turning points—for each country in our sample.²⁶ The other data used in the estimations are those used in the baseline regressions, albeit taken over a smaller sample period.

In Table 8, we report OLS-based mean-group estimates obtained when estimating eq.(6) with d^{rec} for $x = \Delta y$ and $x = \Delta c$ (we leave out the results for $x = (\Delta y - \Delta c)$ to save space). For these results, however, we cannot reject the null hypothesis of no autocorrelation based on the Cumby-Huizinga test. As such, we also look at the results obtained when estimating eq.(8) where the lagged dependent variable is included as a regressor. By adding this regressor, the autocorrelation issue can be tackled to some extent as can be seen from the improved autocorrelation tests. The reported results suggest that, during ordinary recessions, the predictive ability of $c - y$ is significantly higher for Δy but not for Δc . The impact of $c - y$ on Δy , while in accordance with the results found for Δy in disasters, is quantitatively smaller, however, and less robust. An example of this lack of robustness is given by the CCEMG estimates that we also report in the table. The CCEMG estimator corrects for cross-sectional dependence as detailed above. Based on these CCEMG estimates, we do not find an increase in the predictive ability of $c - y$ during ordinary recessions, neither for Δc nor for Δy .²⁷

Hence, while our previous results show that the consumption-income ratio has more predictive ability for future income and consumption growth rates during disaster episodes, we cannot robustly draw the same conclusion when looking at ordinary recessions, even severe ones like the Great Recession (2007–2009).

²⁶We first calculate a monthly recession dummy per country that is set to one for the months after the peak and up to and including the trough. A quarterly recession dummy for that country then equals one if the monthly dummy equals one during at least two months of the quarter under consideration. An annual recession dummy for that country then equals one if the quarterly dummy equals one during at least two quarters of the year under consideration.

²⁷This is also true when estimating the regressions with d^{rec} using IV (results are unreported but available upon request).

Table 8: Results for ordinary recessions: OLS-and CCE-based mean-group estimates

	Dependent variable $x_{i,t+1}$					
	$\Delta y_{i,t+1}$			$\Delta c_{i,t+1}$		
	OLS		CCE	OLS		CCE
	(1)	(2)		(1)	(2)	
d_{it}^{rec}	-0.015***	-0.011***	-0.009***	-0.009***	-0.004***	-0.007***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)
$(c_{it} - y_{it})$	0.041	0.054	0.005	-0.020	-0.045	-0.089***
	(0.046)	(0.035)	(0.023)	(0.047)	(0.031)	(0.022)
$(c_{it} - y_{it})d_{it}^{rec}$	0.049**	0.062**	0.024	0.011	0.037	0.002
	(0.024)	(0.026)	(0.025)	(0.033)	(0.030)	(0.032)
x_{it}		0.270***			0.421***	
		(0.054)			(0.052)	
Cumby-Huizinga AC	7.939	3.421	4.206	9.427	3.039	4.103
	[0.019]	[0.181]	[0.122]	[0.009]	[0.219]	[0.129]

Notes: Reported are the mean-group results based on either OLS or static CCE estimation of eq.(6) or eq.(8) with either $x_{i,t+1} = \Delta y_{i,t+1}$ or $x_{i,t+1} = \Delta c_{i,t+1}$ and with recession dummy d_{it}^{rec} instead of disaster dummy d_{it} . The recession dummy d_{it}^{rec} is constructed from the OECD Composite Leading Indicator (CLI) of activity. Estimation is based on panel data for sixteen countries over the period 1960–2016. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation tests, testing the null hypothesis of no autocorrelation against the alternative that the autocorrelations of the error term are nonzero at lags greater than zero (with maximum lag equal to two).

3.6 Major historical disasters

We also investigate whether all disaster episodes magnify the predictive impact of the log consumption-income ratio or whether only particular episodes do so. To look at this issue, we investigate the separate impact of the major disaster episodes that occurred during the sample period according to Barro and Ursúa (2008): World War I (WW1), the Spanish flu pandemic of the late 1910s/early 1920s (PAN), the Great Depression (GRD) and World War II (WW2). The estimation details and the obtained results of this exercise are provided in Appendix E. For all major disasters considered, we find that the predictive power of the log consumption-income ratio becomes significantly higher during the occurrence of these major crises. Hence, the reduction in decoupling between consumption and income and the implied reduction in consumption smoothing is not limited to one particular disaster type, instead seemingly characterizing every major crisis type that we consider in our historical dataset.

3.7 The COVID-19 pandemic

We now take a look at the impact of the COVID-19 pandemic, a contemporaneous macroeconomic disaster, on the predictive ability of the log consumption-income ratio. To this end, we use quarterly

data over the period 1995Q1–2021Q4 for twenty industrial economies (i.e., $N = 20$).²⁸ In line with our previous estimations, our specification is given by,

$$x_{i,t+1} = \alpha_i d_t + \beta_i (c_{it} - y_{it}) + \gamma_i (c_{it} - y_{it}) d_t + \epsilon_{i,t+1} \quad (11)$$

with dependent variable $x_{i,t+1}$ and where d_t denotes the COVID-19 dummy that is set to one over the period 2020Q1–2021Q4 for all countries. For c_{it} , we use the log of per capita real private final consumption expenditures, while for y_{it} we use the log of per capita real GDP.²⁹

The OLS-based mean-group results of estimating eq.(11) are presented in Table 9 (column No lag dep. var.). Results are reported only for $x_{i,t+1} = (\Delta y_{i,t+1} - \Delta c_{i,t+1})$ because of space considerations and because the results obtained for $\Delta y_{i,t+1}$ and $\Delta c_{i,t+1}$ separately are considerably less precise. In line with our previous discussion and findings, we observe that this period’s log consumption-income ratio $c - y$ has a positive impact on the next period’s income-consumption differential $\Delta y - \Delta c$, and that this predictive ability is significantly higher during the COVID-19 pandemic. This suggests that during the COVID-19 pandemic there is also less decoupling between consumption and income, which points to a reduction in consumption smoothing. The reported results are robust to adding the lagged dependent variable as a regressor to eq.(11) (column Lag dep. var.), to detrending the predictor variable $c - y$ (column Detrended $c - y$), and to using log per capita real disposable income instead of log per capita real GDP as a proxy for y (column Disp. inc.).^{30,31}

²⁸These are Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, the Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, the UK and the US.

²⁹Real private final consumption expenditures and real GDP are taken from OECD Economic Outlook (No.110) and we calculate per capita measures using quarterly population data from Datastream.

³⁰In contrast to what we find for the historical period 1870–2016 (see Appendix C), panel unit root tests applied to the variable $c - y$ over the period 1995Q1–2021Q4 do not reject that $c - y$ is stochastically trended. To deal with this, we consider $c - y$ in deviation from its stochastic trend $\overline{c - y}$ where the latter is approximated by a twenty-quarter moving average as $\overline{c - y} = \frac{1}{20} \sum_{j=0}^{19} (c_{-j} - y_{-j})$. Our findings are also robust if instead we proxy the stochastic trend using a moving average calculated over either ten or forty quarters.

³¹Data for nominal disposable income of households and non-profit institutions serving households are taken from OECD Economic Outlook (No.110) and are available for seven countries (i.e., Australia, Canada, France, Germany, Japan, the UK and the US). They are put in per capita real terms using the deflator of private final consumption expenditures from OECD Economic Outlook and population data from Datastream.

Table 9: Results for the COVID-19 pandemic: OLS-based mean-group estimates

	Dependent variable $x_{i,t+1} = (\Delta y_{i,t+1} - \Delta c_{i,t+1})$					
	Excluding ordinary recessions				Including ordinary recessions	
	No lag dep. var.	Lag dep. var.	Detrended $c - y$	Disp. inc.	No lag dep.var.	Lag dep. var.
d_t	0.624*** (0.082)	0.641*** (0.082)	0.030** (0.014)	0.175*** (0.023)	0.708*** (0.082)	0.691*** (0.078)
$(c_{it} - y_{it})$	0.044*** (0.009)	0.042*** (0.009)	0.050*** (0.012)	0.161*** (0.022)	0.039*** (0.011)	0.033*** (0.010)
$(c_{it} - y_{it})d_t$	0.903*** (0.095)	0.926*** (0.100)	0.873*** (0.097)	1.031*** (0.063)	1.012*** (0.097)	0.988*** (0.101)
x_{it}		0.014 (0.045)				-0.032 (0.045)
d_{it}^{rec}					0.007 (0.011)	0.010 (0.010)
$(c_{it} - y_{it})d_{it}^{rec}$					0.016 (0.018)	0.022 (0.017)
Cumby- Huizinga AC	4.247 [0.120]	3.144 [0.208]	3.619 [0.164]	1.517 [0.468]	4.689 [0.096]	2.870 [0.238]

Notes: Reported are the mean-group results based on OLS estimation of eqs.(11) and (12). d_t denotes the COVID-19 dummy which equals one over the period 2020Q1–2021Q4. d_{it}^{rec} denotes the recession dummy which is constructed from the OECD Composite Leading Indicator (CLI) of activity. Columns two to five present the results of the estimation of eq.(11). Column 'No lag dep. var.' presents the baseline results of the estimation of eq.(11). In column 'Lag dep. var.', the first lag of the dependent variable is added as a regressor to eq.(11). In column 'Detrended $c - y$ ', the detrended log consumption-income ratio is used for $c - y$ in eq.(11). In column 'Disp. inc.', log per capita real disposable income is used for y in eq.(11) instead of log per capita real GDP. Both final columns present the results of the estimation of eq.(12) where in column 'Lag dep. var.' the first lag of the dependent variable is added as a regressor to eq.(12). Estimation is based on panel data for twenty countries (columns 2, 3 and 4), seven countries (column 5) or nineteen countries (columns 6 and 7) over the period 1995Q1–2020Q4. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation test, testing the null hypothesis of no autocorrelation against the alternative that the autocorrelations of the error term are nonzero at lags greater than zero (with maximum lag equal to two).

As before, we ask ourselves whether the increased predictive ability of the log consumption-income ratio during the COVID-19 pandemic is specific to this disaster episode or whether it occurs also during more conventional recessions that have taken place over the considered sample period. To investigate this, we estimate,

$$(\Delta y_{i,t+1} - \Delta c_{i,t+1}) = \alpha_i d_t + \beta_i (c_{it} - y_{it}) + \gamma_i (c_{it} - y_{it}) d_t + \alpha_i^{rec} d_{it}^{rec} + \gamma_i^{rec} (c_{it} - y_{it}) d_{it}^{rec} + \epsilon_{i,t+1} \quad (12)$$

where, as before, d_t is the common COVID-19 dummy and d_{it}^{rec} denotes the country-specific recession dummy. The latter is calculated from the OECD Composite Leading Indicator (CLI) of activity, which

provides monthly data on recession dates—that is, turning points—for each country in our sample.³² If ordinary recessions also increase the predictive power of the log consumption-income ratio, we should not only find a significantly positive impact of the regressor $(c_{it} - y_{it})d_t$, but also of the regressor $(c_{it} - y_{it})d_{it}^{rec}$. The OLS-based mean-group results of the estimation of eq.(12)—without and also with the inclusion of the lagged dependent variable—are presented in the final two columns of Table 9. In line with the findings for annual historical data reported and discussed in Section 3.5 above, there is no evidence that suggests that conventional recessions have an impact on the predictive ability of the consumption-income ratio, that is, the coefficient on the regressor $(c_{it} - y_{it})d_{it}^{rec}$ is never significantly different from zero. As such, in terms of its impact on the long-run IBC-implied relationship between consumption and income, the COVID-19 pandemic is more akin to historical disaster episodes and has less in common with more typical recessions (including the Great Recession of 2007–2009). This conclusion supports the sentiment of Goldstein et al. (2021) who argue that the COVID-19 crisis is not just “another” large-scale shock, but is instead fundamentally different from previous financial and economic crises, including the Great Recession.

4 Theoretical interpretation of the results

The previous section provides robust evidence that consumption and income are more closely linked (or, less decoupled) during rare macroeconomic disasters, implying that consumption smoothing opportunities are reduced. In this section, we investigate how this evidence relates to consumption theory. To this end, we impose additional structure on the predictive relationship obtained from the IBC by specifying consumer behavior.

4.1 Consumption growth

We consider a savers-spenders set-up where one consumer type is optimizing intertemporally and the other type follows a rule-of-thumb and consumes current income in every period (see e.g., Campbell and Mankiw, 1989; Mankiw, 2000). Mankiw (2000) suggests that rule-of-thumb consumer behavior may stem both from consumers who deviate from rational expectations and/or from consumers who face a binding liquidity constraint. This gives the following expression for total consumption growth,

$$\Delta c_{t+1} = \lambda \Delta y_{t+1} + (1 - \lambda) \Delta c_{t+1}^* \tag{13}$$

³²See footnote 25 for details. Since these data are missing for New Zealand in 2020 and 2021, estimations with the recession dummy—that is, both final columns of Table 9—are based on a panel of nineteen countries instead of twenty.

where λ reflects the fraction of income going to rule-of-thumb consumers (with $0 \leq \lambda < 1$) and Δc_{t+1}^* is the consumption growth rate of intertemporally optimizing consumers. The latter is derived in Appendix A and is given by,

$$\Delta c_{t+1}^* = -\frac{1}{\theta}\delta + \frac{1}{\theta}E_t r_{t+1} + \frac{\sigma}{\theta}E_t \nu_{t+1} + \omega_{t+1} \quad (14)$$

where $\theta > 0$ is the coefficient of relative risk aversion, $\delta > 0$ is the rate of time preference and r_{t+1} is the real rate of return on wealth. The term ω_{t+1} for which $E_t(\omega_{t+1}) = 0$ reflects the part of consumption growth related to the arrival of new information. The component $\frac{1}{\theta}E_t r_{t+1}$ is related to intertemporal substitution in consumption in response to expected changes in the rate of return; that is, a period t expected increase (resp. decrease) in the rate of return of period $t+1$ implies an increase (resp. decrease) in consumption growth from t to $t+1$ as consumption is shifted from t to $t+1$ (resp. from $t+1$ to t). The component $\frac{\sigma}{\theta}E_t \nu_{t+1}$ is the part of consumption growth of the optimizing consumer that reflects a precautionary saving motive (i.e., the precautionary component) (see e.g., Parker and Preston, 2005). As precautionary saving reduces period t consumption and augments period $t+1$ consumption, thereby raising consumption growth from t to $t+1$, we show in Appendix A that this component is positive. The parameter $\sigma > 0$ reflects the strength of the precautionary saving motive.

Our model for consumption growth nests several consumption models considered in the literature. Upon setting $r_{t+1} = \delta$ ($\forall t$), $\lambda = 0$ and $\sigma = 0$, we obtain the log-linear version of the standard permanent income model with log consumption following a random walk (see e.g., Campbell and Mankiw, 1989). We then have $c_{t+1} = c_t + \omega_{t+1}$ with $E_t(\omega_{t+1}) = 0$. In this setting, there is maximal consumption smoothing as consumers expect the same consumption in every period; that is, we have $E_t(c_{t+1}) = c_t$ ($\forall t$). The log-linear permanent income model with intertemporal substitution in consumption in response to return variation is obtained for $\lambda = 0$ and $\sigma = 0$ (see e.g., Hall, 1988). If this model is extended with rule-of-thumb consumers and we therefore only restrict our set-up by imposing $\sigma = 0$, we obtain a standard savers-spenders model (see e.g., Campbell and Mankiw, 1989; Mankiw, 2000). Finally, if the only imposed restriction is $\lambda = 0$, we have the consumption growth rate obtained from a typical buffer stock model of saving (see Carroll, 1992; Parker and Preston, 2005).

4.2 The consumption-income ratio

Taking into account consumer behavior, the log consumption-income ratio can be obtained by substituting eqs.(13) and (14) into the IBC given by eq.(1) to obtain,

$$c_t - y_t = (1 - \lambda) \sum_{j=1}^{\infty} \rho^j \left[E_t(\Delta y_{t+j}) + \frac{1}{\theta}\delta - \frac{1}{\theta}E_t(r_{t+j}) - \frac{\sigma}{\theta}E_t(\nu_{t+j}) \right]. \quad (15)$$

From this equation, we note that, since $0 \leq \lambda < 1$, the consumption-income ratio depends on expected future income changes, on expected future rates of return on wealth and on the expected future precautionary components. Under the standard (log-linearized) permanent income model for which we have $r_{t+1} = \delta$ ($\forall t$), $\lambda = 0$ and $\sigma = 0$, eq.(15) then reduces to $c_t - y_t = \sum_{j=1}^{\infty} \rho^j E_t(\Delta y_{t+j})$, which is the log-linear version of Campbell (1987)'s "saving for a rainy day" expression (i.e., if income is expected to fall, the consumer saves). With respect to the other determinants of $c_t - y_t$, we note that it is negatively affected by expected rates of return $E_t r_{t+j}$ and by the expected precautionary components $E_t \nu_{t+j}$; that is, saving increases when $E_t r_{t+j}$ increases (i.e., intertemporal substitution) and when $E_t \nu_{t+j}$ increases (i.e., precautionary saving).

4.3 Implications for predictability

With respect to our findings of Section 3, as it turns out, both deviations from the standard (log-linearized) permanent income model with time-varying returns discussed above—that is, rule-of-thumb consumption and precautionary saving—can explain our documented changes in the predictive impact of the log consumption-income ratio for income and consumption growth rates during disasters.

First, a reduction in consumption smoothing can occur because of an increase in rule-of-thumb consumer behavior. This is captured by the parameter λ , where an increase in λ implies—all else constant—a more positive predictive impact of $c_t - y_t$ on future income growth rates Δy_{t+j} . This can immediately be observed from eq.(15) above by multiplying both sides of the equation by $\frac{1}{1-\lambda}$ so that future income growth rates can be written as a function of the current log consumption income ratio $c_t - y_t$ times $\frac{1}{1-\lambda}$. An increase in λ , however, cannot explain the observed more negative impact of $c_t - y_t$ on future consumption growth rates Δc_{t+j} . Indeed, a rise in λ , by increasing the positive predictive impact of $c_t - y_t$ on Δy_{t+j} , tends to also lead to a less negative or even positive predictive impact of $c_t - y_t$ on future Δc_{t+j} as, from eq.(13), Δc_{t+1} is driven by Δy_{t+1} .³³

Second, a reduction in consumption smoothing can occur because of a stronger precautionary saving motive of the optimizing consumers. This is captured by the parameter σ where an increase in σ may lead—all else constant—to a more negative predictive impact of $c_t - y_t$ for future consumption growth rates Δc_{t+j} . To see this, suppose initially that $\sigma = 0$ (i.e., there is no precautionary component in consumption growth). In this case, if $c_t - y_t$ has a negative predictive impact for future consumption growth, it must stem from its negative relationship with $E_t r_{t+j}$ (i.e., it is due to intertemporal substitution). If the precautionary component in consumption growth then becomes more important so that $\sigma > 0$, then the

³³For a large λ , for instance, income and consumption growth rates are highly positively correlated so that the positive impact of $c_t - y_t$ on Δy_{t+j} more than likely implies a positive impact of $c_t - y_t$ on Δc_{t+j} .

predictive ability of $c_t - y_t$ for future consumption growth increases—that is, becomes more negative—as $c_t - y_t$ then has predictive power not only for r_{t+j} but also for ν_{t+j} .³⁴

4.4 Empirical evidence

We shed light on the theoretical channels underlying the results of Section 3 by focussing on the predictive relationships implied by eq.(15). As such, we avoid the direct estimation of the specification for consumption growth given by eqs.(13)-(14). Apart from the theoretical objections that can be formulated against attempting to estimate structural parameters such as risk aversion from aggregate data, there are also practical considerations that complicate this estimation. A major issue concerns the use of instruments for the potentially endogenous regressors. The variables Δy_{t+1} and r_{t+1} , for instance, are notoriously hard to instrument, which renders the instrumental variables estimation of a regression for consumption growth largely unreliable.³⁵

4.4.1 Approach

According to eq.(15), the log consumption-income ratio $c_t - y_t$ may predict Δy_{t+j} , r_{t+j} and ν_{t+j} . Evidence of the predictive ability of $c_t - y_t$ for Δy_{t+1} has been provided in Section 3. Given the theory presented in this section, the finding that $c_t - y_t$ has a more positive predictive impact on future income growth Δy_{t+1} during disaster episodes can be attributed to an increase in rule-of-thumb consumption behavior during these episodes. This, in turn, may be the result of liquidity constraints becoming more binding during disasters. In what follows, we present evidence of the predictive ability of $c_t - y_t$ for ν_{t+1} as this channel constitutes our explanation for the finding reported above that, during disasters, $c_t - y_t$ has a more negative predictive impact on future consumption growth Δc_{t+1} . The problem, however, is that the component ν_{t+1} is unobserved. To deal with this, our approach is twofold. First, we look at the predictive impact of $c_t - y_t$ for future returns r_{t+1} in normal times and during disasters. In doing so, we investigate whether we can rule out the alternative explanation for observing a more negative predictive impact of $c_t - y_t$ on consumption growth Δc_{t+1} during disasters, namely that it is due to a more negative predictive impact of $c_t - y_t$ on r_{t+1} . Second, the precautionary saving motive is induced by uncertainty about the future (see e.g., Deaton, 1992; Carroll, 1992; Carroll et al., 2019). Hence, we proxy the precautionary

³⁴It can be shown that, all else constant, the relationship between σ and the predictive ability of $c_t - y_t$ for the next period's consumption growth Δc_{t+1} is in fact non-linear. For lower values of σ , an increase in σ renders the predictive power of $c_t - y_t$ for Δc_{t+1} more negative, while for higher values of σ , a further increase in σ renders the predictive power of $c_t - y_t$ for Δc_{t+1} less negative.

³⁵This is confirmed by the values obtained for the Cragg-Donald weak instrument test calculated when estimating regressions of consumption growth on income growth and returns using our historical dataset. Using a variety of instrumental variables for income growth and returns, we find values for the Cragg-Donald weak instrument test are typically below one (whereas the rule-of-thumb value for this test equals ten). These results are not reported but are available upon request.

component ν_{t+1} using an uncertainty measure. Then, we investigate whether $c_t - y_t$ has predictive power for this proxy and whether this predictive power is higher during disasters.

4.4.2 Data

The estimations are conducted with the historical dataset used in most previous estimations and detailed in Section 3.1. Additionally, for real returns on wealth r_{t+1} , we use the real rate of return on equity. Historical data for the nominal rate of return on equity are reported by Jordà et al. (2019).³⁶ We deflate nominal returns using the inflation rate calculated from the Consumer Price Index (CPI), which is obtained from the Jordà-Schularick-Taylor Macroeconomic Database.³⁷

To proxy the precautionary component ν_{t+1} using an uncertainty measure, there are few possibilities as, over the historical period considered, data are often unavailable, in particular during the disaster periods that we investigate. A viable option is to follow Mody et al. (2012) who, in their paper on precautionary saving during the Great Recession, consider the variance of per capita real GDP growth as an uncertainty measure. To this end, we estimate a first-order GARCH process for per capita real GDP growth $\Delta y_{i,t+1}$ for every country included in our historical dataset. Details on the GDP data used are provided in Section 3.1. From these estimations, we calculate the conditional variance series h_{t+1} of shocks to per capita real GDP growth. Graphs of these series are presented in Appendix F and they clearly show a pattern of increased uncertainty during the delineated historical disaster episodes.

4.4.3 Results

Table 10 presents the results of estimating the predictive impact of the log consumption-income ratio $c_{it} - y_{it}$ on the real rate of return on equity $r_{i,t+1}$ and on the conditional variance $h_{i,t+1}$ of shocks to per capita real GDP growth; that is, we estimate eq.(6) above with $x_{i,t+1} = r_{i,t+1}$ and with $x_{i,t+1} = h_{i,t+1}$. We report both OLS-based and IV-based mean-group estimates, where the latter control for measurement error as discussed in Section 3.3 above. The Sargan-Hansen OR and Cragg-Donald WI test statistics suggest that the instruments used—that is, lags of the regressors—are valid and of good quality. The results reported for the conditional variance $h_{i,t+1}$ include estimates obtained from estimating a specification that includes the lagged dependent variable as a regressor (i.e., the estimation of eq.(8) above with $x_{i,t+1} = h_{i,t+1}$). This is necessary as the conditional variance series are highly persistent so that excluding the lagged dependent variable in these instances implies poor results for the Cumby-Huizinga

³⁶The data can be found at <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/GGDQGJ> where the nominal equity returns have code eq-tr. Details on the data sources are discussed in the online Appendix of Jordà et al. (2019)'s paper.

³⁷The website is <http://www.macroeconomic.net/data>. The data used has code cpi.

autocorrelation test statistic; that is, the null hypothesis of no autocorrelation is strongly rejected.

The results for the returns on equity suggest that $c_{it} - y_{it}$ has a significant negative impact on $r_{i,t+1}$. This finding supports the theory of intertemporal substitution, where high (expected) returns coincide with a low consumption-income ratio or, conversely, with a high saving ratio. This relationship is unaffected by macroeconomic disasters, however, as can be concluded from the positive but insignificant impact of the regressor $(c_{it} - y_{it})d_{it}$ on $r_{i,t+1}$. As such, it seems that the more negative predictive ability of the log consumption-income ratio for consumption growth during disasters that we document in Section 3 cannot be attributed to a more negative predictive impact of the log consumption-income ratio for real returns.

Table 10: Results for the determinants of consumption growth: OLS- and IV-based mean-group estimates

	Dependent variable $x_{i,t+1}$					
	$r_{i,t+1}$		$h_{i,t+1}$			
	OLS	IV	OLS		IV	
			(1)	(2)	(1)	(2)
d_{it}	-0.049 (0.043)	-0.033 (0.054)	0.006* (0.003)	0.000 (0.002)	0.007 (0.004)	0.000 (0.002)
$(c_{it} - y_{it})$	-0.058* (0.035)	-0.078** (0.035)	0.010 (0.008)	0.007 (0.005)	0.014* (0.008)	0.006 (0.005)
$(c_{it} - y_{it})d_{it}$	0.201 (0.178)	0.368 (0.251)	-0.024* (0.015)	-0.029* (0.016)	-0.048* (0.029)	-0.030** (0.015)
x_{it}				0.769*** (0.035)		0.774*** (0.033)
Cumby-Huizinga AC	4.014 [0.134]	3.528 [0.171]	15.559 [0.000]	3.054 [0.217]	14.331 [0.001]	2.475 [0.290]
Sargan-Hansen OR		7.724 [0.562]			11.900 [0.219]	15.472 [0.217]
Cragg-Donald WI		9.080			9.495	7.692

Notes: Reported are the mean-group results based on either OLS or IV estimation of eq.(6) or eq.(8) with either $x_{i,t+1} = r_{i,t+1}$ or $x_{i,t+1} = h_{i,t+1}$. The variable $r_{i,t+1}$ is the real rate of return on equity. The variable $h_{i,t+1}$ is the conditional variance of shocks to per capita real GDP growth $\Delta y_{i,t+1}$ as estimated from a first-order GARCH process. Estimation is based on panel data for sixteen countries over the period 1870–2015 for the results with $x_{i,t+1} = r_{i,t+1}$ and over the period 1870–2016 for the results with $x_{i,t+1} = h_{i,t+1}$. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation tests, testing the null hypothesis of no autocorrelation against the alternative that the autocorrelations of the error term are nonzero at lags greater than zero (with maximum lag equal to two). The Sargan-Hansen OR test reported is the average of the country-specific Sargan-Hansen overidentifying restrictions statistics that test the null hypothesis of the joint validity of the instruments used (see Sargan, 1958; Hansen, 1982). The Cragg-Donald WI test is the average of the country-specific Cragg-Donald weak instrument test statistics (see Cragg and Donald, 1993). For the critical values, we refer to the notes to Table 4 and to Stock and Yogo (2004). The instrument set used both for $x_{i,t+1} = r_{i,t+1}$ and $x_{i,t+1} = h_{i,t+1}$ consists of a constant and lags one to four of the regressors d_{it} , $(c_{it} - y_{it})$ and $(c_{it} - y_{it})d_{it}$.

The results obtained for the conditional variance in the relevant cases where a lagged dependent variable is included as a regressor suggest that, during normal times, there is no link between the log consumption-income ratio $c_{it} - y_{it}$ and our uncertainty measure $h_{i,t+1}$. During macroeconomic disaster episodes, however, a significant negative relationship is uncovered between $c_{it} - y_{it}$ and $h_{i,t+1}$, where high (expected) uncertainty coincides with a low consumption-income ratio or, conversely, with a high saving ratio. While our uncertainty measure is only an (imperfect) proxy for the theoretical precautionary component in aggregate consumption growth discussed in Section 4.1, this result nonetheless suggests that the precautionary saving motive of optimizing consumers is significantly higher during disasters. Not only is there more uncertainty during disaster episodes, *a given degree* of uncertainty also matters more for saving during these times of turmoil. Importantly, this result supports a precautionary saving interpretation of the empirical finding documented in Section 3 that, during disasters, the log consumption-income ratio has a more negative predictive impact on aggregate consumption growth; that is, during disasters, the log consumption-income ratio has a more negative predictive impact on consumption growth because it has a negative predictive impact on *the precautionary component in* consumption growth.

5 Conclusions

This paper uses a large historical dataset (1870–2016) for sixteen industrial economies to investigate whether rare macroeconomic disasters (e.g., wars, pandemics, depressions) affect consumption smoothing opportunities. Through the estimation of predictive panel regressions, we find that the current log consumption-income ratio has more predictive power for future income and consumption growth rates during rare macroeconomic disasters. This result survives a battery of robustness checks and also holds for the ongoing COVID-19 pandemic, though not for more conventional postwar recessions. It implies that the intertemporal budget constraint (IBC) holds more strictly and that consumption and income are significantly less decoupled during disaster episodes. This, in turn, points to a structural reduction in consumption smoothing. Using a savers-spenders type of model, we show that this reduction can be interpreted as stemming from an increase during disasters of rule-of-thumb consumer behavior, which potentially reflects tighter credit constraints as well as from a stronger precautionary saving motive of those consumers who do optimize. Hence, our evidence supports an incomplete markets-based interpretation of the reduction in consumption smoothing observed during rare disasters.

References

- Adema, Y. and Pozzi, L. (2015). Business cycle fluctuations and household saving in OECD countries: a panel data analysis. *European Economic Review*, 79:214–233.
- Alan, S., Crossley, T., and Low, H. (2012). Saving on a rainy day, borrowing for a rainy day. *IFS Working Paper*, W12/11.
- Barro, R. and Ursúa, J. (2008). Macroeconomic crises since 1870. *Brookings Papers on Economic Activity*, 38(1):255–335.
- Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866.
- Barro, R. J. (2009). Rare disasters, asset prices and welfare costs. *American Economic Review*, 99:243–264.
- Barro, R. J. and Ursúa, J. F. (2012). Rare macroeconomic disasters. *Annual Review of Economics*, 4:83–109.
- Campbell, J. (1987). Does saving anticipate declining labor income? An alternative test of the permanent income hypothesis. *Econometrica*, 55(6):1249–1273.
- Campbell, J. and Mankiw, N. (1989). Consumption, income, and interest rates: reinterpreting the time series evidence. *NBER Macroeconomics Annual*, 4:185–216.
- Carroll, C. D. (1992). The buffer-stock theory of saving: some macroeconomic evidence. *Brookings Papers on Economic Activity*, 1992(2):61–156.
- Carroll, C. D., Slacalek, J., and Sommer, M. (2019). Dissecting saving dynamics: measuring wealth, precautionary, and credit effects. *NBER Working Paper 26131*.
- Chudik, A. and Pesaran, M. (2015). Common correlated effects estimation of heterogeneous dynamic panel data models with weakly exogenous regressors. *Journal of Econometrics*, 188:393–420.
- Cochrane, J. (2005). *Asset Pricing*. Princeton University Press.
- Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). The cost of the Covid-19 crisis: lockdowns, macroeconomic expectations, and consumer spending. *CESifo Working Paper series*, 8292.
- Cragg, J. and Donald, S. (1993). Testing identifiability and specification in instrumental variables models. *Econometric Theory*, 9(2):222–240.

- Cumby, R. E. and Huizinga, J. (1992). Testing the autocorrelation structure of disturbances in ordinary least squares and instrumental variables regressions. *Econometrica*, 60(1):185–195.
- Deaton, A. (1992.). *Understanding Consumption*. Oxford University Press: Oxford, UK.
- Dossche, M. and Zlatanos, S. (2020). Covid-19 and the increase in household savings: precautionary or forced? *ECB Economic Bulletin*, 6/2020.
- Engle, R. and Granger, C. (1987). Co-integration and error correction: representation, estimation and testing. *Econometrica*, 55:251–76.
- Farhi, E. and Gabaix, X. (2016). Rare disasters and exchange rates. *Quarterly Journal of Economics*, 131(1):1–52.
- Frey, B. and Kucher, M. (2000). World War II as reflected on capital markets. *Economics Letters*, 69:197–191.
- Gillman, M., Kejak, M., and Pakoš, M. (2015). Learning about rare disasters: implications for consumption and asset prices. *Review of Finance*, 19:1053–1104.
- Goldstein, I., Kojien, R., and Mueller, H. (2021). Covid-19 and its impact on financial markets and the real economy. *Review of Financial Studies*, 34:5135–5148.
- Gourio, F. (2012). Disaster risk and business cycles. *American Economic Review*, 102(6):2734–2766.
- Hall, R. (1988). Intertemporal substitution in consumption. *Journal of Political Economy*, 96:339–57.
- Hamilton, J. (2008). Macroeconomics and ARCH. *NBER Working Papers*, 14151.
- Hansen, L. (1982). Large sample properties of generalized method of moment estimators. *Econometrica*, 50:1029–1054.
- Im, K., Pesaran, M. H., and Shin, Y. (2003). Testing for unit roots in heterogeneous panels. *Journal of Econometrics*, 115:53–74.
- Jordà, Ò., Knoll, K., Kuvshinov, D., Schularick, M., and Taylor, A. (2019). The rate of return on everything, 1870-2015. *Quarterly Journal of Economics*, 134(3):1225–1298.
- Jordà, Ò., Schularick, M., and Taylor, A. (2016). Macrofinancial history and the new business cycle facts. In Eichenbaum, M. and Parker, J., editors, *NBER Macroeconomics Annual 2016*, volume 31. University of Chicago Press, Chicago.

- Jordà, Ò., Singh, S., and Taylor, A. (2020). Longer-run economic consequences of pandemics. *NBER Working Papers*, 26934.
- Levine, R., Lin, C., Tai, M., and Xie, W. (2021). How did depositors respond to Covid-19? *Review of Financial Studies*, 34:5438–5473.
- Ludvigson, S., Ma, S., and Ng, S. (2020). Covid-19 and the macroeconomic effects of costly disasters. *NBER Working Papers*, 26987.
- Mankiw, N. (2000). The savers-spenders theory of fiscal policy. *American Economic Review: Papers and Proceedings*, 90(2):120–125.
- Mody, A., Ohnsorge, F., and Sandri, D. (2012). Precautionary savings in the Great Recession. *IMF Economic Review*, 60(1):114–138.
- Nakamura, E., Sergeyev, D., and Steinsson, J. (2017). Growth-rate and uncertainty shocks in consumption: cross-country evidence. *American Economic Journal: Macroeconomics*, 9(1):1–39.
- Nakamura, E., Steinsson, J., Barro, R., and Ursúa, J. (2013). Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics*, 5(3):35–74.
- Newey, W. and West, K. (1987). A simple positive semi-definite heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55:703–708.
- Parker, J. and Preston, B. (2005). Precautionary savings and consumption fluctuations. *American Economic Review*, 95(4):1119–1143.
- Pesaran, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4):967–1012.
- Pesaran, M. H. (2007). A simple panel unit root test in the presence of cross-section dependence. *Journal of Applied Econometrics*, 22(2):265–312.
- Pesaran, M. H. and Smith, R. (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics*, 68(1):79–113.
- Pesaran, M. H., Smith, R., and Im, K.-S. (1996). Dynamic linear models for heterogeneous panels. In Matyas, L. and Sevestre, P., editors, *The econometrics of panel data*, pages 145–195. Kluwer Academic Publishers, Dordrecht, second edition.
- Piketty, T. and Zucman, G. (2014). Capital is back: wealth-income ratios in rich countries 1700-2010. *Quarterly Journal of Economics*, 129(3):1255–1310.

- Richardson, G. (2013). Banking panics of 1930-31. *Federal Reserve History Essays*.
- Rietz, T. A. (1988). The equity risk premium: a solution. *Journal of Monetary Economics*, 22(1):117–131.
- Sargan, J. (1958). The estimation of economic relationships using instrumental variables. *Econometrica*, 26:393–415.
- Sarno, L. and Schmeling, M. (2014). Which fundamentals drive exchange rates? A cross-sectional perspective. *Journal of Money, Credit and Banking*, 46(2-3):267–292.
- Silber, W. (2005). What happened to liquidity when World War I shut the NYSE? *Journal of Financial Economics*, 78:685–701.
- Sommer, M. (2007). Habit formation and aggregate consumption dynamics. *The B.E. Journal of Macroeconomics - Advances*, 7(1).
- Stock, J. H. and Yogo, M. (2004). Testing for weak instruments in linear IV regression. In Andrews, D. and Stock, J., editors, *Identification and inference in econometric models: essays in honor of Thomas J. Rothenberg*. Cambridge University Press, Cambridge, UK.
- Vandenbroucke, G. (2021). Personal saving during the Covid-19 recession. *Federal Reserve Bank of St. Louis Economic Synopses*, 2/2021.

Appendices

Appendix A Derivation of eqs.(1) and (14)

A.1 Derivation of eq.(1)

This appendix describes the steps in the derivation of eq.(1) in the main text. For details, we refer to Campbell and Mankiw (1989). When total wealth is tradeable, the period-by-period budget constraint of a consumer can be written as,

$$W_{t+1} = R_{t+1}(W_t - C_t) \quad (\text{A-1})$$

where W_t is real total wealth, C_t is real consumption and R_t is the gross real return on total wealth. Dividing both sides by W_t , we can write $\frac{W_{t+1}}{W_t} = R_{t+1} \left(1 - \frac{C_t}{W_t}\right)$. After taking logs, this gives

$$\Delta w_{t+1} = r_{t+1} + \ln(1 - \exp(c_t - w_t)) \quad (\text{A-2})$$

with $w_t = \ln W_t$, $r_t = \ln R_t$ and $c_t = \ln C_t$. We linearize the term $\ln(1 - \exp(c_t - w_t))$ by taking a first-order Taylor approximation which gives,

$$\ln(1 - \exp(c_t - w_t)) \approx -\frac{C}{W-C}(c_t - w_t) = \left(1 - \frac{1}{\rho}\right)(c_t - w_t) \quad (\text{A-3})$$

where we ignore the linearization constant and where W and C are the average or steady state values of W_t and C_t .¹ The second step replaces $-\frac{C}{W-C}$ by $1 - \frac{1}{\rho}$ with $\rho \equiv 1 - \frac{C}{W}$ where $0 < \rho < 1$. Substituting eq.(A-3) into eq.(A-2), we obtain $\Delta w_{t+1} = r_{t+1} + \left(1 - \frac{1}{\rho}\right)(c_t - w_t)$. Note that we can write Δw_{t+1} as $\Delta w_{t+1} = \Delta c_{t+1} + (c_t - w_t) - (c_{t+1} - w_{t+1})$. Upon combining these results and rearranging terms, we obtain,

$$c_t - w_t = \rho(r_{t+1} - \Delta c_{t+1}) + \rho(c_{t+1} - w_{t+1}). \quad (\text{A-4})$$

Solving eq.(A-4) forward ad infinitum, imposing the transversality condition $\rho^\infty(c_{t+\infty} - w_{t+\infty}) = 0$ and taking expectations at period t then gives,

$$c_t - w_t = \sum_{j=1}^{\infty} \rho^j E_t(r_{t+j} - \Delta c_{t+j}) \quad (\text{A-5})$$

with E_t the expectations operator conditional on period t information.

Following Campbell and Mankiw (1989), we derive an income-based budget constraint by assuming total wealth W_t consists of N_t shares with ex-dividend price given by P_t and where Y_t is real income (i.e., the real dividend) obtained from total wealth. As such, we have $W_t = N_t(P_t + Y_t)$ where $P_t + Y_t$ is the

¹Note that the linearization occurs around the point $c_t - w_t = c - w$ with $c - w = \ln\left(\frac{C}{W}\right)$.

cum-dividend price of a share. The gross real return on total wealth is given by $R_{t+1} = \frac{P_{t+1} + Y_{t+1}}{P_t}$. By combining these results and rearranging terms, we obtain,

$$W_{t+1}^* = R_{t+1} (W_t^* - Y_t) \quad (\text{A-6})$$

with $W_t^* \equiv \frac{W_t}{N_t}$. Eq.(A-6) is in the same form as eq.(A-1) so the same steps (linearization, defining the discount factor, forward solving) can be applied to obtain,

$$y_t - w_t^* = \sum_{j=1}^{\infty} \kappa^j E_t (r_{t+j} - \Delta y_{t+j}) \quad (\text{A-7})$$

where $w_t^* = \ln W_t^*$ and $y_t = \ln Y_t$. The discount factor κ is given by $\kappa \equiv 1 - \frac{Y}{W^*}$ where $0 < \kappa < 1$.

We then combine eqs.(A-5) and (A-7) where, after imposing the normalization $N_t = 1$ or $\ln N_t = 0$, we obtain,

$$c_t - y_t = \sum_{j=1}^{\infty} [\kappa^j E_t (\Delta y_{t+j} - r_{t+j}) - \rho^j E_t (\Delta c_{t+j} - r_{t+j})]. \quad (\text{A-8})$$

We note that the link between $c_t - y_t$ and expected future returns on wealth r_{t+j} is ambiguous and not substantial if, as can be expected, the discount factor for future income growth rates κ is close to that of future consumption growth rates ρ . Hence, we follow Campbell and Mankiw (1989), and set $\rho = \kappa$ to obtain eq.(1) in the main text.

A.2 Derivation of eq.(14)

This appendix describes the steps in the derivation of eq.(14) in the main text. Consider the following first-order condition for a utility-maximizing consumer who faces uncertainty about future labor income and returns, that is,

$$E_t \left(\frac{(1 + r_{t+1}) U'(C_{t+1}^*)}{(1 + \delta) U'(C_t^*)} \right) = 1 \quad (\text{A-9})$$

where r_t denotes the real return on wealth and $U(C_t^*)$ denotes utility as a function of the level of real consumption of the optimizing consumer C_t^* and, where δ is the rate of time preference. This equation can also be written as,

$$\left(\frac{(1 + r_{t+1}) U'(C_{t+1}^*)}{(1 + \delta) U'(C_t^*)} \right) = 1 + \chi_{t+1} \quad (\text{A-10})$$

where χ_{t+1} is an expectation error uncorrelated with period t information (i.e., we have $E_t \chi_{t+1} = 0$). Using the isoelastic utility function $U(C^*) = \frac{C^{*1-\theta}}{1-\theta}$ with coefficient of relative risk aversion $\theta > 0$, we can rewrite eq.(A-10) as,

$$\left(\frac{(1 + r_{t+1}) C_{t+1}^{*-\theta}}{(1 + \delta) C_t^{*-\theta}} \right) = 1 + \chi_{t+1}. \quad (\text{A-11})$$

After taking logs of both sides of this expression and solving for the growth rate in consumption Δc_{t+1}^* , we obtain,

$$\Delta c_{t+1}^* = -\frac{1}{\theta}\delta + \frac{1}{\theta}r_{t+1} + \frac{1}{\theta}v_{t+1} \quad (\text{A-12})$$

where $v_{t+1} \equiv -\ln(1 + \chi_{t+1})$ and where we use the approximation $\ln(1 + x) \approx x$ for δ and r . The variables r_{t+1} and v_{t+1} can be decomposed into the expected parts $E_t r_{t+1}$ and $E_t v_{t+1}$ and the unexpected parts $(r_{t+1} - E_t r_{t+1})$ and $(v_{t+1} - E_t v_{t+1})$ to obtain,

$$\Delta c_{t+1}^* = -\frac{1}{\theta}\delta + \frac{1}{\theta}E_t r_{t+1} + \frac{1}{\theta}E_t v_{t+1} + \omega_{t+1} \quad (\text{A-13})$$

where $\omega_{t+1} \equiv \frac{1}{\theta} [(r_{t+1} - E_t r_{t+1}) + (v_{t+1} - E_t v_{t+1})]$ with $E_t(\omega_{t+1}) = 0$ and $\frac{1}{\theta}E_t v_{t+1}$ is the part of consumption growth related to the precautionary saving motive of the optimizing consumer (i.e., the precautionary component) (see e.g., Parker and Preston, 2005). This component is positive; that is, we have $E_t v_{t+1} > 0$.² For expositional purposes, we rescale the precautionary component by dividing it by its standard deviation σ . As such, we have $E_t v_{t+1} = \sigma E_t \nu_{t+1}$ where $E_t \nu_{t+1}$ is the standardized precautionary component with unit variance and where $\sigma > 0$ captures the strength of the precautionary saving motive. We substitute this expression for $E_t v_{t+1}$ into eq.(A-13) to obtain eq.(14) in the main text.

Appendix B Historical disaster episodes and dummies

Table B-1 presents the disaster periods used in the construction of the disaster dummies. The reported periods are obtained by combining the consumption and GDP disasters reported in Tables 6 and 8 in Barro and Ursúa (2008). The grouping of consumption and GDP disasters according to principal world economic crises (e.g., World War I, Spanish flu pandemic, Great Depression, World War II) is based on Tables 7 and 9 in Barro and Ursúa (2008).^{3,4,5}

²This can be shown by noting that $\ln(E_t(1 + \chi_{t+1})) = \ln(1) = 0$ (this follows from $E_t(\chi_{t+1}) = 0$). For the concave log function, we have that $\ln(E_t(\cdot)) > E_t(\ln(\cdot))$ so that $E_t(\ln(1 + \chi_{t+1})) < 0$ and $-E_t(\ln(1 + \chi_{t+1})) = E(v_{t+1}) > 0$.

³To illustrate, the UK experienced a consumption disaster over the period 1915–18 attributed to World War I and a GDP disaster over the period 1918–21 attributed to the Spanish flu pandemic. Hence, the overall disaster period is 1915–21 and the general dummy d_{it} for the UK takes on the value of one during this period. Additionally, the episode-specific dummies d_{it}^{WW1} and d_{it}^{PAN} take on the value of one during the periods 1915–18, respectively 1918–21.

⁴We slightly deviate from the grouping considered in Barro and Ursúa (2008) by allocating a number of their post-World War II disaster episodes, occurring in the immediate aftermath of World War II, to our World War II category. This is the case for Denmark (the 1946–48 consumption disaster), Spain (the 1946–49 consumption disaster, the UK (the 1943–47 output disaster) and the US (the 1944–47 output disaster). This minor change has a minimal impact on the estimates and no impact on the conclusions of the paper.

⁵The Spanish flu pandemic is based on the 1920s grouping of Barro and Ursúa (2008) where we include an episode if the first year of the GDP or consumption disaster is either 1918 or 1919. Hence, some episodes from Barro and Ursúa (2008)'s 1920s grouping are unrelated to the pandemic and are therefore not included in our Spanish flu pandemic group. Examples are Germany (1922–23) and Portugal (late 1920s).

Table B-1: Disaster periods used in the construction of disaster dummies

	Episodes						Episodes				
	All	WW1	PAN	GRD	WW2		All	WW1	PAN	GRD	WW2
Australia	1889-95 1910-18 1926-32 1938-46	1910-18		1926-32	1938-46	Netherlands	1889-93 1912-18 1929-34 1939-44	1913-18		1929-34	1939-44
Belgium	1913-18 1930-34 1937-43	1913-18		1930-34	1937-43	Norway	1916-21 1939-44	1916-18	1919-21		1939-44
Denmark	1914-21 1939-41 1946-48	1914-18	1919-21		1939-41 1946-48	Portugal	1913-19 1927-28 1934-36 1939-42 1974-76	1913-19			1939-42
Finland	1876-81 1913-15 1913-18 1928-32 1938-44 1989-93	1913-18		1928-32	1938-44	Spain	1892-96 1913-15 1929-33 1935-38 1940-49	1913-15		1929-33	1940-49
France	1870-71 1874-79 1882-86 1912-18 1929-35 1938-44	1912-18		1929-35	1938-44	Sweden	1913-18 1920-21 1939-45	1913-18			1939-45
Germany	1912-19 1922-23 1928-32 1939-46	1912-19		1928-32	1939-46	Switzerland	1870-72 1875-79 1881-83 1885-88 1912-18 1939-45	1912-18			1939-45
Italy	1918-20 1939-45		1918-20		1939-45	UK	1915-21 1938-47	1915-18	1918-21		1938-47
Japan	1937-45				1937-45	US	1906-08 1913-14 1917-21 1929-33 1944-47		1917-21	1929-33	1944-47

Notes: The periods in the table correspond to periods reported by Barro and Ursúa (2008) as either GDP disaster episodes, consumption disaster episodes or both. The grouping of episodes according to principal world economic crises in columns WW1 (World War I), PAN (Spanish flu pandemic), GRD (Great Depression) and WW2 (World War II) follows the grouping reported by Barro and Ursúa (2008).

The episodes in column All are used to construct the general dummy d_{it} , which is equal to one over the reported periods in the column. The episodes in columns WW1 (World War I), PAN (Spanish Flu pandemic), GRD (Great Depression) and WW2 (World War II) are used to construct the episode-specific dummies d_{it}^j with $j = WW1, PAN, GRD, WW2$, which are equal to one over the reported periods in the respective columns. The episode-specific dummies are used in the estimations discussed in Section 3.6.

Appendix C Panel unit root test consumption-income ratio

The table below reports panel unit root tests applied to the log consumption-income ratio $c_{it} - y_{it}$ constructed using the historical panel data discussed in Section 3.1. Reported are the Im et al. (2003) heterogeneous panel unit root test that does not control for cross-sectional dependence in the data (the IPS statistic) and the Pesaran (2007) heterogeneous panel unit root test that does control for cross-sectional dependence. We report both the standard CIPS statistic and the truncated CIPS* statistic (see Pesaran, 2007, for details). The statistics are reported both for the case without and with a deterministic linear time trend included in the underlying country-specific augmented Dickey-Fuller regressions. We find that in all cases the null hypothesis of a unit root in all countries is strongly rejected in favor of the alternative hypothesis of no unit root in at least one country. Besides the reported p-values (in square brackets), the table also reports the number of countries for which the null hypothesis of a unit root is rejected (in rounded brackets). For these countries, $c_{it} - y_{it}$ is not stochastically trended but rather is stationary, possibly around a deterministic trend. For the CIPS (or CIPS*) test with deterministic linear time trend, the ten countries for which the null hypothesis of a unit root is rejected are Australia, Belgium, France, Italy, the Netherlands, Spain, Sweden, Switzerland, the UK and the US. For a majority of countries and for the overall panel, we therefore reject the presence of a unit root in the log consumption-income ratio.

Table C-1: Heterogeneous panel unit root tests applied to the log consumption-income ratio $c_{it} - y_{it}$

	Panel unit root test		
	IPS	CIPS	CIPS*
Without linear time trend	-2.791	-3.104	-3.093
	[< 0.010]	[< 0.010]	[< 0.010]
	(5)	(9)	(9)
With linear time trend	-6.819	-3.560	-3.554
	[< 0.010]	[< 0.010]	[< 0.010]
	(7)	(10)	(10)

Notes: Estimation is based on panel data for the log consumption income ratio $c_{it} - y_{it}$ for sixteen countries over the period 1870–2016. Reported are the Im et al. (2003) heterogeneous panel unit root test that does not control for cross-sectional dependence (IPS statistic) and the Pesaran (2007) heterogeneous panel unit root tests that do control for cross-sectional dependence (the CIPS statistic and the truncated version, CIPS*). Test statistics are reported both for the case without and with a deterministic linear time trend included in the underlying country-specific augmented Dickey-Fuller regressions. The number of lags included in these regressions is based on the Schwarz information criterion. P-values for testing the null hypothesis of a unit root are between square brackets. Between rounded brackets is the number of countries (out of sixteen) for which the null hypothesis of a unit root is rejected at the 10% level of significance based on the conducted country-specific augmented Dickey-Fuller tests.

Appendix D Per country baseline estimates

The following table reports the per country OLS estimates of the coefficients β_i and γ_i obtained from estimating the baseline specification eq.(6). These estimates are used in the calculation of the mean-group estimates reported in Table 1 in the text. Also reported, between brackets, are heteroskedasticity- and autocorrelation-consistent standard errors (see Newey and West, 1987).

Table D-1: Per country OLS estimates of β_i and γ_i in the baseline specification eq.(6)

Country	Regressor	Dependent variable			Country	Regressor	Dependent variable		
		$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$			$\Delta y_{i,t+1}$	$\Delta c_{i,t+1}$	$(\Delta y_{i,t+1} - \Delta c_{i,t+1})$
Australia	$(c_{it} - y_{it})$	0.001	-0.023	0.024	Netherlands	$(c_{it} - y_{it})$	0.187	0.092	0.095
		(0.015)	(0.018)	(0.015)			(0.173)	(0.095)	(0.083)
	$(c_{it} - y_{it})d_{it}$	-0.053	-0.100	0.047		$(c_{it} - y_{it})d_{it}$	0.129	-0.339	0.467
		(0.063)	(0.075)	(0.095)			(0.213)	(0.345)	(0.239)
Belgium	$(c_{it} - y_{it})$	0.280	0.072	0.208	Norway	$(c_{it} - y_{it})$	-0.015	-0.020	0.005
		(0.109)	(0.043)	(0.106)			(0.010)	(0.010)	(0.007)
	$(c_{it} - y_{it})d_{it}$	-0.148	-0.080	-0.068		$(c_{it} - y_{it})d_{it}$	0.156	-0.042	0.199
		(0.171)	(0.203)	(0.178)			(0.181)	(0.167)	(0.070)
Denmark	$(c_{it} - y_{it})$	0.023	0.020	0.004	Portugal	$(c_{it} - y_{it})$	-0.179	-0.259	0.080
		(0.012)	(0.018)	(0.011)			(0.064)	(0.049)	(0.042)
	$(c_{it} - y_{it})d_{it}$	0.231	-0.142	0.374		$(c_{it} - y_{it})d_{it}$	0.157	0.240	-0.082
		(0.113)	(0.227)	(0.159)			(0.164)	(0.134)	(0.084)
Finland	$(c_{it} - y_{it})$	-0.035	-0.090	0.055	Spain	$(c_{it} - y_{it})$	0.001	-0.021	0.022
		(0.036)	(0.033)	(0.023)			(0.037)	(0.044)	(0.014)
	$(c_{it} - y_{it})d_{it}$	-0.081	-0.083	0.002		$(c_{it} - y_{it})d_{it}$	-0.214	-0.592	0.379
		(0.097)	(0.098)	(0.112)			(0.089)	(0.177)	(0.110)
France	$(c_{it} - y_{it})$	0.004	-0.038	0.042	Sweden	$(c_{it} - y_{it})$	0.021	0.017	0.004
		(0.023)	(0.022)	(0.020)			(0.015)	(0.013)	(0.011)
	$(c_{it} - y_{it})d_{it}$	0.158	-0.033	0.191		$(c_{it} - y_{it})d_{it}$	-0.129	-0.335	0.206
		(0.063)	(0.103)	(0.132)			(0.219)	(0.242)	(0.093)
Germany	$(c_{it} - y_{it})$	-0.018	-0.085	0.066	Switzerland	$(c_{it} - y_{it})$	0.086	-0.020	0.106
		(0.065)	(0.054)	(0.060)			(0.045)	(0.052)	(0.045)
	$(c_{it} - y_{it})d_{it}$	0.307	0.181	0.126		$(c_{it} - y_{it})d_{it}$	-0.092	-0.969	0.877
		(0.238)	(0.061)	(0.209)			(0.175)	(0.214)	(0.155)
Italy	$(c_{it} - y_{it})$	-0.046	-0.058	0.011	UK	$(c_{it} - y_{it})$	-0.041	-0.074	0.033
		(0.020)	(0.021)	(0.015)			(0.024)	(0.021)	(0.018)
	$(c_{it} - y_{it})d_{it}$	0.601	0.314	0.287		$(c_{it} - y_{it})d_{it}$	0.121	-0.051	0.172
		(0.290)	(0.166)	(0.186)			(0.057)	(0.079)	(0.117)
Japan	$(c_{it} - y_{it})$	-0.040	-0.070	0.030	US	$(c_{it} - y_{it})$	-0.010	-0.005	-0.005
		(0.018)	(0.018)	(0.013)			(0.023)	(0.017)	(0.020)
	$(c_{it} - y_{it})d_{it}$	0.533	-0.021	0.554		$(c_{it} - y_{it})d_{it}$	0.105	-0.122	0.227
		(0.100)	(0.219)	(0.297)			(0.049)	(0.047)	(0.058)

Notes: Reported estimates are for β_i and γ_i in equation (6). Heteroskedasticity- and autocorrelation-robust Newey-West standard errors are in parentheses (see Newey and West, 1987). The OLS estimates reported are used to calculate the baseline mean-group estimates reported in Table 1.

Appendix E Results for major historical disaster episodes

In this appendix, we investigate the separate impact of the major disaster episodes that occurred during the sample period according to Barro and Ursúa (2008): World War I (WW1), the Spanish flu pandemic of the late 1910s/early 1920s (PAN), the Great Depression (GRD) and World War II (WW2). To this end, we estimate predictive regression equations of the following form,

$$x_{i,t+1} = \alpha_i d_{it}^j + \alpha_i^{-j} d_{it}^{-j} + \beta_i (c_{it} - y_{it}) + \gamma_i (c_{it} - y_{it}) d_{it}^j + \gamma_i^{-j} (c_{it} - y_{it}) d_{it}^{-j} + \epsilon_{i,t+1} \quad (\text{E-1})$$

where, as before, we have $x_{i,t+1} = \Delta y_{i,t+1}$, $x_{i,t+1} = \Delta c_{i,t+1}$ or $x_{i,t+1} = (\Delta y_{i,t+1} - \Delta c_{i,t+1})$. We estimate the equation per major disaster episode j (with $j = WW1, PAN, GRD, WW2$) while controlling for

all other disasters. To achieve this, we include the specific disaster dummy variable d^j that equals one during disaster period j , but also a dummy variable d^{-j} that takes on the value of one when disasters other than j occur. We note that the dummy d^{-j} equals $d - d^j$, where d is the disaster dummy used in the main text. Both dummies d^j and d^{-j} enter the equation separately as well as interacted with the log consumption-income ratio. As not all specific disasters occur in all sixteen countries of our sample, the estimations are conducted with a different number of countries for each particular disaster episode j . Refer to Appendix B for an overview of the exact dates of the major disaster episodes in each country. In particular, estimation is based on panel data for thirteen countries when $j = WW1$, for five countries when $j = PAN$, for eight countries when $j = GRD$ and for fifteen countries when $j = WW2$.⁶

In Table E-1, we report mean-group estimates obtained from estimating eq.(E-1) for every major disaster episode j . To control for measurement error, we report not only OLS-based but also IV-based estimates. We refer to Section 3.3 in the main text for details. Results are reported only for $x_{i,t+1} = (\Delta y_{i,t+1} - \Delta c_{i,t+1})$ because the results obtained for $\Delta y_{i,t+1}$ and $\Delta c_{i,t+1}$ separately are considerably less precise.⁷ From looking at the results in the table, we note that the largest disasters also tend to have the largest impact on the predictive ability of the log consumption-income ratio; that is, the estimates on the regressor $(c - y)d^j$ (with d^j the dummy for the major disaster episode j under scrutiny) are generally larger in magnitude than those on the regressor $(c - y)d^{-j}$ (with d^{-j} the dummy for the other major disasters but also all the minor ones). Furthermore, we find that for all major disasters considered (with the exception of $j = PAN$ in the IV case), the predictive power of the log consumption-income ratio becomes significantly higher during the occurrence of these major crises. Hence, the reduction in decoupling between consumption and income and the implied reduction in consumption smoothing is not limited to one particular disaster type, but seemingly characterizes every major crisis type that we consider in our historical dataset. Finally, we acknowledge that IV estimation does not necessarily improve on OLS estimation here. While a priori it can be expected that the IV-based estimates control for measurement error as detailed in Section 3.3 in the main text, we find, based on the reported Cragg-Donald statistics, that the instruments used in the estimations reported in Table E-1 are not very strong.

⁶We note that since estimations occur at the country level, a country can only be included in the panel estimation if both dummies d^j and d^{-j} are defined for that country (i.e., if both dummies take on the value of one at least once over the sample period for that country). For example, even though for $j = WW2$ the dummy variable d^j is defined for all sixteen countries, we cannot include Japan in the sample as the dummy d^{-j} is not defined for Japan (i.e., the only disaster identified by Barro and Ursúa (2008) for Japan is $WW2$). Hence, for $j = WW2$, we have $N = 15$ instead of $N = 16$. If we do not include the dummy d^{-j} in the estimations, we can add Japan to the sample when $j = WW2$ and we find that the results with respect to the impact of $WW2$ on the predictive impact of $c - y$ are very similar to those reported in Table E-1. These results are not reported, but are available upon request.

⁷These results are not reported but are available from the authors upon request.

Table E-1: Results for major disaster episodes: OLS- and IV-based mean-group estimates


	Dependent variable $x_{i,t+1} = (\Delta y_{i,t+1} - \Delta c_{i,t+1})$							
	OLS				IV			
	Disaster episode j				Disaster episode j			
	WW1	PAN	GRD	WW2	WW1	PAN	GRD	WW2
d_{it}^j	-0.178 (0.123)	-0.137*** (0.041)	-0.146*** (0.056)	-0.049 (0.035)	-0.637 (0.579)	0.016 (0.208)	-0.241* (0.140)	-0.048 (0.039)
d_{it}^{-j}	-0.056*** (0.019)	-0.029 (0.031)	-0.019 (0.018)	-0.093*** (0.025)	-0.056*** (0.021)	-0.006 (0.055)	-0.046 (0.029)	-0.106*** (0.026)
$(c_{it} - y_{it})$	0.057*** (0.016)	0.010 (0.006)	0.064*** (0.023)	0.050*** (0.014)	0.046*** (0.017)	0.109 (0.096)	0.020 (0.025)	0.040** (0.016)
$(c_{it} - y_{it})d_{it}^j$	0.718*** (0.210)	0.615*** (0.089)	0.529*** (0.129)	0.487*** (0.057)	1.470* (0.793)	0.192 (0.558)	0.848*** (0.294)	0.338*** (0.098)
$(c_{it} - y_{it})d_{it}^{-j}$	0.245*** (0.077)	0.268*** (0.080)	0.187*** (0.069)	0.324*** (0.087)	0.330*** (0.123)	0.044 (0.085)	0.362*** (0.096)	0.342*** (0.094)
Cumby-Huizinga AC	1.735 [0.420]	3.213 [0.201]	1.507 [0.471]	2.543 [0.280]	1.910 [0.385]	3.588 [0.166]	1.533 [0.465]	2.196 [0.334]
Sargan-Hansen OR					13.325 [0.577]	18.253 [0.250]	13.859 [0.536]	14.344 [0.500]
Cragg-Donald WI					3.963	0.588	2.643	3.693

Notes: Reported are the mean-group results based on OLS and IV estimation of eq.(E-1). The dummy variable d^j (with $j = WW1, PAN, GRD, WW2$) equals one during the considered major disaster episode (World War I, Spanish flu pandemic, Great Depression, World War II). We refer to Appendix B for details on the exact dates of these disasters. The dummy d^{-j} takes on the value of one when disasters other than j occur (i.e., it equals $d - d^j$ where d is the general disaster dummy used in previous sections). Estimation is based on panel data for thirteen countries (WW1), five countries (PAN), eight countries (GRD) or fifteen countries (WW2) over the period 1870–2016. Standard errors are in parentheses, p -values are in square brackets. *, **, *** indicate significance at the 10%, 5% and 1% level respectively. The Cumby-Huizinga test shows the average of the individual countries' Cumby and Huizinga (1992) autocorrelation test, testing the null hypothesis of no autocorrelation against the alternative that the autocorrelations of the error term are nonzero at lags greater than zero (with maximum lag equal to two). The Sargan-Hansen OR test reported is the average of the country-specific Sargan-Hansen overidentifying restrictions statistics that test the null hypothesis of the joint validity of the instruments used (see Sargan, 1958; Hansen, 1982). The Cragg-Donald WI test is the average of the country-specific Cragg-Donald weak instrument test statistics (see Cragg and Donald, 1993). For the critical values, we refer to the notes to Table 4 and to Stock and Yogo (2004). The instrument set used for IV estimation consists of a constant and lags one to four of the regressors of eq.(E-1).

Appendix F Per country uncertainty measures

The following figure presents the conditional variance series $h_{i,t+1}$ of shocks to GDP growth for all sixteen countries in our historical dataset. These are obtained from the per country estimation of a first-order GARCH process for per capita real GDP growth. These conditional variance series capture uncertainty and are used as proxies for the precautionary component in aggregate consumption growth as detailed in Section 4. It is clear from the figure that the delineated historical disaster episodes are characterized by substantially higher uncertainty.

Figure F-1: The conditional variance of shocks to per capita real GDP growth



Notes: The blue line denotes the conditional variance $h_{i,t+1}$ of shocks to per capita real GDP growth. Shaded areas correspond to disaster episodes as identified by Barro and Ursúa (2008). Refer to Sections 3.1 and 4.4.2 for more details on the data used in this figure.