Downside-Risk Sharing*

Federico Lenzi† Luca Rossi‡

Abstract

The empirical risk sharing literature has traditionally relied on models that implicitly estimate risk sharing using both tails of the income growth distribution. Drawing from Markowitz (1959), we measure risk by focusing on unexpected *negative* income falls, and develop a methodology that estimates channels of downside-risk sharing. Leveraging the heterogeneity of the Panel Study of Income Dynamics, we show that previous literature has significantly underestimated the importance of the fiscal channel in the US. We generalize our framework to estimate downside-risk sharing for increasingly severe negative shocks, an issue of particular relevance to policymakers that prior methodologies could not address. Our results show that also in these cases the federal government hedges substantial shares of income growth variation. Moreover, we find that fiscal risk sharing is larger for relatively lower-income households, and that its relevance increased during our sample.

JEL classification: E64, F40, F45, H29.
Keywords: Risk sharing, Downside risk, Federal taxes and transfers.

*We thank Pierfederico Asdrubali, Barry Eichengreen, Daniel Gros, and Anna Rogantini Picco for their discussion of a previous version of this paper, and seminar participants at Banca d'Italia, Banco de España, CEPS-ECFIN, ESCB Research Cluster 2 Workshop, the 3^{rd} International Conference on European Studies, the 20^{th} LAGV International Conference in Public Economics, the XXXII Villa Mondragone International Economic Seminar, the 23^{rd} INFER Annual Conference, and the Fourth Conference on Law and Macroeconomics for useful comments and suggestions. We are also grateful to Marco Del Negro, Paolo D'Imperio, Andrea Finicelli, Evi Pappa, Javier Pérez, Ivan Petrella, Luigi Federico Signorini, and Fabrizio Venditti. The views expressed in this paper are ours and do not necessarily reflect those of Banca d'Italia. All remaining errors are our own. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

†Kellogg School of Management, Northwestern University, federico.lenzi@kellogg.northwestern.edu.
‡Banca d'Italia, luca.rossi@bancaditalia.it.

1 Introduction

How much does the existence of a federal government help members of a fiscal union share risks with each other? This question sparked the literature on empirical risk sharing around the time of European monetary unification (Sala-i-Martin and Sachs, 1992). In this period, economists asked whether a monetary union could function properly without the existence of a central government that collects higher taxes from members that are experiencing idiosyncratic expansions and sends transfers to those that are temporarily struggling relatively to the broader federation.[1] By taking the US as an example, the literature provided insights about what a mature federation can achieve in terms of fiscal risk sharing. The seminal contribution by Asdrubali et al. (1996) proposed an econometric model for the joint estimation of three channels of risk sharing, namely the capital market channel, the credit channel, and the fiscal channel, finding a relatively modest role for the latter. Their work marked a turning point, for much of the subsequent and current literature has built on its foundations, generally corroborating the original results.

We highlight two characteristics common to all the most recent contributions studying channels of risk sharing in the US. First, they do not distinguish between risk sharing driven by a reduction in upside risks versus that driven by downside risks. In other words, they implicitly average the effects of the insurance provided by given risk sharing channels against negative income shocks with those stemming from positive income shocks. Second, those papers use state-level data, but - as we argue below and as our results will show - data at the state level suffer from excessive aggregation. We depart from most of the previous literature and use micro data from the University of Michigan Panel Study of Income Dynamics (PSID), as we believe using micro data is a necessary

[1] Differently from monetary policy, fiscal policy can indeed be targeted or designed to affect specific regions, thereby helping to balance the economic outcomes of union members (Kenen, 1969). Farhi and Werning (2017) provide theoretical support for the usefulness of government intervention within a currency union, whereby the full benefits of risk sharing can be reaped only when such an intervention takes place. In their own words, fiscal integration and financial integration are not perfect substitutes.

step if one wants to properly compute fiscal risk sharing.[2] Indeed, the within-state distribution of US income is quite unequal: according to the 2019 PSID wave, average per capita income is representative of incomes at the 67^{th} percentile of the distribution. This means that using aggregate state-level data implicitly provides risk sharing estimates that focus on relatively affluent individuals, who are also those that depend relatively less on the proper functioning of risk sharing channels, since they can usually count on stronger savings buffers. In addition, state-level data conceal the effects both of federal interventions that are typically targeted at lower-income households, and of federal tax payments that affect disproportionately more higher-income ones. In 2019, 50% of all social security benefits was received by households in the bottom 23% of the income distribution, and the top 11% of income earners paid 50% of all federal taxes. The above descriptive evidence is a useful starting point to explain why as of today the literature has recovered relatively small effects of the fiscal risk sharing channel. The results emerging from our framework challenge these previous findings.

This paper makes several contributions to the empirical risk sharing literature. First, we introduce a novel methodology to estimate the extent of risk sharing specifically associated with downside risks, which we dub "downside-risk sharing". Importantly, the usefulness of our framework is not confined to estimates of fiscal risk sharing, rather it can successfully be applied to any risk sharing channel and at any level of aggregation. For example, while our interest is in the fiscal channel, we also estimate consumption risk sharing. Second, thanks to available household-level data we show that using state-level data as has been done by the previous literature introduces a downward bias in fiscal risk sharing estimates; we argue that this arises because aggregate data inherently abstract from within-state inequality. Third, we provide downside-risk sharing estimates by income subgroups, for policymakers are likely more interested in gauging the fiscal risk sharing effects on the relatively lower-income segments of the population. Fourth,

[2] Asdrubali et al. (2020) use *Italian* micro data to explore risk sharing channels.

we generalize our framework to estimate risk sharing of progressively larger negative shocks, something that we would not be able to achieve with standard risk sharing metrics. This is of interest to policymakers governing sub-areas that are facing particularly severe idiosyncratic downturns.

We find that computing risk sharing with downside risks only rather than using both tails of the residual income growth distribution leads to significantly different - and notably higher - fiscal risk sharing estimates. We find that the federal government hedges a remarkably higher share of income growth than what the previous literature has found: while Asdrubali et al. (1996) (ASY henceforth) and later contributions obtained fiscal risk sharing estimates at approximately 15% after two years, our findings in fact indicate a tightly estimated value of 78.9%.[3] Importantly, we show that virtually all risk sharing capabilities are indeed driven by a reduction in the left tail thanks to the federal system of taxes and transfers, while the right tail proves to be remarkably stable after accounting for government intervention. Because of the problems discussed above, state-level data completely conceals this result as downside-risk sharing estimates obtained with aggregate data are not significantly different from risk sharing ones, and positive risk sharing estimates stem from a reduction in both tails.

We also estimate downside-risk sharing separately for different quartiles of the distribution and find that poorer households are those that enjoy higher income insurance from government stabilizers. Moreover, our baseline measure shows in a clearer way that public risk sharing has increased over time, especially after the Great Financial Crisis.

To the best of our knowledge, we are the first in the risk sharing literature to exclusively and explicitly focus on left tails and gauge the capability of a given mechanism to share *downside* risks.[4] The recent financial, macro-econometric and forecasting literature

[3] We use the two-year time interval since the PSID is a biennial survey, which precludes us from providing estimates for non-biennial time intervals.

[4] Dynarski and Gruber (1997) is the only work that acknowledges the importance of income contractions versus expansions, but their paper does not develop a downside-risk sharing methodology. The authors indeed only perform a brief exercise where they estimate their model separately for observations with negative versus positive income growth.

has clearly recognized the importance of downside risks, see Giglio et al. (2016), Adrian et al. (2019), Plagborg-Møller et al. (2020), Brownlees and Souza (2021), Adrian et al. (2022), Carriero et al. (2024) and Delle Monache et al. (2024) among others. As recalled by Estrada (2006) and in the references therein, investors associate risk with negative outcomes. When the distribution of stock returns is perfectly symmetric this distinction is irrelevant, yet the concept of standard deviation in the presence of significant skewness is of little use. In such an environment, an investor that aims at minimizing the effects of negative events should estimate the volatility of the left tail only. Markowitz (1959) first provides a way to do this by proposing his concept of semi-variance. In Section 2.3, we adapt his definition to our framework in order to obtain estimates of downside risk in aggregate income measures.

Related literature. Our work is closely related to Athanasoulis and Van Wincoop (2000) and Athanasoulis and Van Wincoop (2001) (AVW henceforth) who develop a novel framework that defines risk sharing as the reduction in the volatility of idiosyncratic shocks after a given risk sharing channel is put in place. This is in contrast with previous contributions - most notably ASY - who compute risk sharing with the covariance between idiosyncratic income growth before and after risk sharing occurs. Crucially, the basic *concept* of risk sharing we use in this paper is the same as AVW, but we differ from them in concentrating on downside risks and developing the first methodology that gauges downside-risk sharing. Section 2.5 discusses the conceptual differences between AVW and ASY, while Section 2.6 explains why we deem necessary to shift to a downside-risk framework for risk sharing.

By using micro data, we get closer to what done by the earlier literature on empirical risk sharing (Cochrane (1991), Mace (1991), Nelson (1994), Hayashi et al. (1996), Dynarski and Gruber (1997)). Asdrubali et al. (2020) note that in order to pin down risk sharing channels at the household level the econometrician necessarily has to control for more confounding factors than what would be required or even possible at the aggregate level,

and suggest that this could explain why the bulk of the recent literature opted to use aggregate data. In our paper, we overcome this difficulty by controlling for more than 30 variables and estimating the model with regularization methods. Since the approach developed by AVW and that we build upon is based on predictive regressions, we can deliberately bias parameters estimates in the conditional mean in order to minimize residual variance, our main object of interest in the models we estimate.

As recalled above, the empirical literature on fiscal risk sharing dates back to Sala-i-Martin and Sachs (1992), who perform regressions of either taxes or transfers on disposable income, and find that between one third and one half of an idiosyncratic shock to the average US region is cushioned by the federal government within the same year. Von Hagen (1992) uses GDP instead of personal income and finds very low elasticities, thereby concluding that the US fiscal system does not provide a significant stabilizing function. Obstfeld and Peri (1998) and Poghosyan et al. (2016) obtain similar results using different empirical strategies. Bayoumi and Masson (1995) is an exception, as the paper finds a high degree of stabilization (approximately 30%) granted by the US federal government.

Sørensen and Yosha (1998) use the method in ASY to estimate risk sharing channels among EMU countries as well as OECD ones; Hepp and Von Hagen (2012) and Hepp and Von Hagen (2013) estimate risk sharing channels among German states, while Burriel et al. (2020) discuss the federal structures of Austria, Belgium, Germany, Spain and the United States and run ASY-like regressions for each of them (except for Belgium whose federal system only consists of three regions).

Milano (2017) extends the approach in ASY and dissects the credit market channel of risk sharing into the private channel, national borrowing, and official credit provided by European institutions. She finds that European countries smooth consumption mainly thanks to counter-cyclical public borrowing, yet after the Great Financial Crisis this concealed heterogeneity between so-called "periphery" and "core" countries, as austerity measures in the former meant that national fiscal policies provided dis-smoothing during

the European sovereign debt crisis, which was more than offset by aid from European institutions. The opposite occurred in core countries. Cimadomo et al. (2020) also find that official loans were important drivers of increased risk sharing in the euro area in the aftermath of the crisis. Methodologically, this paper explicitly models time variation in risk sharing by adding interaction terms to the ASY framework.

Ferrari and Rogantini Picco (2023) use a synthetic control method and find that the introduction of the euro led to reduced risk sharing among member countries. Recently, some papers explored heterogeneities in risk sharing in the US (Stempel, 2021) and among OECD countries (Asdrubali et al., 2023). Cimadomo et al. (2023) extends Asdrubali and Kim (2004) and adopts a full-fledged Bayesian approach to construct time-varying estimates of risk sharing. The authors also propose a test to gauge whether different risk sharing channel are substitutes or complements.

Finally, Alcidi et al. (2023) extends US and EA analyses based on ASY using the most recent data, generally confirming previous results.

2 Methodology

The starting point of our modeling approach is the framework in AVW, which we review in Section 2.1. Section 2.2 proposes two simple but in our opinion necessary modifications to their framework. Then, Section 2.3 advances AVW's methodology by proposing a strategy to assess the extent to which the channel under consideration is capable of reducing semi-variances - namely, downside risk - as defined in Markowitz (1959) rather than standard deviations as in AVW. Finally, Sections 2.5 and 2.6 discuss AVW as well as our methodology with a broader perspective. Like other papers, also AVW use state-level data. Here we present the framework for an household-level analysis, yet our methodology can be indistinctly be applied to state-level data with minor modifications.

2.1 AVW - a review

Let $i = 1, \ldots, N$ denote household i, ω_t^i a generic income measure, and let $g_{t,t+h}^i = \log \omega_{t+h}^i - \log \omega_t^i$ be the growth rate of income from period t to $t+h$. While households can insure idiosyncratic risk, they can do little to diversify away aggregate risk. AVW write the difference between state and aggregate growth as the sum of a predictable and an unpredictable component, where the residual corresponds to the growth component which is diversifiable. In our application, we use the equivalent approach of not subtracting for aggregate growth but rather to include time fixed effects. Therefore, the model to be estimated is

$$\begin{aligned} g_{t,t+h}^{b,i} &= \boldsymbol{\beta}_h'^b \mathbf{x}_{t,h}^{b,i} + \epsilon_{t,t+h}^{b,i}, \\ g_{t,t+h}^{a,i} &= \boldsymbol{\beta}_h'^a \mathbf{x}_{t,h}^{a,i} + \epsilon_{t,t+h}^{a,i}, \end{aligned} \tag{2.1}$$

where the superscripts b and a refer to variables from regressions whose dependent variable is income measured before and after risk sharing, respectively. $\mathbf{x}_{t,h}^{\cdot,i}$ is a $(K+1) \times 1$ vector of control variables which includes a constant term, fixed effects, dummy and continuous variables, and we assume that $\epsilon_{t,t+h}^{\cdot,i} \overset{iid}{\sim} (0, \sigma_h^{2,\cdot})$. This system is estimated using non-overlapping intervals of length h, which is necessary in order to obtain state-specific independent innovations.[5]

In this setting, $\sigma_h^{2,b}$ has to be interpreted as a measure of diversifiable risk. AVW define the amount of risk sharing accomplished by a given channel as the relative reduction in the residual standard deviation of income growth after risk sharing as compared to the one before it. Therefore, their estimates of risk sharing at horizon h are given by

$$RS_h^* = 100 \cdot \frac{\hat{\sigma}_h^b - \hat{\sigma}_h^a}{\hat{\sigma}_h^b}. \tag{2.2}$$

[5]The reliance on non-overlapping intervals is the reason why we need a subscript h in the control variables. Moreover, AVW assume Gaussianity in the error terms and obtain confidence bands with analytic approximations. Here, we instead drop this assumption and perform inference with bootstrapping procedures. Indeed, standard Gaussianity tests on estimated forecast errors reject the null of normality.

Crucially, this modeling approach does not require to recover structural parameters: all that is needed is to obtain reliable estimates of the conditional mean. Thus, in principle one can use Ordinary Least Squares (OLS) to estimate equation (2.1) and its variants below. Nevertheless, given the considerations above, we prefer to include in our regressions many predictors and not use OLS, which is known to be prone to over-fitting issues in the presence of many regressors. We instead estimate our model with Lasso (Tibshirani, 1996), which shrinks and selects variables with relatively low pseudo-out-of-sample predictive ability.

Finally, given that our goal is not to test the risk sharing theory but rather to obtain *aggregate* measures of *idiosyncratic* risk, we give higher weight to more representative households.[6]

2.2 AVW - modifying the basic framework

In this section, we acknowledge that the way risk sharing is computed in AVW can first be improved over some dimensions, and we therefore first examine those intermediate steps. First, equation (2.1) does not allow for state fixed effects, something that is shared by other contributions in the literature. AVW essentially claim that controlling for state-and-time specific information should be sufficient to capture idiosyncratic differences in growth rates. We disagree with this statement: when we run AVW without state fixed effects we find clear evidence of non-zero-mean state-specific residuals, and we thereby add state and division fixed effects in our baseline specification. As we explain in Section 2.5, variance reduction stemming from persistent differences in growth rates is due to redistribution rather than risk sharing, as those differences are indeed predictable.

Then, we believe that computing risk sharing with variances rather than standard

[6] In the following exposition, we aim at being as intuitive as possible and therefore outline how we would calculate risk sharing if we were *not* to weight observations by population or - in the PSID application - by household weights. Our baseline weighted regressions need slightly different computations, which we describe in Appendix A.

deviations as in AVW is more appropriate, for two reasons. First, variances are additive while standard deviations are not: if we normalize risk reduction by total variation rather than the residual one (as we suggest in the next paragraph) we then necessarily have to recur to variances. In addition, calculating risk sharing with ratios of second moments rather than their squared roots is more consistent with the rest of the literature, which in different ways always pins down risk sharing with regression coefficients which by their very nature are the result of operations involving (mixed) second moments as well. We therefore first modify equation (2.2) as follows:

$$RS_h^{**} = 100 \cdot \frac{\hat{\sigma}_h^{2,b} - \hat{\sigma}_h^{2,a}}{\hat{\sigma}_h^{2,b}}. \tag{2.3}$$

Now, by means of a numerical example we argue that computing risk sharing as in equations (2.2) and (2.3) can lead to misleading estimates. Let's first assume that estimating the model in equation (2.1) leads to point estimates $\hat{\sigma}_h^b = 20\%$ and $\hat{\sigma}_h^a = 10\%$ for some h. In this case, equation (2.3) suggests that the channel under consideration is able to hedge $100 \cdot \frac{0.2^2 - 0.1^2}{0.2^2} = 75\%$ of idiosyncratic risk in income growth. In economic terms, this result points to this channel being particularly useful, as the risk reduction is sizeable not only in percentage terms but also in absolute ones (a ten *percentage points* reduction in income growth risk). Now, suppose instead that fitting the same model on the same variables yields estimates $\hat{\sigma}_h^b = 0.5\%$ and $\hat{\sigma}_h^a = 0.25\%$. Also in this case, equation (2.3) would imply the same 75% risk sharing value. *However*, it is clear that risk reduction in absolute terms is now much lower, equal to a mere fourth of a percentage point as compared to ten in the previous scenario. In other words, the latter high risk sharing estimate is of little economic relevance as idiosyncratic risk was already low *before* risk sharing took place. Generally speaking, the same values of both RS_h^* and RS_h^{**} in equations (2.2)-(2.3) are consistent with a whole range of potentially very different values of the amount of risk that agents would like to insure against.

Given all the above, we therefore deem that risk sharing should be computed not as a

percentage of the before-risk-sharing *residual standard deviation* as is done in AVW, but as a share of *total variance*. Thus, we further modify the calculation of risk sharing:

$$RS_h = 100 \cdot \frac{\hat{\sigma}_h^{2,b} - \hat{\sigma}_h^{2,a}}{\hat{\sigma}_h^{2,b,tot}} = RS_h^{**} \frac{\hat{\sigma}_h^{2,b}}{\hat{\sigma}_h^{2,b,tot}} < RS_h^{**}, \tag{2.4}$$

where $\hat{\sigma}_h^{2,b,tot}$ is the sample variance of our before-risk-sharing dependent variable at every horizon. Note that this adjustment leads to risk sharing estimates that are mechanically *lower* than those obtained in equation (2.3).

Continuing with the example above, assume that $\hat{\sigma}_h^{b,tot} = 30\%$. In the first case, we would now obtain risk sharing estimates equal to $75 \cdot \left(\frac{0.2}{0.3}\right)^2 \approx 33.3\%$, while in the second case they would drop to $75 \cdot \left(\frac{0.005}{0.3}\right)^2 \approx 0.02\%$, reflecting the much lower absolute risk reduction. In the extreme case where $\hat{\sigma}_h^b = \hat{\sigma}_h^{b,tot}$, we would obtain "adjusted" risk sharing values equal to those in equation (2.3). On the other hand, in the opposite scenario where $\hat{\sigma}_h^b \to 0$ and we keep $\hat{\sigma}_h^b = 2\hat{\sigma}_h^a$, our formula would suggest zero risk sharing, while equation (2.3) would always yield (with the example above) a 75% hedging capability.

2.3 Downside-risk sharing

This paper aims at providing estimates of *downside*-risk sharing. In the following we therefore modify the framework above to obtain them. Markowitz (1959) first acknowledges the limitations of standard deviations as a way to compute the underlying risk of stock prices: large positive returns should not contribute to increase risk estimates of a certain asset. Nevertheless, the mathematical definition of variance inherently neglects signs and treats positive and negative returns in the same way. Therefore, Markowitz (1959) introduced the concept of semi-variance: for a given series $\{r_t\}_{t=1}^T$ of stock returns, he defines conditional returns $r_t^0 = \min(r_t, 0)$ and computes the sample variance of $\{r_t^0\}_{t=1}^T$ as a measure of the downside risk of that asset.

In our paper we apply the same definition, but we do this on the series of *residuals*

we obtain from estimating the system in equation (2.1). The underlying reasoning is the same as the one in Markowitz, simply converted to fit in our macro-forecasting setting: we are interested in gauging how much the risk sharing channel we are considering is able to hedge against unexpected left-tail events. Computing risk sharing as in AVW would bias the risk measure we are interested in towards the volatility of unexpected *positive* idiosyncratic income growth events.

Operationally, we then define

$$\begin{aligned}
\hat{\epsilon}_s^{b,i,0} &= \min(\hat{\epsilon}_s^{b,i}, 0), \\
\hat{\epsilon}_s^{a,i,0} &= \min(\hat{\epsilon}_s^{a,i}, 0), \\
\hat{\sigma}_h^{2,\cdot,0} &= \frac{1}{S_h \cdot N - K - 1} \sum_{s=1}^{S_h} \sum_{i=1}^{N} \left(\hat{\epsilon}_s^{\cdot,i,0}\right)^2, \\
g_s^{b,i,0} &= \min(g_s^{b,i}, 0), \\
\hat{\sigma}_h^{2,b,0,tot} &= \frac{1}{S_h \cdot N - K - 1} \sum_{s=1}^{S_h} \sum_{i=1}^{N} \left(g_s^{b,i,0}\right)^2,
\end{aligned} \tag{2.5}$$

where the index s refers to each non-overlapping interval for a given horizon h, and S_h is the total number of intervals for that given horizon.

We therefore define our baseline downside-risk sharing measure as:[7]

$$DRS_h^0 = 100 \cdot \frac{\hat{\sigma}_h^{2,b,0} - \hat{\sigma}_h^{2,a,0}}{\hat{\sigma}_h^{2,b,0,tot}}. \tag{2.6}$$

2.4 Hedging against larger shocks

The simplicity and flexibility of our framework enable us to extend our downside-risk analysis even further. One might argue that households (or states) are also interested in understanding the federal government's capacity to respond to *large* negative shocks.

[7] In this case, it could happen that models where the control variables explain little variation could lead to residual *semi*-variance being slightly *higher* than its total counterpart, thereby leading adjusted risk sharing values to be higher than non-adjusted ones.

Therefore, we generalize our framework and define downside risk by setting progressively negative thresholds in the conditional forecast errors distribution. In particular, equation (2.5) now becomes:

$$
\begin{aligned}
\hat{\epsilon}_s^{b,i,c} &= \min(\hat{\epsilon}_s^{b,i}, c), \\
\hat{\epsilon}_s^{a,i,c} &= \begin{cases} \min(\hat{\epsilon}_s^{a,i}, 0) & \text{if} \quad \hat{\epsilon}_s^{b,i} < c, \\ c & \text{if} \quad \hat{\epsilon}_s^{b,i} \geq c, \end{cases} \\
\hat{\sigma}_h^{2,b,c} &= \left(|c| + \sqrt{\frac{1}{S_h \cdot N - K - 1} \sum_{s=1}^{S_h} \sum_{i=1}^{N} \left(\hat{\epsilon}_s^{b,i,c} \right)^2} \right)^2, \\
\hat{\sigma}_h^{2,a,c} &= \frac{1}{S_h \cdot N - K - 1} \sum_{s=1}^{S_h} \sum_{i=1}^{N} \left(\hat{\epsilon}_s^{a,i,c} \right)^2, \\
g_s^{b,i,c} &= \min(g_s^{b,i}, c), \\
\hat{\sigma}_h^{2,b,c,tot} &= \left(|c| + \sqrt{\frac{1}{S_h \cdot N - K - 1} \sum_{s=1}^{S_h} \sum_{i=1}^{N} \left(g_s^{b,i,c} \right)^2} \right)^2,
\end{aligned}
\tag{2.7}
$$

where c defines the minimum income growth loss for which we want to compute downside-risk sharing. Note that when $c = 0$ the system in (2.7) reduces to that in (2.5). Downside-risk sharing of larger shocks is thus computed as:

$$
DRS_h^c = 100 \cdot \frac{\hat{\sigma}_h^{2,b,c} - \hat{\sigma}_h^{2,a,c}}{\hat{\sigma}_h^{2,b,c,tot}}. \tag{2.8}
$$

2.5 On the importance of using residual volatility

We develop our methodology by building on AVW, who estimate risk sharing within a very different conceptual framework than ASY, in particular by gauging the extent to which a given risk sharing channel reduces the *residual* volatility of income growth. Interestingly, their definition of risk predates later contributions such as Jurado et al. (2015)

and Fernández-Villaverde et al. (2015) who define uncertainty as the volatility of the *unexpected* variation of a given variable.[8] As exemplified in AVW, if the correlation between income before and after risk sharing is zero with the two series having the same residual standard deviation, ASY would imply the existence of perfect risk sharing, while AVW would point to no risk sharing at all. Moreover, the framework in AVW does not require to assume a particular process for income, whereas ASY needs to impose that GDP is exogenous.

Empirical work focusing on the estimation of structural parameters has to deal with potentially biased estimates stemming from - among other things - omitted variable bias and contemporaneous causality. ASY abstracts from this issue by assuming idiosyncratic GDP growth to be exogenous; under this assumption, there is no need to control for the predictable variance of income growth. *Nevertheless*, later empirical contributions - see, for example, Mélitz and Zumer (1999) and Asdrubali and Kim (2004) - acknowledge the problems of previous static empirical models and control for potentially relevant variables.[9] As Asdrubali and Kim (2004) state: "*It is somewhat surprising - given the economic structure of the issue at hand - that the literature has so far neglected the use of dynamic and simultaneous econometric models [...] our approach makes clear that a simple static regression of a smoothing channel on [GDP] [...] is subject to a bias, due to the omitted variables [...] and their lagged values, and due to the neglect of the dynamics and interaction among smoothing channels*".

In general, and differently from ASY, the aim of AVW - and our paper - is not to test the theory of risk sharing, but rather to properly estimate ex-ante risk and gauge the extent of its reduction thanks to different risk sharing channels. From AVW: "*ASY do not control for the predictable component of growth rates. It is important to do so if one wants to measure risk sharing, which is about reducing risk exposure rather than smoothing a predictable income path*".

[8] In general, the literature on uncertainty has clarified the importance of distinguishing between predictable variations and the volatility surrounding the unpredictable component of any series. See Jurado et al. (2015) for an in-depth discussion.

[9] ASY themselves have a short paragraph (p. 1103) where they show that in their application results do not change significantly when they control for lagged income growth rates.

In addition to the above considerations - which apply to all risk sharing channels - we argue that controlling for the predictable variation is even more important if one is interested in estimating fiscal risk sharing. As for what concerns the fiscal channel, we argue that *all* the reduction in the variance of income growth that is due to the predictable component is *not* due to risk sharing. Rather, it is a by-product of fiscal redistribution.[10] We explain this statement with a simple example. Suppose the existence of a federation which consists of only two states A and B with the same population which is constant over time. Both states have per capita GDP *levels* equal to \$100 in the base period, and it is *known* with certainty that - thanks to fixed investments made in the past - A will grow by 6% during the year, whereas - because of structural inefficiencies in the economy - B will contract by the same amount in the same period, which implies zero growth nationwide and state growth rates equal to idiosyncratic ones. At the end of the period, the federal government steps in by taxing A \$2, transferring this amount to B. Therefore, growth rates after government intervention will equal 4% (as opposed to 6%) in A, and -4% (as opposed to -6%) in B. In other words, federal government intervention reduced the variance of income growth. However, we know that states in this federation were not hit by idiosyncratic income shocks (we are assuming a deterministic world). Rather, their different growth rates were due to structural *and therefore predictable* reasons (past investments in one country, inefficiencies in the other). Thus, the federal tax and transfers payments and the consequent reduction in the variance of income growth were due to redistribution as opposed to risk sharing. In general, and as we said in the paper, variance reduction stemming from persistent (structural) differences in growth rates between states is rather due to fiscal redistribution than risk sharing, as those differences are indeed predictable. Any model that is estimated without accounting for this will obtain biased estimates (most likely overestimates) of the risk sharing capabilities of the tax and transfer system. The distinction between predictable and unpredictable variation is also

[10] In this paper, we use the definition of redistribution available for example in Malkin and Wilson (2013).

extremely relevant for policy purposes, for instance in the European Union where the long debate regarding risk sharing, redistribution, and fiscal discipline needs to be properly informed.

2.6 On the need for a downside-risk approach

To fully understand why one should focus on downside risks, it is useful to make a step back and ask why uncertainty matters at all.[11] One of the most important reasons why macroeconomists are concerned with uncertainty is that higher uncertainty leads to lower economic activity, something that is well established in the literature.[12] As Bloom (2014) highlights, the channels through which this phenomenon occurs are mainly two, namely a real options channel and a financial channel.[13],[14] In the first, it is assumed that a substantial fraction of agents in the economy face investment opportunities that have three crucial features: i) they are partially or completely irreversible; ii) their payoffs are uncertain; iii) they can be postponed to obtain more information about future outcomes. This is the situation, for example, of a firm considering a capital expenditure, or a household deciding whether to increase discretionary spending on goods and services, or under-

[11] In the following, we use the words "risk" and "uncertainty" interchangeably. This is because the way AVW and we compute risk in the paper is not conceptually different than what done e.g. in Jurado et al. (2015) when they compute macroeconomic uncertainty.

[12] See, for example, Bloom (2009), Christiano et al. (2014), Jurado et al. (2015), Fernández-Villaverde et al. (2015), Baker et al. (2016), Bloom et al. (2018), Carriero et al. (2018), Anzuini et al. (2020), Husted et al. (2020), Anzuini and Rossi (2021).

[13] For a thorough presentation of the fundamentals of real options, see Dixit and Pindyck (1994).

[14] Bloom (2014) also cites two channels through which higher uncertainty can incentivize investments, i.e. a growth options channel and the so-called Oi-Hartman-Abel effects. Those effects are nonetheless not particularly relevant for households, while being confined to a relatively small subset of firms, namely those that devote a high share of their balance sheet to investments having long time-to-build and/or those that enjoy very high short-run flexibility in their ability to expand or contract production capacity. In general, there exists scant empirical evidence on their relevance. Segal et al. (2015) separate good and bad uncertainty and indeed find that spikes in uncertainty have opposite effects on the economy depending on whether it is the left tail or the right tail to expand. Nevertheless, one can see the studies that quantify the effect of uncertainty on the macroeconomy as ones that implicitly compute the *net* effects of uncertainty (namely, the adverse effects of downside uncertainty net of the positive effects of upside uncertainty) and, to the best of our knowledge, none of them find an aggregate positive effect.

take a large investment such as purchasing or renovating a house.[15] In a seminal paper, Bernanke (1983) lays out what is known as the "bad news principle", whereby he shows that whenever an agent is contemplating to undertake an irreversible investment, only the size and probability of a bad outcome matter for his or her decision to delay the investment rather than paying the sunk cost today. Intuitively, when it becomes more likely that the future payoffs of an investment opportunity will not be worth the investment cost, the agent will have a greater incentive to wait (a higher "option value of waiting") until the uncertain outcome is realized, and invest only if the good state of the world materializes. On the other hand, if the probability that the investment will yield the same high returns increases or if the expected magnitude of those favorable future returns rises - leading in both cases to higher upside uncertainty - this will not drive agents to change their behavior, the reason being that if they knew they will find themselves in the good state of the world, they would invest immediately anyway. More specifically, whenever upside uncertainty increases, agents will regret not to have invested earlier. Since this is not feasible, the shape of the right tail simply does not factor into their optimization problem. As Bernanke (1983) says: "*[...] given the current return of the most profitable investment* i, *the willingness to invest in* t *depends only on the average expected severity of bad news for* i *that may arrive in the next period. Potential good news for the investment does not matter at all*". Finally, as recalled by Bernanke (1983) and by Dixit and Pindyck (1994), this result does not depend on agents being risk-averse. Therefore, as for the real options channel, neither upside uncertainty nor risk preferences substantially matter for the basic mechanism to work.

The second channel that links uncertainty to economic growth - the financial channel, useful to explain the relevance of e.g. the capital market channel of risk sharing - arises

[15] In the cases of household purchases of goods and services, payoffs are uncertain not necessarily because the stream of utility expected to be enjoyed after the purchase is not known ex-ante, but rather because the law of motion of the cost of the purchase is stochastic. Rossi (2013) shows that the canonical models in real options theory yield the same qualitative suggestions if one substitutes the hypothesis of a stochastic value of the investment with that of a stochastic cost. On the relevance of uncertainty for consumption decisions see also Romer (1990) and Eberly (1994).

from the fact that increased uncertainty leads investors to ask for higher risk premia, both because of risk aversion and because of heightened default probabilities. *However*, this mechanism mainly works through changes in the left tail only. Apart from Markowitz (1959) there exists a wealth of literature that either assumes, finds, or argue that it is optimal for agents to place more weight on losses rather than gains. Roy (1952) formalizes the "Safety first" principle, postulating that it is reasonable for individuals to "*seek to reduce as far as is possible the chance of [...] a catastrophe occurring*"; in their experiments, Kahneman and Tversky (1979) find that individuals value more a monetary loss than an equivalently sized gain, and that this tendency becomes increasingly more pronounced as the magnitude of the payoffs symmetrically increases. Gul (1991) establishes a theoretical framework for disappointment aversion, which was later generalized by Routledge and Zin (2010). Ang et al. (2006) find the existence of a large downside-risk premium driven by investors who require a higher compensation for holding stocks that co-vary more with the market when the latter is falling; Estrada (2006) documents that "*many funds use* [downside risk] *to calculate risk-adjusted returns and include this measure of risk in annual reports. [...] Morningstar uses* [downside risk] *to create its influential star ratings*". Moreover, "*Investors think of risk differently from the way it is defined in modern portfolio theory. Both the standard deviation and beta give equal weight to upside and downside fluctuations. Investors, however, do not. It is perhaps for this reason that the downside risk framework has been rapidly gaining acceptance among academics and practitioners. Semideviation [...] captures the downside volatility that investors want to avoid and not the upside volatility investors are seeking*". Lettau et al. (2014) provide empirical evidence that only the downside-risk CAPM (as opposed to its traditional counterpart) can explain the cross-sectional dispersion of currency returns, and that it can jointly explain the cross section of the returns of a variety of asset classes. Farago and Tédongap (2018) find evidence that - in addition to market returns and volatility - markets price three left-tail-related factors in the cross section of asset returns.

In addition to what discussed in the above-mentioned literature, widely used financial risk metrics such as Value-at-Risk or Expected Shortfall corroborate that agents are mainly concerned with the exact shape of the left tail and usually aim at shrink it towards less harmful outcomes.

In this work, we acknowledge the above theoretical insights as well as the overwhelming evidence in favor of agents' disproportionate interest towards hedging against downside uncertainty, and build the analysis accordingly. In this regard, we are aligned with Lettau et al. (2014): "*Given its demonstrated empirical relevance, it is important to gain a deeper theoretical understanding of the sources and time variation of downside risk. It remains an open question whether downside risk comes from preferences or from micro-founded constraints*".

3 Data and estimation strategy

We apply the model developed in Section 2 to the fiscal channel of downside-risk sharing in the US, and we exploit household-level data available from the PSID since 1968. During the years, the survey included individuals that joined the family and excluded those leaving it. Moreover, starting from 1997 the PSID also started to properly account for the flow of new immigrants. We restrict our attention to the Survey Research Center (SRC) sample which selected 2930 families representing the entire US population. The dataset provides a set of weights which allows researchers to obtain a sample that is representative of the US society in any given year. More specifically, we weight all observations according to the procedure illustrated in Gouskova et al. (2008). These weights reproduce the US population, compensating both for unequal selection probabilities and differential attrition.

For our purposes, in order to construct the dependent variables to be used in equation 2.1 - recall that to obtain risk sharing estimates within this framework we need to run two regressions - we need total taxable income, federal transfers received and federal

taxes paid. The transfers series we use includes Temporary Assistance for Needy Families (TANF), Social Security income, Supplemental Security income, veteran pensions, other retirement pensions, Individual Retirement Arrangements (IRA), unemployment compensation, worker compensation, child support income, heating subsidies, Supplemental Nutrition Assistance Program (SNAP), and other welfare-related transfers. As for taxes, the PSID does not directly provide them, but we compute them using total taxable income from the PSID and historical income tax rates obtained from the Tax Foundation.

Our income growth measure before risk sharing is equal to taxable income growth, while the corresponding measure after risk sharing first adds federal transfers and subtracts federal taxes before computing growth rates.[16]

Given the high number of regressors included, we estimate our baseline model with regularization methods, which have the advantage of allowing the researcher to include many variables and letting the algorithm decide which of them to maintain and whether to shrink them or not. We therefore include a constant term, the lagged dependent variable, household size, age of the head, education level of the head of household, education level of the wife, number of children, age of the youngest child, interview length, total working hours of head and wife, average hourly salary for head and wife, working days lost for sick leave of head and wife. We also control for the most relevant variables at the state level that were used in AVW, namely the log of per capita income, the growth rate of per capita income, the five-year lagged population growth rate, and the ratio of state government expenditure to GDP. In addition, we include the following set of dummy variables: year, state of residence, year and state of residence interaction, Division, family composition change, sex of the head, marital status of the head, dummy variable for new wife, primary race of the head and wife, employment status of the head and wife, work for pay of head, wife, or both of them, dummy for self-employed head, corporate job for

[16] A very low number of individuals enter the job market for the first time, thereby recording an extremely high income growth rate. To avoid our results to be driven by those observations, we censor these values at 600%.

the head, public job for the head, unionized job for the head, mode of interview.

Section 4 compares estimates obtained from PSID data with those based on data from the Bureau of Economic Analysis (BEA). As for the latter, the closest definition of income we can use is per capita personal income net of transfers. We measure state-level federal taxes as the sum of employee and self-employed contributions, and personal current taxes to the federal government. Appendix B describes how we compute federal transfers. State-level regressions include few regressors and we thereby estimate them with Weighted Least Squares, using population weights.

We estimate equation (2.1) on PSID data and compute downside-risk sharing for the 1997-2019 period. The reason why we focus on this time span is twofold. First - as Figure 3.1 shows - federal government intervention in the form of direct or indirect transfers to individuals has substantially increased over time. If we look at a sample starting in 1963 (the starting date for many past studies) one can see that those transfers weighed 5% of US GDP at the beginning of the sample, rapidly increased during the '70s, fluctuated around 10-12% during the Great Moderation, and then rose in the aftermath of the Great Financial Crisis and especially after the pandemic. Limiting our analysis to 1997-2019 avoids using the earlier sample when the role of the federal government was significantly lower, and ensures our results are not driven by the unprecedented federal aid disbursed to US households especially during 2020 and 2021. Moreover, below we compare results obtained with PSID data with those derived using state-level data from the BEA. Therefore, the second reason why we start our estimation from 1997 is that in this year the BEA replaced SIC codes with NAICS ones, thus by using only post-1997 we ensure accounting consistency in state-level regressions.

Finally, we perform inference following a procedure similar to the one in Kapetanios (2008). Specifically, we first sample an equal number of households within each income quartile, compute the risk sharing measure, and repeat the procedure 20,000 times in order to obtain confidence intervals. Sampling by percentiles allows mimicking the dis-

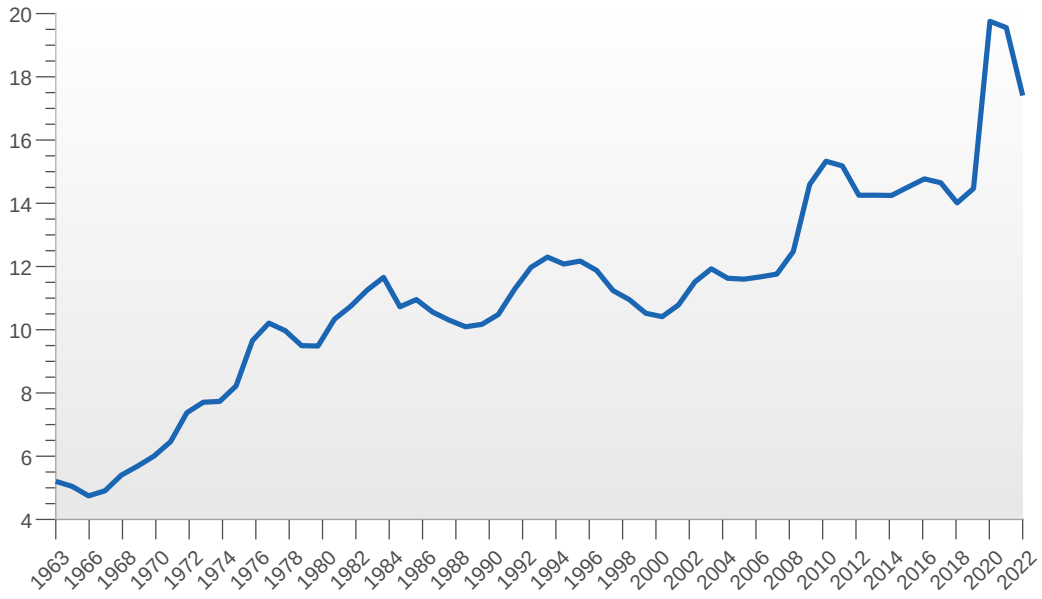


Figure 3.1: Total payments for individuals. The figure shows the evolution of the sum of direct payments for individuals and grants to state and local governments aimed at the personal sector, in percent of GDP. Data are taken from the Historical Tables of the Office of Management and Budget (OMB).

tribution of an original population characterized by non-negligible economic inequalities.

4 Results

The upper part of Table 1 reports results obtained by estimating fiscal downside-risk sharing with PSID data at the household level. We find that, contrary to many past contributions, fiscal risk sharing cushions a very large share of idiosyncratic income fluctuations. After two years, it is able to reduce downside risks by 78.9%, with this share progressively lowering as the horizon increases, until it reaches 54.1% after six years, still a sizable estimate. These values are higher than those found in many past research, strongly questioning the conclusion that the fiscal channel plays a relatively limited role in hedging income shocks. Moreover, the table indicates that the uncertainty around the point estimate is small.

In order to gauge the usefulness of using household-level data as opposed to aggre-

Table 1: Fiscal downside-risk sharing: household- and state-level data from the PSID

			68%		90%		95%	
	Horizon	Point estimate	Lower	Upper	Lower	Upper	Lower	Upper
Household								
	$h = 2$	78.9	77.0	80.3	75.9	81.5	75.9	81.9
	$h = 4$	59.6	57.6	63.0	55.6	63.5	54.5	66.3
	$h = 6$	54.1	51.2	55.2	50.2	60.9	50.2	60.9
State-level								
	$h = 2$	30.1	12.4	23.3	7.6	27.6	4.9	28.7
	$h = 4$	26.2	14.6	32.9	8.6	42.8	5.9	49.0
	$h = 6$	20.8	2.5	27.5	-6.5	35.0	-11.9	38.8

The table reports results obtained by estimating our downside-risk sharing model on household-level PSID data as well as PSID data aggregated at the state-level. 68%, 90%, and 95% empirical confidence bands are reported. Figures are expressed as percentages. Horizons are expressed in years.

gate data, Table 1 also reports the results from an additional exercise where we estimate our baseline model using household data aggregated at the state level. Estimates based on state-level data always lie well below those based on micro data, by 33 to 49 percentage points. This is suggestive evidence that points to the fact that estimating risk sharing with state-level data - namely, what most researchers have done in the last three decades - introduces a downward bias that we argue is due to the fact that all within state inequality is not considered; rather, in those papers the high heterogeneity in households' economic situations is substituted with a single representative agent for each state.

Figure 4.1 dissects our findings - based on PSID data - by showing which is the contribution of each of the variations we impose on AVW's methodology. In the top-left panel we provide estimates of risk sharing by using the methodology in AVW. Interestingly, while state-level estimates - not shown - are qualitatively similar to what found in AVW, when exploiting the heterogeneity of micro data we find significantly higher estimates, above 30% in the first four years. The top-right panel reports results obtained by using variances rather than standard deviations. This simple difference is alone responsible for a 15 to 20 p.p. increase in risk sharing estimates. The bottom-left panel shows that

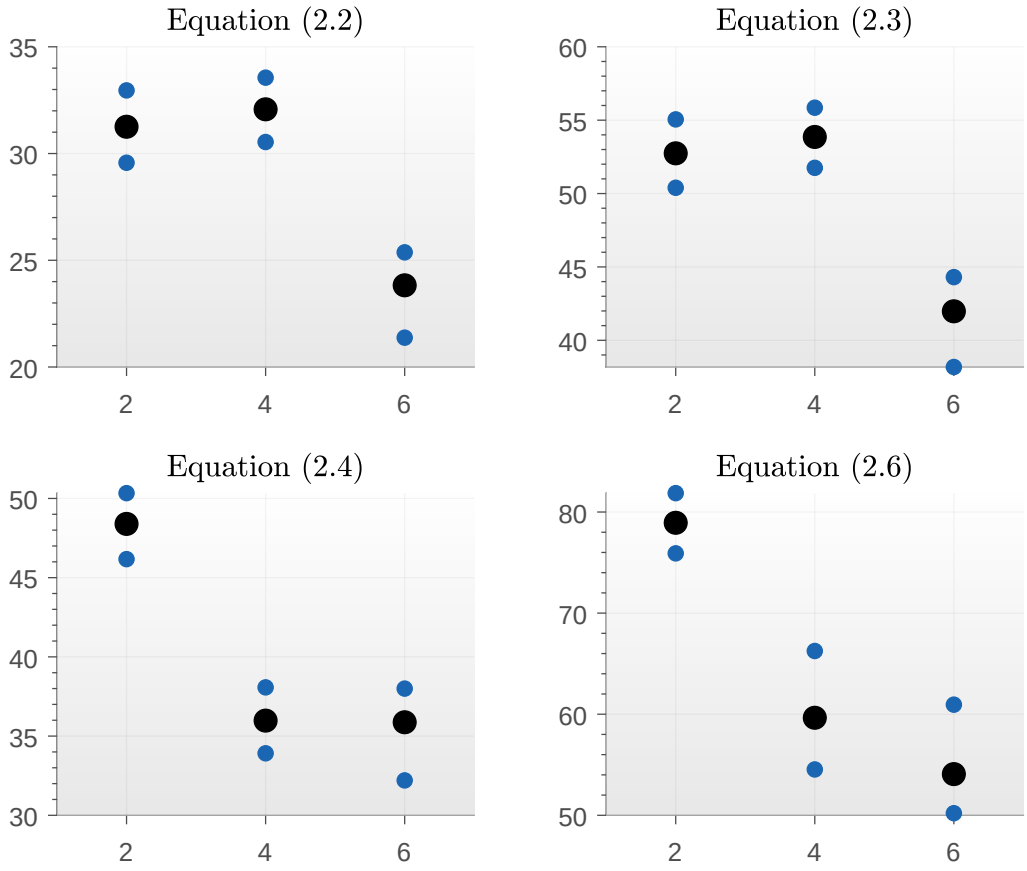


Figure 4.1: From AVW to our baseline. The figure provides estimates obtained by computing risk sharing with different metrics, as explained in Section 2. The x-axis represents the horizon (in years); the y-axis represents the risk sharing estimate (in percent). Black circles are our baseline estimates, taken as the median of the distribution; blue circles are 95% confidence intervals.

- as explained in Section 2.2 - normalizing risk reduction by total variance rather than residual variance leads to a mechanical reduction in risk sharing estimates. Finally, the comparison between the bottom-left and bottom-right panels shows the main empirical contribution of this paper: when focusing on downside risks - as in equation (2.6) - as opposed to the whole distribution - as in equation (2.4) - risk sharing estimates starkly rise at all horizons: after two years, risk sharing a-la AVW (but with our modifications) estimates that the federal government provides a 48.4% risk reduction, whereas our downside-risk sharing metric suggests 78.9%. After four (six) years, the former points to 36.0% (35.9%) whereas the latter to a markedly higher 59.6% (54.1%).

Table 2: Fiscal risk sharing and downside-risk sharing: state-level data from the BEA

			68%		90%		95%	
	Horizon	Point estimate	Lower	Upper	Lower	Upper	Lower	Upper
Risk sharing								
	$h = 2$	20.7	17.1	23.7	15.0	25.4	14.0	26.1
	$h = 4$	17.9	14.6	21.0	12.5	23.1	11.5	24.0
	$h = 6$	21.3	17.9	24.7	15.7	26.9	14.8	27.8
Downside-risk sharing								
	$h = 2$	22.4	19.5	24.7	17.6	26.2	16.7	26.8
	$h = 4$	17.5	13.5	21.6	11.0	24.2	9.7	25.5
	$h = 6$	23.4	18.6	28.3	16.1	31.4	14.9	33.0

The table reports results obtained by estimating fiscal risk sharing and downside-risk sharing as in equations (2.4) and (2.6) using BEA state-level data. 68%, 90%, and 95% empirical confidence bands are reported. Figures are expressed as percentages. Horizons are expressed in years.

Table 2 reports the results we obtain by estimating equations (2.4) and (2.6) on state-level data - described in Section 3 - from the BEA. First, note that estimates based on BEA data lie close to those obtained by previous papers. Moreover, aggregating PSID data at the state level (as in Table 1) leads to estimates that are close to those we obtain with BEA data in Table 2. This is reassuring as it makes us confident that the household data we use and build - recall we compute federal taxes based on historical income tax rates - correctly represent the US population. More interestingly however, Table 2 shows that differences between risk sharing and downside-risk sharing at the state-level are small and not significantly different between them. Therefore, relying solely on state-level data not only introduces a downward bias in standard risk-sharing metrics, but also obscures the significance of downside-risk sharing, making the two measures appear artificially similar.

Figure 4.2 presents residuals from models estimated using both household-level data from the PSID and state-level data from the BEA. Notably, the top panel reveals that for shorter horizons, government intervention effectively mitigates a substantial portion of

the left tail, as the left bound of the yellow distribution is significantly closer to zero than that of the blue distribution - particularly at the two-year horizon. However, the right tail remains remarkably stable across all cases. This implies that virtually *all* non-zero risk-sharing estimates derived from standard risk-sharing metrics using household-level data are driven solely by the compression of the left tail. In other words, incorporating the right tail in the computation effectively pushes risk-sharing estimates toward zero. Interestingly, the bottom panel shows that this asymmetry does not appear in state-level data, where reductions in the left and right tails are more balanced.

Finally, Figure 4.3 shows the results we obtain by estimating our model over eight-year rolling windows. Interestingly, when we look at risk sharing estimates that use both tails of the income shocks distribution - namely, the first three panels - we find no economically significant changes in the insurance role of the US federal fiscal system over time. This aligns with findings from previous research. Nevertheless, we observe a clear upward trend in *downside*-risk sharing estimates starting after the Great Financial Crisis. This is consistent with what we showed in Figure 3.1, namely that federal transfers to individuals markedly increased after 2009.

4.1 Downside-risk sharing by income quartiles

Thanks to the availability of micro data, we can potentially explore a very wide range of heterogeneities. An interesting one to explore is that based on taxable income differences. In a sense, it is not clear a priori that poorer households should benefit more from risk sharing within a fiscal federation: on the one hand - and as we showed in Section 1 - the government undoubtedly disburses the majority of social security benefits to low-income households. Nevertheless, richer households are those that pay the largest share of federal income taxes: when they are hit by idiosyncratic negative income shocks, they suddenly start to pay less taxes.

Figure 4.4 reports risk sharing estimates separately by taxable income quartiles. We

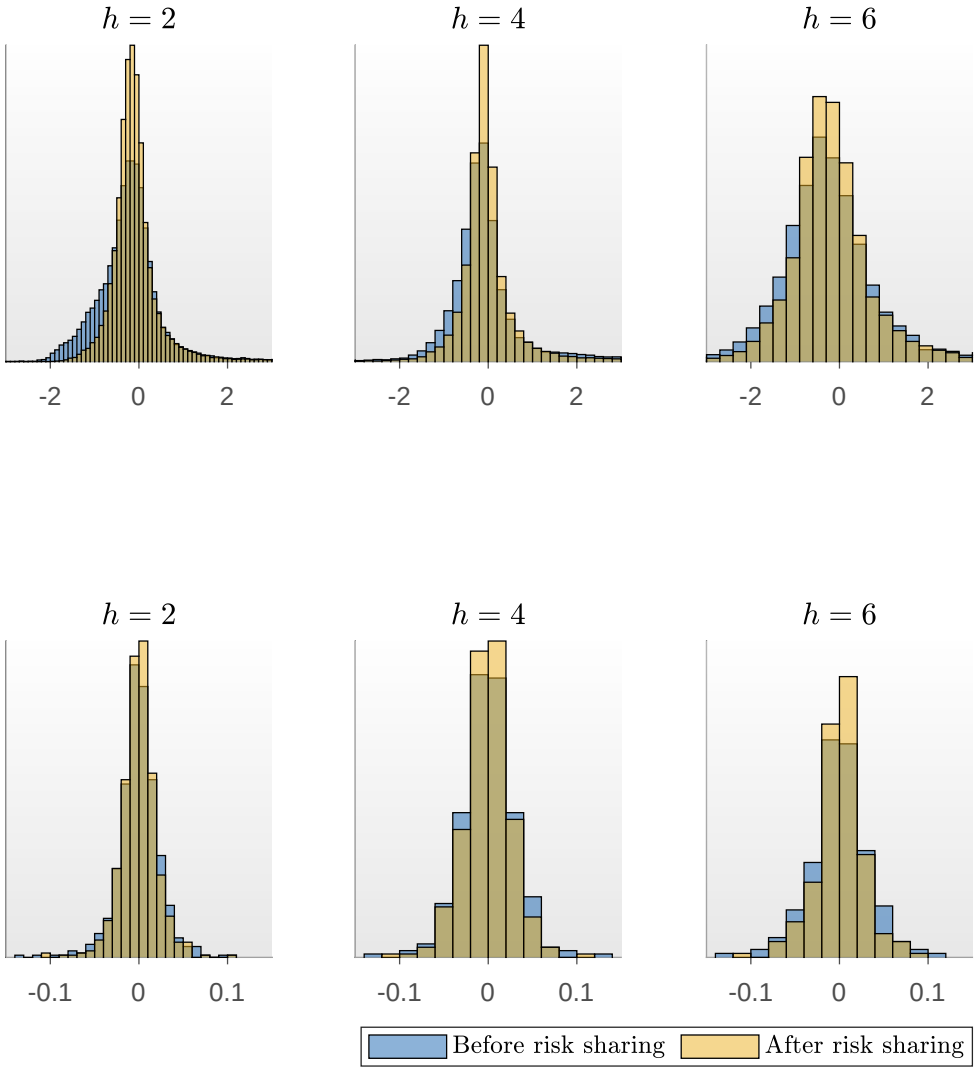


Figure 4.2: Histograms of residuals. The top panels display histograms of residuals estimated from the model using household-level data. Each panel corresponds to a different forecasting horizon. Light blue histograms represent residuals from the model using pre-government-intervention income, while yellow histograms correspond to residuals from the model using post-government-intervention income. The bottom panels present histograms of residuals estimated from the model using state-level data from the BEA.

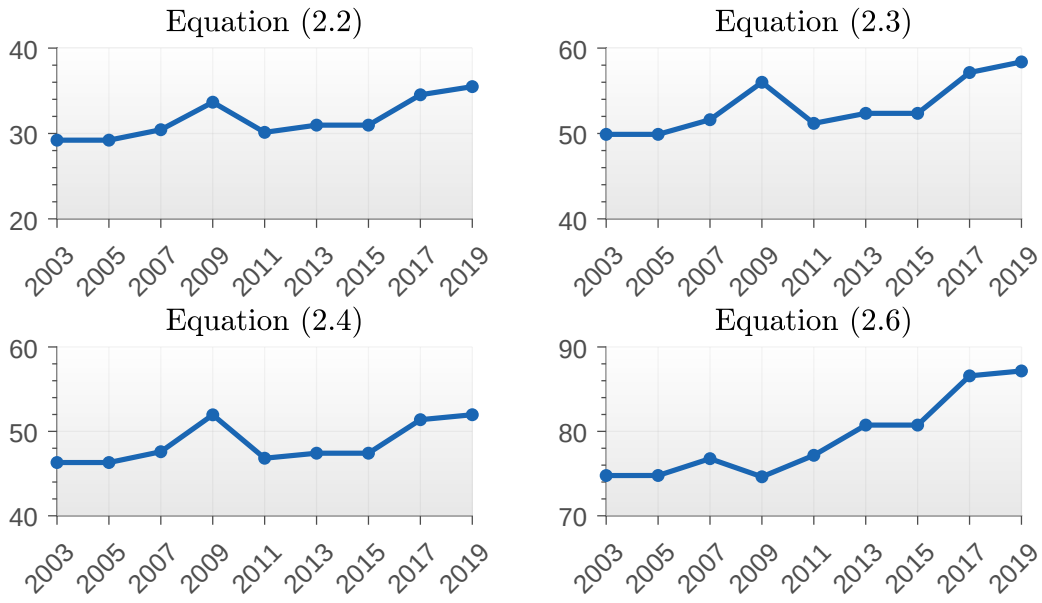


Figure 4.3: Rolling windows. Each panel shows risk sharing estimates obtained through different metrics explained in Section 2. The time series are obtained through eight-year rolling regressions.

find clear evidence that lower-income households benefit disproportionately more than higher-income ones: for those in the first quartile, the federal government cushions more than the loss they suffer during the first four years, and 78.9% after six years. On the other hand, households in the fourth quartile enjoy lower insurance, especially at longer horizons. After two years they are still able to hedge a sizable 57.8% downside risk. Nevertheless, in the fourth and sixth year this share drops to 23.2% and 21.4% respectively. More generally, our estimates are monotonically decreasing with quartiles.

5 Concluding remarks

In this paper, we develop a methodology that allows to estimate downside-risk sharing. Building on the framework in Athanasoulis and Van Wincoop (2001) and adopting a simple definition of (stocks') downside risk from Markowitz (1959), we adapt the latter to our setting to measure downside-risk sharing. We apply our methodology to PSID data to gauge the extent to which the federal government of the United States is able to insure

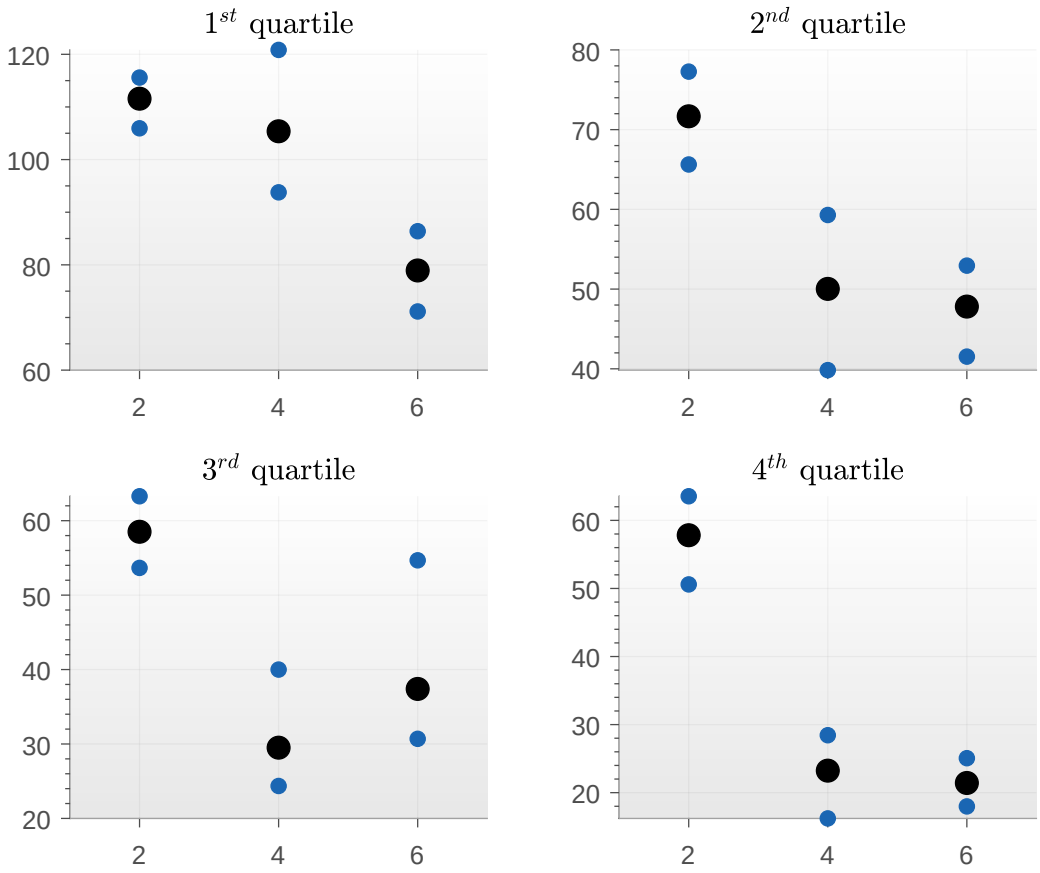


Figure 4.4: Downside-risk sharing by income quartiles. The figure shows downside-risk sharing estimates obtained by running regressions separately for each quartile of the taxable income distribution available in the PSID. The x-axis represents the horizon (in years); the y-axis represents the risk sharing estimate (in percent). Black circles are our baseline estimates, taken as the median of the distribution; blue circles are 95% confidence intervals obtained with the bootstrap procedure described in Kapetanios (2008).

households against idiosyncratic income uncertainty, and we re-evaluate the importance of the fiscal channel by showing that previous estimates vastly understate its capabilities. This is primarily due to two factors. First, using state-level data - as done in most papers - leads to a downward bias which is driven by the implicit reliance on a state-specific representative household onto which all income concentration is allocated. This masks key underlying inequalities and heterogeneities. Second, computing risk sharing by implicitly using both tails of the income growth distribution also compresses risk sharing estimates. Indeed, we show that virtually all the risk reduction achieved through federal transfers and taxes results from lower volatility of the left tail, while the right tail remains remarkably stable. Interestingly, this phenomenon is not visible when estimating our model with state-level data.

Our work challenges the foundations of many policy debates that - based on past research - assert that public risk sharing is significantly weaker than private risk sharing. In fact, not only the fiscal risk-sharing channel plays an increasingly important stabilizing role for US households' incomes, but is even more pronounced for lower-income individuals.

A Appendix - weighted estimates

Here, we show that estimating our model with weighted regressions requires a slight adjustment to the non-weighted formulas, and therefore first recall some relevant basic facts. For ease of notation, define $y^i_{t,h} = g^i_{t,t+h} - g^{US}_{t,t+h}$, $\mathbf{z}^i_{t,h} = \mathbf{x}^i_{t,h} - \mathbf{x}^{US}_{t,h}$, and $u^i_{t,h} = \epsilon^i_{t,t+h}$. We can therefore rewrite equation (2.1) as:

$$\mathbf{y}_h = \mathbf{Z}_h \boldsymbol{\beta}_h + \mathbf{u}_h, \tag{A.1}$$

where we have stacked all $j = 1, \ldots, S_h \cdot N$ observations and used matrix notation. We then define $\mathbf{W}_h$ to be our (diagonal) weighting matrix, with diagonal elements $w_{jj,h}$ normalized so as their sum equals the total number of observations. We also define $\mathbf{P}_h = \text{chol}(\mathbf{W}_h)$ the matrix obtained by applying the Cholesky decomposition on $\mathbf{W}_h$, which in our case simply yields a diagonal matrix with each element in $\mathbf{P}_h$ being equal to the square root of the corresponding ones in $\mathbf{W}_h$. The weighted version of our model is then performed by estimating the following transformed regression:

$$\mathbf{y}^*_h = \mathbf{Z}^*_h \boldsymbol{\beta}_h + \mathbf{u}^*_h, \tag{A.2}$$

where $\mathbf{y}^*_h = P_h \mathbf{y}_h$, $\mathbf{Z}^*_h = P_h \mathbf{Z}_h$, and $\mathbf{u}^*_h = P_h \mathbf{u}_h$. Morever, define $\hat{\mathbf{u}}_h = \mathbf{y}_h - \mathbf{Z}_h \hat{\boldsymbol{\beta}}^*_h$, with $\hat{\boldsymbol{\beta}}^*_h$ being the weighted estimator.

Because of the fact that also the intercept gets weighted in the transformed regression, only *weighted* residuals $\mathbf{W}_h \hat{\mathbf{u}}_h = \mathbf{P}_h \hat{\mathbf{u}}^*_h$ have a mean equal to zero. This implies that computing conditional residuals as in equation (2.5) or even with the general threshold in equation (2.7) could be misleading.

Therefore, our baseline specification modifies the above equations and define condi-

tional forecast errors in the following way:

$$\hat{u}_{j,h}^{*,c} = \begin{cases} \hat{u}_{j,h}^{*} & \text{if} \quad w_{jj,h} \cdot \hat{u}_{j,h} < c, \\ 0 & \text{otherwise,} \end{cases} \tag{A.3}$$

where c is again a general threshold. Finally, we compute residual semi-variances as

$$\hat{\sigma}_{h}^{2,*,\cdot,c} = \frac{(\hat{\mathbf{u}}_{h}^{*,c})' \hat{\mathbf{u}}_{h}^{*,c}}{S_h \cdot N - K - 1}. \tag{A.4}$$

B Appendix - federal transfers

The main source we use to obtain federal transfers data is the BEA. More specifically, we sum up retirement and disability insurance benefits, Medicare benefits, military medical insurance benefits, income maintenance benefits, veterans' benefits (excluding the category "other assistance to veterans"), education and training assistance, other transfer receipts of individuals from governments, current transfer receipts of nonprofit institutions from the federal government.

Public assistance medical care benefits are a large component and almost exclusively consist of Medicaid and the Children's Health Insurance Program (CHIP), which are administered by states and jointly funded by the federal government and the states. The amount of federal funding for Medicaid is determined by the Federal Medical Assistance Percentage (FMAP), which is based on a formula that provides higher reimbursement to states with lower per capita incomes relative to the national average (and viceversa). FMAPs have been defined in the Social Security Act Amendments of 1965, which created Medicaid by adding Title XIX to the Social Security Act. They grant a statutory minimum of 50 percent and a maximum of 83 percent. The amount of federal funding for CHIP is determined by the enhanced FMAPs (E-FMAP), which are higher than standard FMAPs, although in this case the federal funding share cannot exceed 85 percent. Therefore, we

adopt a conservative approach and include only 50% of public assistance medical care benefits.

Unemployment insurance compensation also is jointly funded by federal and state governments, but the BEA only provides aggregate benefits figures. For this reason, many studies simply include them in their entirety, thereby inducing an upward bias in the actual amount of *federal* unemployment insurance. As pointed out in Alcidi and Thirion (2017), federal funding for unemployment compensation is very low during normal times. Nevertheless, it increases both during high state-specific unemployment spells through the permanent Extended Benefits program, and through discretionary federal Emergency Unemployment Compensation schemes that have historically been approved when large symmetric shocks hit the US economy.

In general, it is therefore clear that using BEA's all-government unemployment benefits is unsatisfactory, and we thus depart from all previous contributions by using US Department of Labor state-level data for the federal share of unemployment insurance benefits paid at all levels, namely regular, extended, and discretionary ones.[17]

[17] Past discretionary programmes have been funded by the federal government only.

References

Adrian, T., Boyarchenko, N., Giannone, D., 2019. Vulnerable growth. American Economic Review 109, 1263–1289.

Adrian, T., Grinberg, F., Liang, N., Malik, S., Yu, J., 2022. The term structure of growth-at-risk. American Economic Journal: Macroeconomics 14, 283–323.

Alcidi, C., D'Imperio, P., Thirion, G., 2023. Risk-sharing and consumption-smoothing patterns in the US and the Euro Area: a comprehensive comparison. Structural Change and Economic Dynamics 64, 58–69.

Alcidi, C., Thirion, G., 2017. Fiscal risk sharing and resilience to shocks: lessons for the Euro Area from the US. CEPS Working Document 7.

Ang, A., Chen, J., Xing, Y., 2006. Downside risk. The Review of Financial Studies 19, 1191–1239.

Anzuini, R., Rossi, L., 2021. Fiscal policy in the US: a new measure of uncertainty and its effects on the American economy. Empirical Economics 61, 2613–2634.

Anzuini, R., Rossi, L., Tommasino, P., 2020. Fiscal policy uncertainty and the business cycle: time series evidence from Italy. Journal of Macroeconomics 65.

Asdrubali, P., Kim, S., 2004. Dynamic risk sharing in the United States and Europe. Journal of Monetary Economics 51, 809–836.

Asdrubali, P., Kim, S., Pericoli, F., Poncela, P., 2023. Risk sharing channels in OECD countries: a heterogeneous panel VAR approach. Journal of International Money and Finance 131.

Asdrubali, P., Sørensen, B.E., Yosha, O., 1996. Channels of interstate risk sharing: United States 1963-1990. The Quarterly Journal of Economics 111, 1081–1110.

Asdrubali, P., Tedeschi, S., Ventura, L., 2020. Household risk-sharing channels. Quantitative Economics 11, 1759–7331.

Athanasoulis, S.G., Van Wincoop, E., 2000. Growth uncertainty and risk sharing. Journal of Monetary Economics 45, 477–505.

Athanasoulis, S.G., Van Wincoop, E., 2001. Risk sharing within the United States: what do financial markets and fiscal federalism accomplish? The Review of Economics and Statistics 83, 688–698.

Baker, S., Bloom, N., Davis, S., 2016. Measuring economic policy uncertainty. The Quarterly Journal of Economics 131, 1593–1636.

Bayoumi, T., Masson, P.R., 1995. Fiscal flows in the United States and Canada: lessons for monetary union in Europe. European Economic Review 39, 253–274.

Bernanke, B., 1983. Irreversibility, Uncertainty, and Cyclical Investment. The Quarterly Journal of Economics 98, 85–106.

Bloom, N., 2009. The impact of uncertainty shocks. Econometrica 77, 623–685.

Bloom, N., 2014. Fluctuations in Uncertainty. Journal of Economic Perspectives 28, 153–176.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., Terry, S., 2018. Really uncertain business cycles. Econometrica 86, 1031–1065.

Brownlees, C., Souza, A.B., 2021. Backtesting global growth-at-risk. Journal of Monetary Economics 118, 312–330.

Burriel, P., Chronis, P., Freier, M., Hauptmeier, S., Reiss, L., Stegarescu, D., Van Parys, S., 2020. A fiscal capacity for the Euro Area: lessons from existing fiscal-federal systems. Banco de España Occasional Paper 2009.

Carriero, A., Clark, T., Marcellino, M., 2018. Measuring uncertainty and its impact on the economy. The Review of Economics and Statistics 100, 799–815.

Carriero, A., Clark, T.E., Marcellino, M., 2024. Capturing macroeconomic tail risks with Bayesian Vector Autoregressions. Journal of Money, Credit and Banking 56, 1099–1127.

Christiano, M., Motto, R., Rostagno, M., 2014. Risk Shocks. American Economic Review 104, 27–65.

Cimadomo, J., Ciminelli, G., Furtuna, O., Giuliodori, M., 2020. Private and public risk sharing in the euro area. European Economic Review 121.

Cimadomo, J., Giuliodori, M., Lengyel, A., Mumtaz, H., 2023. Changing patterns of risk-sharing channels in the United States and the euro area. European Central Bank Working Papers .

Cochrane, J.H., 1991. A simple test of consumption insurance. Journal of Political Economy 99, 957–976.

Delle Monache, D., De Polis, A., Petrella, I., 2024. Modeling and forecasting macroeconomic downside risk. Journal of Business and Economic Statistics 42, 1010–1025.

Dixit, A., Pindyck, R., 1994. Investment under uncertainty. Princeton: Princeton University Press .

Dynarski, S., Gruber, J., 1997. Can families smooth variable earnings? Brookings Papers on Economic Activity 28, 229–303.

Eberly, J., 1994. Adjustment of consumers' durables stocks: evidence from automobile purchases. Journal of Political Economy 102, 403–436.

Estrada, J., 2006. Downside risk in practice. Journal of Applied Corporate Finance 18, 117–125.

Farago, A., Tédongap, R., 2018. Downside risks and the cross-section of asset returns. Journal of Financial Economics 129, 69–86.

Farhi, E., Werning, I., 2017. Fiscal unions. American Economic Review 107, 3788–3834.

Fernández-Villaverde, J., Guerrón-Quintana, P., Kuester, K., Rubio-Ramírez, J., 2015. Fiscal volatility shocks and economic activity. American Economic Review 105, 3352–3384.

Ferrari, A., Rogantini Picco, A., 2023. Risk sharing and the adoption of the euro. Journal of International Economics 141.

Giglio, S., Kelly, B., Pruitt, S., 2016. Systemic risk and the macroeconomy: an empirical evaluation. Journal of Financial Economics 119, 457–471.

Gouskova, E., Heeringa, S., McGonagle, K., Schoeni, R., Stafford, F., 2008. Panel Study of Income Dynamics revised longitudinal weights 1993-2005. Survey Research Center - Institue for Social Research, University of Michigan Technical Series Paper .

Gul, F., 1991. A theory of disappointment aversion. Econometrica 59, 667–686.

Hayashi, F., Altonji, J., Kotlikoff, L., 1996. Risk-sharing between and within families. Econometrica 64, 261–294.

Hepp, R., Von Hagen, J., 2012. Fiscal federalism in Germany: stabilization and redistribution before and after unification. Publius: The Journal of Federalism 42, 234–259.

Hepp, R., Von Hagen, J., 2013. Interstate risk sharing in Germany: 1970-2006. Oxford Economic Papers 65, 1–24.

Husted, L., Rogers, J., Sun, B., 2020. Monetary policy uncertainty. Journal of Monetary Economics 115, 20–36.

Jurado, K., Ludvigson, S.C., Ng, S., 2015. Measuring uncertainty. American Economic Review 105, 1177–1216.

Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. Econometrica 47, 263–292.

Kapetanios, G., 2008. A bootstrap procedure for panel data sets with many cross sectional units. Econometrics Journal 11, 377–395.

Kenen, P., 1969. The theory of optimum currency areas: an eclectic view. In *Monetary Problems of the International Economy*, edited by Robert A. Mundell and Alexander K. Swoboda. Chicago: Chicago University Press .

Lettau, M., Maggiori, M., Weber, M., 2014. Conditional risk premia in currency markets and other asset classes. Journal of Financial Economics 114, 197–225.

Mace, B.J., 1991. Full insurance in the presence of aggregate uncertainty. Journal of Political Economy 99, 928–956.

Malkin, I., Wilson, D.J., 2013. Taxes, transfers, and State economic differences. Economic Letter, Federal Reserve Bank of San Francisco 36.

Markowitz, H., 1959. Portfolio selection: efficient diversification of investments. Yale University Press, New Haven and London .

Mélitz, J., Zumer, F., 1999. Interregional and international risk-sharing and lessons for EMU. Carnegie-Rochester Conference Series on Public Policy 51, 149–188.

Milano, V., 2017. Risk sharing in the eurozone: the role of European institutions. CeLEG Working Papers Series - LUISS Guido Carli .

Nelson, J.A., 1994. On testing for full insurance using Consumer Expenditure Survey data. Journal of Political Economy 102, 384–394.

Obstfeld, M., Peri, G., 1998. Regional non-adjustment and regional shocks. Economic Policy 13, 205–259.

Plagborg-Møller, M., Reichlin, L., Ricco, G., Hasenzagl, T., 2020. When is growth at risk? Brookings Papers on Economic Activity , 167–229.

Poghosyan, T., Senhadji, A., Cottarelli, C., 2016. The role of fiscal transfers in smoothing regional shocks: evidence from existing federations. IMF Working Papers .

Romer, C., 1990. The great crash and the onset of the great depression. The Quarterly Journal of Economics 105, 597–624.

Rossi, L., 2013. Real options valuations in oligopolistic markets. MSc Thesis, Università Cattolica del Sacro Cuore .

Routledge, B., Zin, S., 2010. Generalized disappointment aversion and asset prices. The Journal of Finance 65, 1303–1332.

Roy, A., 1952. Safety first and the holding of assets. Econometrica 20, 431–449.

Sala-i-Martin, X., Sachs, J., 1992. Fiscal policies and optimum currency areas: evidence for Europe from the United States. In Matthew Canzoneri, Vittorio Grilli, and Paul Masson (eds.), Establishing a Central Bank: Issues in Europe and Lessons from the United States. Cambridge: Cambridge University Press , 195–219.

Segal, G., Shaliastovich, I., Yaron, A., 2015. Good and bad uncertainty: macroeconomic and financial market implications. Journal of Financial Economics 117, 369–397.

Sørensen, B.E., Yosha, O., 1998. International risk sharing and European monetary unification. Journal of International Economics 45, 211–238.

Stempel, D., 2021. Risk sharing heterogeneity in the United States. Economics Bulletin, Heinrich Heine University Düsseldorf 41, 1223–1240.

Tibshirani, R., 1996. Regression shrinkage and selection via the Lasso. Journal of the Royal Statistical Society. Series B (Methodological) 58, 267–288.

Von Hagen, J., 1992. Fiscal arrangements in a monetary union: evidence from the US. In D. Fair and C. de Boissieux (eds), Fiscal Policy, Taxes, and the Financial System in an Increasingly Integrated Europe, Kluwer , 337–359.