The Role of Risk Sharing in Attenuating Business Cycles Within Currency Unions

Alberto Pavia*[1] and Christian Proebsting †[1]

[1]KU Leuven

March 20, 2025

Abstract

How does risk sharing across member states lower the cost of being in a currency union? We decompose the response of risk sharing to state-specific income fluctuations into two components: the direct response of consumption to income changes and the indirect feedback of the shock to income via general equilibrium effects. Risk sharing not only directly smooths consumption but also reduces the volatility of income by stabilizing demand. We estimate both direct and indirect effects by exploiting regional variation in military buildups across U.S. states. A $1 increase in external demand raises state-level income by $1.43, with 47% of this increase passed through to consumption. We rationalize these findings using a small open economy model that incorporates multiple risk-sharing channels. The model allows us to perform policy counterfactuals, showing that risk sharing between U.S. states reduces state-level consumption volatility by a factor of 3.5, with indirect effects accounting for 40% of this reduction. These results imply that the benefits of risk sharing are 40% larger than previously estimated.

JEL Codes: F44, F45, E32, F15, F41, F36.

Keywords: risk sharing, currency unions, consumption volatility.

*alberto.paviasoto@kuleuven.be, Naamsestraat 69, 3000 Leuven, Belgium.

†christian.probsting@kuleuven.be, Naamsestraat 69, 3000 Leuven, Belgium.

We would like to thank Chris Boehm, Christophe Blot, Mathias Hoffmann, Raphaël Huleux, Gernot Müller, Fabrizio Perri, Anna Rogantini Pico, Bent Sørensen, Burak Uras, as well as numerous seminar and conference participants for useful comments. We gratefully acknowledge funding from the Research Foundation – Flanders (FWO), project no. G040923N and fellowship no. 1132725N.

1 Introduction

When countries enter a currency union, they forgo their ability to conduct independent monetary policy, leaving them vulnerable to country-specific economic shocks. While a common monetary policy can stabilize union-wide business cycles, idiosyncratic fluctuations in output pose significant risks to the stability of currency unions, whose success therefore hinges on the availability of alternative stabilization mechanisms, as emphasized in the Optimum Currency Area literature (Mundell, 1961).

The United States is often regarded as a well functioning currency union, where countercyclical fiscal transfers, income diversification, high labor mobility, and integrated credit markets mitigate the impact of state-specific shocks on consumption. In contrast, the euro area has faced persistent challenges due to large differences in business cycles among its member states, amplifying the costs of the monetary union. Figure 1 illustrates this by showing the differences in the growth rates of GDP and consumption between U.S. states and euro area countries, respectively. In the euro area, per-capita consumption is almost as volatile as aggregate output, whereas U.S. consumption volatility is remarkably lower—less than 40% of output volatility—suggesting a higher degree of risk sharing.

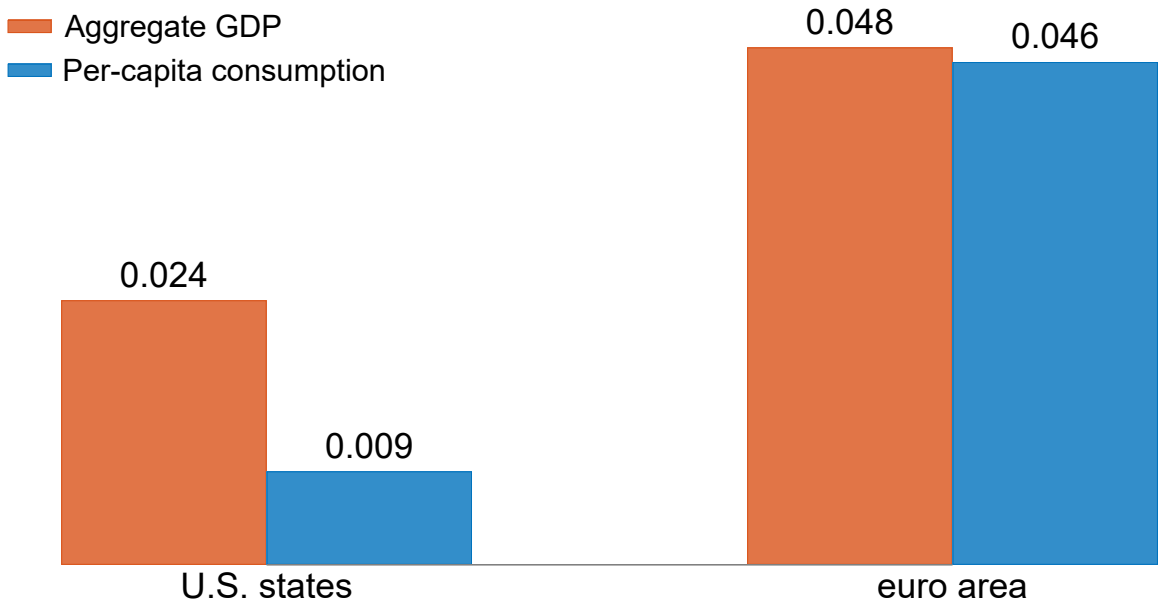


Figure 1: CROSS-SECTIONAL STANDARD DEVIATION OF GROWTH RATES

Note: The bars display the cross-sectional standard deviations of nominal per-capita consumption and aggregate nominal GDP for U.S. states and euro area countries, averaged over 1999-2019.

This observation aligns with a long strand of literature, pioneered by Asdrubali, Sørensen and

Yosha (1996) (hereafter ASY), which decomposes empirically the pass-through of output fluctuations to consumption fluctuations using a variance decomposition framework (Sørensen and Yosha, 1998; Kalemli-Ozcan, Sørensen and Yosha, 2003; Balli, Sørensen et al., 2006). Recent estimates by Kohler, Müller and Wellmann (2023) indicate that risk-sharing mechanisms smooth approximately 80% of income fluctuations in the U.S., compared to just 40% in the euro area.

However, while informative, the imperative empirical risk-sharing framework cannot answer a central question in the optimum currency area debate—and the key question of this paper: *To what extent does the lack of risk sharing amplify consumption volatility, thereby increasing the cost of currency union membership?* This is because the ASY framework treats output fluctuations as given, focusing on how much of an income shock is smoothed. However, risk-sharing may not only reduce the volatility of consumption directly but also stabilize the business cycle itself. By stabilizing demand, risk-sharing mechanisms—such as fiscal transfers or credit markets—may dampen the magnitude of output fluctuations. Figure 1 illustrates this point, showing that not only is consumption less volatile than output in the United States, but U.S. state-level output fluctuations are also half as large as those in the euro area (2.4% vs. 4.8%).

In this paper, we argue that the lower output volatility observed across U.S. states is, at least in part, a consequence of smoother consumption. Risk sharing thus reduces consumption volatility through two channels: (1) a **direct effect**, smoothing consumption taking output fluctuations as given, and (2) an **indirect general equilibrium effect**, dampening output fluctuations via stabilized demand. For example, fiscal transfers to a recession-hit region boost local demand, mitigating the downturn's severity. This indirect effect of risk-sharing—through its impact on output volatility—has been largely overlooked in the literature. Our analysis quantifies the magnitude of this indirect effect and its implications for the cost of currency unions.

To formalize this argument, we first introduce a standard small open economy model. We show that the response of consumption to an external demand shock can be decomposed into two components: a direct effect, which captures the elasticity of consumption to income, and an indirect effect, which reflects the response of income to the shock itself. The direct effect has been the primary focus of the risk-sharing literature and, under certain conditions, can be estimated using the variance decomposition framework proposed by ASY. The indirect effect, by contrast, corresponds to the open-economy relative multiplier studied by Nakamura and Steinsson (2014).

While these two components have been extensively analyzed in separate strands of the literature, we establish a tight link between them. In particular, we show that economies with greater risk-sharing exhibit lower multipliers: when consumption is stabilized, demand fluctuations are dampened, reducing the volatility of output. This finding suggests that the total benefits of risk-sharing extend beyond the direct income-smoothing effects emphasized in prior studies, encom-

passing broader macroeconomic stabilization.

Next, we empirically estimate both the direct and indirect effects of risk-sharing by exploiting regional variation in military buildups across U.S. states, following the empirical framework of Nakamura and Steinsson (2014). Unlike the traditional ASY approach, which relies on an unconditional variance decomposition, our strategy identifies how risk-sharing operates in response to specific, external shocks. Our empirical estimates indicate that a \$1 increase in external demand raises a state's output by \$1.43, with 53% of the shock being smoothed through risk-sharing mechanisms, leading to a \$0.70 increase in consumption. As in the risk-sharing literature, we further decompose the extent of risk-sharing across different channels: migration and fiscal transfers each account for 12% and 14% of risk absorption, capital market diversification contributes almost a quarter, and credit markets provide a more moderate 4%.

We use these identified responses to discipline the key risk-sharing parameters in a structural, quantitative model that extends our simple model by explicitly incorporating the various risk-sharing channels. The quantitative model also matches well the open-economy relative multiplier, that is, it captures both the direct and indirect effects of risk sharing. To evaluate the importance of risk-sharing across U.S. states, our main counterfactual exercise shuts down all risk-sharing. This raises consumption volatility by 3.25, around 40% of which is driven by an increase in the volatility of output. In other words, 40% of the benefits of U.S. interstate risk-sharing stem from indirect, general equilibrium effects.

In a second set of counterfactuals, we shut down one risk sharing mechanism at a time. This reveals that the risk-sharing channels generally are substitutes to each other: those risk-sharing channels that remain active partially compensate for the loss in risk sharing. For instance, while shutting down fiscal transfers across states directly raises states' consumption volatility all else being equal, this increase in consumption volatility raises the benefits from migration. The ensuing increase in migration partially dampens the direct rise in consumption volatility.

After reviewing the literature, we propose a simple SOE model to separate the direct and indirect effects of risk sharing and connect them to the risk-sharing coefficients and the open-economy multiplier. Section 3 estimates these direct and indirect effects of risk sharing by exploiting regional variation in military buildups across U.S. states. These estimates discipline our quantitative model that we present in Section 4 and which we then use to quantify the benefits of the observed level of interstate risk sharing. The last section concludes.

1.1 Literature

This paper links two distinct literatures: an older literature that examines empirical measures and tests of risk sharing, and a more recent literature that exploits regional variation to estimate open-economy multipliers. Combining the two literatures has important implications for the literature on optimum currency areas.

Since the seminal work of Cochrane (1991), Mace (1991) and Townsend (1994), a vast literature testing risk sharing at the state and country level has developed, as exemplified by ASY, Sørensen and Yosha (1998); Becker and Hoffmann (2006); Hoffmann (2008); Flood, Marion and Matsumoto (2012) and Kohler, Müller and Wellmann (2023). Some of these papers highlight the lack of improved risk sharing over time despite increased financial globalization, while others contrast the relatively high level of risk sharing across U.S. states with the much lower level observed across euro area countries, thereby connecting to an older literature on optimal currency areas from the 1950s and 1960s (Friedman, 1953; Mundell, 1961). A defining feature of this literature is its structural agnosticism, as it builds on a pure accounting framework. In contrast, our focus is to provide a structural interpretation of risk-sharing estimates, which allows us to quantify the contribution of risk-sharing in reducing consumption volatility through counterfactual analyses.

While several authors propose structural models for specific risk-sharing channels, they do so without connecting their work to the reduced-form risk-sharing regressions (see e.g. House, Proebsting and Tesar (2023) for labor mobility and Evers (2015); Pennings (2016); Beraja (2022) for fiscal transfers). These studies provide a detailed picture of a single risk-sharing mechanism. We instead incorporate all commonly discussed risk-sharing channels and highlight their interdependencies, emphasizing that the relevance of a specific channel depends on the size of the others.

Our approach also connects to a second, more recent literature on cross-regional spending multipliers (Nakamura and Steinsson, 2014; Chodorow-Reich, 2019). Several papers have noted that these multiplier estimates are informative about local general equilibrium effects and can be used to estimate either micro parameters (Guren et al., 2021) or, in conjunction with aggregate fiscal multiplier estimates, to assess the aggregate effects of generic consumption demand shifters (Wolf, 2023). In the spirit of Guren et al. (2021), we show that, under certain conditions, the indirect effects of risk sharing are related to the size of the cross-regional spending multiplier.

More broadly, our research relates to the literature on optimum currency areas, dating back to Friedman (1953). This literature established a set of "criteria" for monetary integration, such as factor market mobility (Mundell, 1961), the degree of economic openness (McKinnon, 1963), fiscal integration (Kenen, 1969), and financial integration through credit markets (Ingram, 1962) and equity markets (Mundell, 1973). We show that the quantitative benefits of these risk-sharing mech-

anisms include not only direct effects but also indirect benefits through lower per-capita output volatility, akin to the effects of a flexible exchange rate. Additionally, we highlight the interactions among these criteria, as the indirect benefits depend on countries' trade openness. For example, risk-sharing through fiscal transfers is particularly potent in economies that are less open to trade.

2 Risk sharing and open-economy multipliers

Consider a small open economy that faces a temporary, and unexpected increase in external demand for its products, denoted by $\Delta G_t^{nom} \equiv G_t^{nom} - G_{t-1}^{nom}$. Conceptually, we can break down the response of nominal consumption into two parts: the elasticity of nominal consumption to nominal output, $\frac{\hat{C}_t^{nom}}{\hat{Y}_t^{nom}}$ (where a hat denotes the percent change of variable relative to the previous period), and the increase in output, ΔY_t^{nom}. That is, the percent change in nominal consumption for an increase in external demand corresponding to 1% of $t-1$ output is given by

$$\frac{\hat{C}_t^{nom}}{\frac{\Delta G_t^{nom}}{Y_{t-1}^{nom}}} = \frac{\hat{C}_t^{nom}}{\hat{Y}_t^{nom}} \times \frac{\Delta Y_t^{nom}}{\Delta G_t^{nom}}.$$

Both of these terms have received substantial attention in the literature. The first term, $\frac{\hat{C}_t^{nom}}{\hat{Y}_t^{nom}}$, has been the focus of the empirical international risk-sharing literature. ASY estimated the elasticity of nominal consumption to nominal output by regressing log changes in nominal consumption on log changes in nominal output across U.S. states. We follow the notation of this literature and denote this elasticity by β_t^U because it captures the share of risk that is not absorbed by any risk-sharing channel and that remains unsmoothed. Assuming that households have log utility, perfect risk sharing would imply equal growth rates in per-capita nominal consumption across the small open economy and the rest of the world (Backus and Smith, 1993). This would imply that the growth rate of the country's nominal consumption is independent of the growth rate of its nominal output: $\beta_t^U = 0$. With imperfect risk sharing $\beta_t^U > 0$ and, in an extreme case of $\beta_t^U = 1$, fluctuations in income would translate one-for-one into fluctuations in consumption.

The second term $\frac{\Delta Y_t^{nom}}{\Delta G_t^{nom}}$ is the open-economy relative multiplier. This follows the terminology of Nakamura and Steinsson (2014) who estimate this relative multiplier by exploiting variation in external demand (federal government spending) across U.S. states. As in their study, we consider the response of nominal output rather than real output because subnational price deflators for output are not available for the United States. We denote this term by m_t.

Taken together, we can therefore write the consumption multiplier as

$$m_t^c \equiv \frac{\hat{C}_t^{nom}}{\frac{\Delta G_t^{nom}}{Y_{t-1}^{nom}}} = \beta_t^U \times m_t. \tag{2.1}$$

The multiplier indicates the change in income for a shock to external demand corresponding to 1 percent of output. The risk-sharing β determines how a change in income translates into a change in consumption.

The literature has treated these two objects separately, but in the following, we show that the multiplier and the risk-sharing β are interrelated. The multiplier is increasing in β_t^U. Risk sharing, a reduction in β_t^U, has a double benefit: It reduces the response of consumption for a given increase in income, but it also reduces the response of output to the increase in external demand. We next develop this argument more formally in a standard open-economy setup.

2.1 A simple open-economy model

We consider a continuum of small open economies that are part of a currency union (Gali and Monacelli, 2005). The individual economies are subject to idiosyncratic fluctuations in government demand for their goods. Government spending at the level of the currency union, however, is constant. We formally derive a link between the open-economy relative multiplier, m_t, and the risk-sharing parameter β.

In our framework, this link is independent of several model features that are known to influence the closed-economy multiplier, such as the degree of nominal price or wage rigidity, the small open economy's monetary policy, whether the production function features decreasing returns to scale, whether input markets are segmented and the elasticity of labor supply. Instead, the multiplier depends on the economy's openness in both financial and goods markets. We therefore discuss these two features of the model more thoroughly and present the remaining model ingredients along with with the quantitative model.

2.1.1 Trade in financial assets

In this simple model, we focus on a single, generic risk sharing mechanism–an assumption we relax in the quantitative model, where we extend the analysis to multiple risk-sharing channels that we can map more easily to the data. We assume that households can trade state-contingent bonds, but these trades are subject to a tax. This distortion–first introduced by Devereux and Yetman (2014)–creates a wedge between the returns to bonds for households living in the small open economy and those in the rest of the union.

More formally, let $B(s^t, s_{t+1})$ be the quantity of state-contingent bonds purchased by the household after history s^t that pay off one unit of the country's currency in state s^{t+1}.[1] The price of these bonds to the household in the small open economy is $(1 + t(s^t, s_{t+1}))Q(s^t, s_{t+1})$, where $Q(s^t, s_{t+1})$ is the net price and $t(s^t, s_{t+1})$ is a tax levied by the government in the rest of the world. The tax rate depends on the state of the economy and the bond to which it applies. For it to be interpreted as a tax, we require that for bond purchases, $t(s^t, s_{t+1}) \geq 0$ (thereby raising the purchase price), whereas for bond sales, $t(s^t, s_{t+1}) \leq 0$ (thereby lowering the income from the sale).

In addition to income from bonds, households receive the nominal income generated from the production of output, $P^Y(s^t)Y(s^t)$, where $P^Y(s^t)$ is the nominal price of the economy's output, and they purchase consumption goods, $C(s^t)$, at price $P(s^t)$. Notice that we are agnostic about how output is produced, but for now assume that any capital stock is fixed and does not require investment. The household's budget constraint is then

$$P(s^t)C(s^t) + \sum_{s^{t+1}} (1 + t(s^t, s_{t+1}))Q(s^t, s_{t+1})B(s^t, s_{t+1}) = P^Y(s^t)Y(s^t) + B(s^{t-1}, s_t).$$

With log utility, the optimal choice of bonds implies the following international risk-sharing condition:

$$\frac{P(s^t)C(s^t)}{P(s^{t-1})C(s^{t-1})}(1 + t(s^{t-1}, s_t)) = \frac{P^*(s^t)C^*(s^t)}{P^*(s^{t-1})C^*(s^{t-1})} \tag{2.2}$$

Following Devereux and Yetman (2014), we assume that the tax rate for bond $B(s^{t-1}, s_t)$ is increasing in the ratio of consumption growth to output growth:

$$1 + t(s^{t-1}, s_t) = \left(\frac{P(s^t)C(s^t)}{P(s^{t-1})C(s^{t-1})} \frac{P^Y(s^{t-1})Y(s^{t-1})}{P^Y(s^t)Y(s^t)} \right)^{\frac{\lambda}{1-\lambda}},$$

where $\lambda \in [0, 1]$ governs the degree of curvature of the tax function. Hence, $t(s^{t-1}, s_t) \geq 0$ if nominal consumption growth exceeds nominal output growth (and households purchase bonds to insure against these "bad" states), and $t(s^{t-1}, s_t) \leq 0$ if nominal output growth exceeds nominal consumption growth (and households sell bonds that pay off in these "good" states). Inserting this into the

[1] We follow the standard notation for models that explicitly account for state-contingent bonds: Each period t, the economy experiences one of finitely many events $s \in S$. Let the transition probability from state s to s' follow a Markov chain denoted by $\pi(s'|s)$. $s^t = (s_0, s_1, ..., s_t)$ represent the history of events up through and including period t. The probability, as of period 0, of any history s^t is $\pi(s^t)$. We denote by $X(s^t)$ the realization of variable X in state s^t. Unless confusion arises, we write X_t for $X(s^t)$.

international risk sharing condition (2.2) we get

$$\frac{P_t C_t}{P_{t-1} C_{t-1}} = \left(\frac{P_t^Y Y_t}{P_{t-1}^Y Y_{t-1}} \right)^{\lambda} \left(\frac{P_t^* C_t^*}{P_{t-1}^* C_{t-1}^*} \right)^{1-\lambda}. \tag{2.3}$$

This equation shows that λ shapes the degree of risk sharing: If $\lambda = 1$, trade is balanced at any point in time, and nominal consumption equals nominal output, $P_t C_t = P_t^Y Y_t$; that is, the economy is in financial autarky. If $\lambda = 0$, the tax is zero and markets are complete, implying a constant level of nominal consumption across countries, $P_t C_t = P_t^* C_t^*$. The parameter λ may therefore be thought of as an inverse index of financial market integration, with higher values indicating a lower degree of integration.

More generally, a 1 percent increase in nominal output raises nominal consumption by λ percent. The risk-sharing parameter β is therefore given by

$$\beta^U = \lambda$$

and is consequently constant over time.

2.1.2 Trade in goods

International trade in goods is modeled along the lines of Armington (1969), as is standard in this literature.[2] More formally, consumption in this economy consists of a CES aggregate of domestically produced goods, H_t, and imports, M_t, with an elasticity of substitution given by $\psi \geq 0$:

$$C_t = \left((\omega)^{\frac{1}{\psi}} (H_t)^{\frac{\psi-1}{\psi}} + (1-\omega)^{\frac{1}{\psi}} (M_t)^{\frac{\psi-1}{\psi}} \right)^{\frac{\psi}{\psi-1}}.$$

Here, ω describes the home bias. The nominal price of the domestic good is P_t^Y, whereas the nominal price of the imported good is normalized to 1. The optimal demand for domestic goods and imported goods is given by

$$H_t = \omega C_t \left(\frac{P_t^Y}{P_t} \right)^{-\psi} \qquad \text{and} \qquad M_t = (1-\omega) C_t \left(\frac{1}{P_t} \right)^{-\psi}.$$

[2] As shown by Arkolakis, Costinot and Rodríguez-Clare (2012), more complex trade models (e.g Eaton and Kortum, 2002; Melitz and Ottaviano, 2008) can be mapped into the Armington model and yield, appropriately calibrated, the same equilbrium conditions up to first order.

Similarly, the rest of the world demands some of the SOE's goods. We assume a similar demand schedule:

$$X_t = \bar{X}^{nom} \left(P_t^Y\right)^{-\psi},$$

where $\bar{X}^{nom}$ is the value of exported goods in the non-stochastic steady state.

By plugging in the optimal demand for domestic goods, H into the market clearing condition, $Y_t = H_t + X_t + G_t$, we obtain a particularly simple relationship between the multiplier and the risk-sharing β, as summarized by Proposition 1.[3]

Proposition 1 (Risk sharing and open-economy multipliers) *Consider the model discussed in this section. The share of unsmoothed income is given by*

$$\beta_t^U \equiv \frac{\hat{C}_t^{nom}}{\hat{Y}_t^{nom}} = \lambda.$$

Assuming a unit Armington elasticity, $\psi = 1$, the open-economy relative output multiplier is increasing in the share of unsmoothed income, β^U:

$$m_t = \frac{1}{1 - \omega c_{t-1}^{nom} \beta^U} = \frac{1}{1 - \omega c_{t-1}^{nom} \lambda}$$

where ω is the share of final expenditure that falls on domestic goods, $c_{t-1}^{nom} \equiv \frac{C_{t-1}^{nom}}{Y_{t-1}^{nom}}$ is the share of nominal consumption in nominal output in $t-1$ and λ a parameter governing financial risk sharing. The corresponding nominal consumption multiplier is given by

[3] Inserting the demand for H into the market clearing condition (with $\psi = 1$) and multiplying by P_t^Y yields

$$Y_t^{nom} = \omega C_t^{nom} + G_t^{nom} + \bar{X}^{nom}.$$

Taking differences and dividing by G_t^{nom} yields

$$m_t = \omega \frac{\Delta C_t^{nom}}{\Delta Y_t^{nom}} m_t + 1,$$

with $m_t \equiv \frac{\Delta Y_t^{nom}}{\Delta G_t^{nom}}$ being the open-economy multiplier. Re-arranging gives

$$m_t = \left(1 - \omega \frac{C_{t-1}^{nom}}{\Delta Y_{t-1}^{nom}} \beta_t^U\right)^{-1},$$

where $\beta_t^U \equiv \frac{\hat{C}_t^{nom}}{\hat{Y}_t^{nom}}$ is the elasticity of consumption to output.

$$m_t^c = \beta^U \times \frac{1}{1-\omega c_{t-1}^{nom} \beta^U} = \lambda \times \frac{1}{1-\omega c_{t-1}^{nom} \lambda}.$$

The multiplier depends on openness in both goods and financial markets (as well as the share of consumption in output). An increase in trade openness, i.e., a reduction in the home bias ω, lowers the multiplier. More importantly for this paper, an increase in financial market integration, i.e. a reduction in λ, lowers the multiplier as well. Hence, higher financial market

To better understand this result, consider a \$1 increase in external demand. With complete financial markets ($\beta^U = 0$), consumption remains fixed such that output rises by \$1. The multiplier in that case is therefore $m|_{\beta^U=0} = 1$. In financial autarky ($\beta^U = 1$), the increase in output translates into a one-for-one increase in consumption. A share ω of this increase in consumption falls on domestic goods, which further stimulates output and consumption, and so on. The total impact on output is therefore larger than in the case of complete financial markets markets:[4]

$$m|_{\beta^U=1} = 1 + \omega + \omega^2 + ... = \frac{1}{1-\omega} \geq 1 = m|_{\beta^U=0}.$$

The continuous, blue line in Figure 2 illustrates this non-linear relationship between β^U and the multiplier, indicating a clear positive slope: Less risk sharing (an increase in β^U) raises the multiplier. Quantitatively, the effect of risk sharing on the multiplier is economically significant, but somewhat modest at high levels of risk sharing. Starting from a multiplier equal to $m = 1$ in the case of perfect risk sharing, the multiplier grows to $m = 1.15$ for $\beta^U = 0.2$ and $m = 1.5$ for $\beta^U = 0.5$.

The dashed red line in Figure 2 depicts the consumption multiplier $m^c = \lambda \times \frac{1}{1-\omega\lambda}$. For high levels of financial market integration (low λ), the direct benefits of risk sharing dominate, and a small change in the degree of financial market integration λ translates almost one for one into changes in the consumption response. For lower levels of financial market integration, however, the indirect benefits in form of a large output multiplier dominate, and small changes in financial market integration can have larger effects on consumption volatility through their effect on output volatility.

A key parameter in the relationship between risk sharing and the multiplier is the home bias. This relationship is stronger the larger the home bias because a larger home bias amplifies the feedback loop. For the purpose of the figure, we choose $\omega = 0.69$, as in Nakamura and Steinsson (2014), who argue that this number describes well the home bias observed in the average U.S. state. Another key parameter is the Armington elasticity. The assumption of a unit Armington elasticity

[4]In this model, consumption equals GDP in steady state. Hence, the elasticity β^U is the same as the aggregate MPC defined as $\frac{\Delta C_t^{nom}}{\Delta Y_t^{nom}}$. Therefore, Proposition 1 can be interpreted as a generalization of the Keynesian multiplier to an open-economy context.

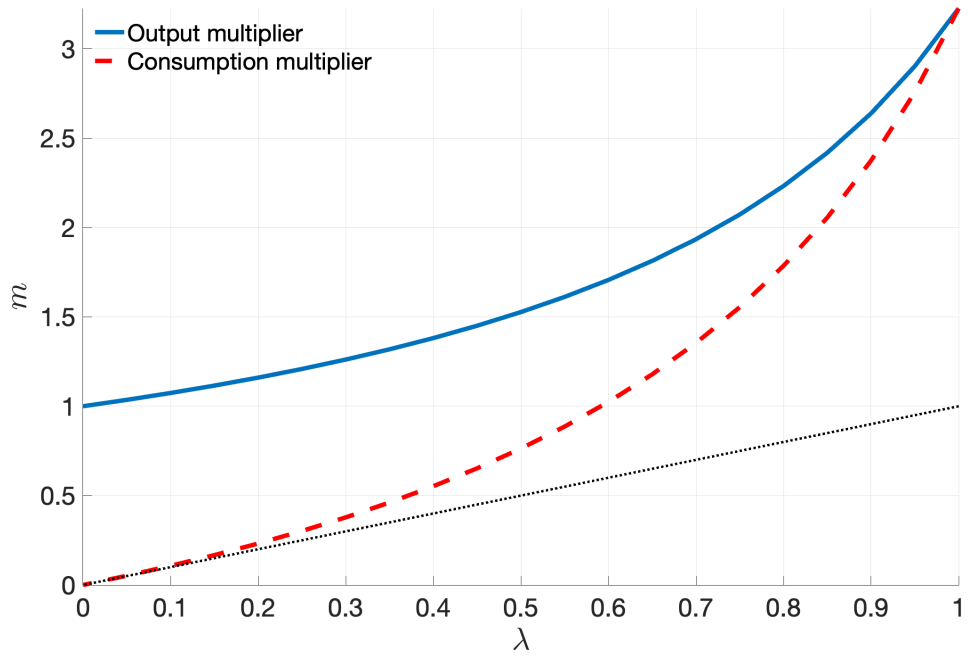


Figure 2: RISK SHARING AND THE OPEN-ECONOMY RELATIVE MULTIPLIER

Notes: The figure displays the relationship between the elasticity of nominal consumption to nominal GDP, $\beta^U = \lambda$, on the x-axis and the open-economy relative output multiplier $m = \frac{1}{1-\omega\lambda}$ and consumption multiplier, $m^c = \lambda m$, on the y-axis. The home bias parameter is set to $\omega = 0.69$. The black dashed line corresponds to the 45 degree line $\lambda = \lambda$ and displays the direct effect of risk sharing.

is convenient in deriving the relationship between the multiplier and the degree of risk sharing because it implies a constant expenditure share on domestic goods. With a higher Armington elasticity, the increase in external demand and the ensuing upward pressure on prices would amplify expenditure switching and cause the expenditure share on domestic goods to decline, thereby weakening the feedback loop and the relationship between risk sharing and the multiplier. Empirical estimates of the Armington elasticity in the literature range from relatively high values in the long run (around 5–8; see, e.g., Broda, Greenfield and Weinstein, 2006) to smaller values that are potentially below one at business cycle frequency (see, e.g., Boehm, Levchenko and Pandalai-Nayar, 2020). Nakamura and Steinsson (2014) set up a DSGE model where households have access to complete financial markets (which corresponds to the case $\lambda = 0$) and set the Armington elasticity to 2. Consequently, their model predicts multipliers below 1: 0.83 for a model variation with sticky prices that mutes the expenditure-switching effect and 0.43 for a model variation with flexible prices (see their Table 2).

3 Empirical Analysis

The previous section emphasized that risk sharing reduces consumption volatility both directly and indirectly through its effect on output volatility. In this section, we present empirical evidence on the degree of risk sharing across U.S. states. In line with our previous analysis, we focus on external demand shocks, which allows us to give our estimates a structural interpretation. Whereas previously, we summarized risk sharing through an imperfection in the trade of state-contingent bonds, we now distinguish between different risk-sharing channels that can be mapped more easily to the data. Our estimates then serve as a key input for our quantitative model, which we use to perform counterfactual analyses.

3.1 Conceptual framework

Following the seminal work by Asdrubali, Sørensen and Yosha (1996) and later extended by Kohler, Müller and Wellmann (2023), we consider the following identity that decomposes aggregate GDP, $Y_{i,t}$:

$$Y_{i,t} = \frac{Y_{i,t}}{y_{i,t}} \times \frac{y_{i,t}}{pi_{i,t}} \times \frac{pi_{i,t}}{di_{i,t}} \times \frac{di_{i,t}}{c_{i,t}} \times c_{i,t},$$

where i refers to a U.S. state and t refers to a year. Here, $y_{i,t}$ is per capita GDP, $pi_{i,t}$ is per capita personal income, $di_{i,t}$ is per capita personal disposable income, and $c_{i,t}$ is personal consumption.[5]

Taking logs and first differences yields

$$\Delta \ln Y_{i,t} = \underbrace{\Delta \ln Y_{i,t} - \Delta \ln y_{i,t}}_{\text{Migration}} + \underbrace{\Delta \ln y_{i,t} - \Delta \ln pi_{i,t}}_{\text{Capital markets}} + \underbrace{\Delta \ln pi_{i,t} - \Delta \ln di_{i,t}}_{\text{Fiscal transfers}} + \underbrace{\Delta \ln di_{i,t} - \Delta \ln c_{i,t}}_{\text{Credit markets}} + \Delta \ln c_{i,t}. \tag{3.1}$$

Fluctuations in aggregate GDP do not translate into fluctuations of per-capita consumption to the extent that they are absorbed by one of four risk-sharing channels. First, an outflow of workers during a recession dampens the fall in per-capita GDP. This is captured by the difference in the log change of aggregate GDP and per-capita GDP, that is, the log change in the working-age population.[6] Second, a fall in per capita GDP in one region might not translate into a decline in household income within that region. The difference between the two reflects factor income flows across regions, including commuting and cross-border dividend and interest payments, capital depreciation,

[5] We transform variables into per-capita terms using working-age population. We focus on working-age population because that part of the population is more likely to move in response to changing economic conditions.

[6] By measuring migration through changes in working-age population, we implicitly assume that cyclical variations in the growth rate of working-age population are driven by net migration.

and corporate retained earnings (that is, net profits less dividends). Following ASY, we denote this channel as capital market smoothing. Third, cross-regional fiscal transfers might further cushion the fall in households' personal income, creating a wedge between the log differences of personal income and disposable income. Finally, to the extent that households have access to credit markets, they can smooth out fluctuations in their disposable income through borrowing or lending. Any risk not absorbed by the four risk-sharing channels remains unsmoothed and leads to a positive comovement between changes in aggregate GDP and per capita consumption.

ASY and the large literature building on their work estimate the strength of each channel using a set of OLS regressions akin to a decomposition of the unconditional variance of output. We depart from this literature by examining the response of the risk-sharing channels to an identified external demand shock—specifically, military spending shocks as in Nakamura and Steinsson (2014). This approach allows us to calibrate the key risk-sharing parameters in our structural model based on the identified responses, which we then employ to perform counterfactual analyses.

More precisely, we estimate a set of equations along the following lines:

$$\frac{Y_{i,t} - Y_{i,t-2}}{Y_{i,t-2}} = \alpha_i + \alpha_t + \gamma^m \frac{G_{i,t} - G_{i,t-2}}{Y_{i,t-2}} + \mathbb{Z}_{i,t} + \varepsilon_{i,t} \tag{3.2}$$

Here, $Y_{i,t}$ is aggregate nominal GDP in state i at time t, $G_{i,t}$ is nominal federal military spending and $\mathbb{Z}_{i,t}$ is a vector of control variables that we discuss below. The estimate for γ^m refers to the open-economy relative multiplier. As in Nakamura and Steinsson (2014), we focus on two-year differences to mitigate any timing mismatches between the military spending prime award date—the type of data that is stored in our dataset—and the effective date of spending (see also Brunet 2022, as well as Briganti and Sellemi 2023). This approach also allows us to partially control for the mismatch between the fiscal year in which the military award data are recorded (starting on October 1st of the previous calendar year) and the calendar year that the other macroeconomic variables follow.

To estimate the effect of military spending shocks on the different risk-sharing channels, we re-run regression (3.2) for different outcome variables: $\Delta \ln Y_{i,t} - \Delta \ln y_{i,t}$ for the migration channel (with coefficient γ^M); $\Delta \ln y_{i,t} - \Delta \ln pi_{i,t}$ for the capital market channel (γ^K); $\Delta \ln pi_{i,t} - \Delta \ln di_{i,t}$ for the fiscal transfer channel (γ^F); $\Delta \ln di_{i,t} - \Delta \ln c_{i,t}$ for the credit market channel (γ^C); and $\Delta \ln c_{i,t}$ for the unsmoothed part (γ^U).

The coefficients of these regressions provide the elasticity of the different risk-sharing channels to a military spending shock. To compare our estimates with those in the literature that decompose the unconditional variance of output, we calculate the ratio of each outcome variable's coefficient to the multiplier coefficient, together with the appropriate standard error calculated from the variance-

covariance matrix of the estimates (e.g., $\hat{\beta}^m = \frac{\hat{\gamma}^M}{\hat{\gamma}^m}$). We estimate the variance-covariance matrix of the estimates by stacking the dataset, saturating it with a sample indicator, and running all regressions together (Angrist, Hull and Walters, 2023; Angrist and Hull, 2023).[7]

Following Nakamura and Steinsson (2014) and Dupor and Guerrero (2017), we instrument state-level military spending using a shift-share design. The instrument is constructed as the product of two components: a *shift* that captures the change in national military spending, normalized by national output, $\frac{G_t - G_{t-2}}{Y_{t-2}}$, and a *share* given by the two-year lagged ratio of state i's government spending to national government spending, divided by the state's share of GDP in national GDP, $\frac{s^G_{i,n}}{s^Y_{i,n}} = \frac{\frac{G_{i,n}}{G_n}}{\frac{Y_{i,n}}{Y_n}}$:

$$\text{instr}^{\text{NS,Bartik}}_{i,t} = \left(\frac{1}{2} \sum_{n=t-4}^{t-3} \frac{s^G_{i,n}}{s^Y_{i,n}} \right) \times \frac{G_t - G_{t-2}}{Y_{t-2}} \tag{3.3}$$

If a given state has a higher share of military spending per capita than the nation as a whole, a change in national military spending will affect it more than another state with a lower exposure. This instrument is valid if the change in national military spending is orthogonal to state-specific shocks, which is likely because military spending is determined by national security considerations rather than by differences in state-specific economic conditions.

Following the literature on shift-share designs (Adao, Kolesár and Morales, 2019; Goldsmith-Pinkham, Sorkin and Swift, 2020; Borusyak, Hull and Jaravel, 2022), we ensure the validity of our instrument through several steps. First, we lag our exposure shares to ensure that the instrument does not capture variation in national military spending across states. Second, since variations in the instrument are driven by both changes in the shifts and changes in the exposure shares, we include the shares as controls. Additionally, the time fixed effects will capture time-varying changes in these exposure shares. These time fixed effects will further control for lagged values of the national shifts, ensuring the instrument cannot be forecasted from prior innovations. Third, we include lags of changes in national military spending to ensure we capture only the innovations in government spending[8]. Finally, we obtain unbiased standard errors by clustering our regressions at the treatment level—the state level in this case. To ensure our regressions account for the relative importance of different states for output fluctuations, we weight our regressions by the share of state-level GDP in national GDP.

[7] We take a nonlinear combination of each risk-sharing channel and the multiplier, with standard errors calculated using the Delta method. See also Erickson and Whited (2002) and Andrews, Gentzkow and Shapiro (2020) for the econometric background behind this approach.

[8] Our results are robust to adding a greater set of control variables, like lags of the dependent variable or further lags.

3.2 Data

Geographical Coverage and Sample Period We analyze risk sharing in the 50 U.S. states plus the District of Columbia. Our sample is annual, starts in 1966, the first year for which state-level military spending data is available, and ends in 2019.

Data Sources Data on state GDP, personal income, disposable income, and consumption are obtained from the Bureau of Economic Analysis (BEA).[9] We focus on consumption of non-durable goods and services because they map better into the standard international risk-sharing framework where purchases of goods coincide with their consumption.[10] Official data on state-level consumption are available only from 1997 onward. For years prior to 1997, we estimate state-level consumption growth by allocating national consumption growth to states based on their wage growth rates for sectors producing non-tradable, non-durable consumption goods.[11] Data on state-level working-age population are obtained from the National Cancer Institute's Surveillance, Epidemiology, and End Results Program. Military spending data are obtained from the universe of U.S. military prime contracts, aggregated to the state level. Data before 2004 is obtained from the Department of Defense's Directorate for Information Operations and Control archives. Data after 2004 is accessed from usaspending.gov. All data are in nominal terms.

3.3 Results

Panel A of Table 1 presents the regression results using our instrumental variable (IV) approach to military spending shocks. Prior to discussing our empirical estimates, we assess the relevance of the instrument using the weak-instrument-robust test proposed by Montiel Olea and Pflueger (2013). The effective first-stage F-statistic is 174 for all regressions, well above the critical value for a 5% worst-case bias (37.4). This indicates that the instrument is very strong.

Turning to the results of regression (3.2), we observe a fairly precisely estimate of 1.43 (0.38) for the open-economy relative multiplier, which is close to the estimate reported by Nakamura and Steinsson (2014). Columns (2)–(6) present the estimates for the risk-sharing channels. To facilitate interpretation and comparison with the literature, we focus on the implied risk-sharing coefficients

[9] Appendix **??** provides more details about data sources, including a table decomposing the relationship between GDP and personal consumption.

[10] Appendix X shows that durable consumption responds somewhat more strongly to an increase in output, consistent with the view that purchases of durables act as an additional saving device.

[11] For the years where both series are available, the average correlation between our proxy and actual state-level consumption growth is 0.83. This correlation is higher than when using retail wages (0.81) or retail sales growth (0.22), two alternative proxies used in the literature (Asdrubali, Sørensen and Yosha, 1996; Hoffmann et al., 2019; Guren et al., 2021).

reported in Panel B. According to the estimates, a $1 increase in aggregate output caused by an increase in federal military spending leads to an increase in per capita consumption of approximately $0.47. The $1 increase in aggregate output is partially offset by an increase in population, resulting in a rise of per capita output by only $0.88. Per-capita personal income increases by $0.65, implying that a significant portion of the additional output is either retained as corporate earnings or distributed as dividends to nonresident households (although the coefficient of 0.26 is statistically significant only at the 10% level).[12] Per-capita disposable income only increases by $0.49, indicating that the federal tax and transfer system smooths out an additional $0.14. Credit markets only play a minor role in risk sharing with a small, positive, but statistically insignificant coefficient of 0.04.

Panel C of Table 1 reports the risk-sharing coefficients obtained by regressing the various channels on the log change of output. This is similar to the regressions in ASY and the subsequent studies except that we consider two-year rather than one-year changes in both output and the dependent variables.[13] We find that the migration and credit market channels are of similar magnitude in both approaches. However, the fiscal transfer channel is almost double, and the capital market channel is more precise and significantly larger at 40%, compared to 23% in the IV framework. The variance decomposition appears to overstate the overall strength of risk sharing and the importance of capital markets, while understating the role of fiscal transfers. This discrepancy is likely due to the fact that the variance decomposition does not account for the endogeneity of the risk-sharing channels to output fluctuations.

Notably, the simple model in Section 2 predicts the relationship between risk sharing and the open-economy multiplier observed in the data. According to Proposition 1 and as displayed in Figure 2, the multiplier is approximately 1.5 in an economy where households smooth about less than half of an income shock. In our case, the multiplier is 1.43, and households smooth a bit more than half of the shock.

4 Model

In the simple model discussed in Section 2, we used a tax on state-contingent bonds to vary the degree of financial risk sharing. In this section, we present a quantitative model that replaces this generic mechanism by several risk-sharing channels that we can map more easily to the data. This does not only allow us to quantify the indirect benefits of overall risk sharing, but also allows us to explore the interdependencies of the various risk-sharing channels.

[12] Further decomposing the capital market channel, we find that a negligible amount, approximately $0.03, is smoothed out through commuting, as some labor income generated in a state is earned by workers residing in other states.

[13] See Appendix XXX for regression results with one-year changes.

Table 1: EMPIRICAL ESTIMATES OF RISK-SHARING CHANNELS

	(1) Output	(2) Migration	(3) Capital market	(4) Fiscal transfers	(5) Credit market	(6) Uninsured
Panel A: IV coefficients						
γ	1.43***	0.17	0.33	0.20***	0.05	0.68***
	(0.38)	(0.13)	(0.23)	(0.06)	(0.14)	(0.25)
F-statistic	174.12	174.12	174.12	174.12	174.12	174.12
Panel B: IV Risk-sharing coefficients						
β	–	0.12*	0.23*	0.14***	0.04	0.47***
		(0.07)	(0.14)	(0.03)	(0.10)	(0.15)
Panel C: Risk-sharing coefficients from variance decomposition (OLS)						
β_{ASY}	–	0.14***	0.40***	0.08***	0.07	0.31***
		(0.02)	(0.06)	(0.01)	(0.06)	(0.03)

Robust standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Column (1) of Panel A displays the estimated coefficient for γ^m from regression (3.2). The remaining columns display the estimated coefficients for alternative dependent variables, as explained in the text. The regressions include state fixed effects and control for the exposure shares, and lagged values of the independent and dependent variables and the military spending shifts. Standard errors are clustered by state. Data before 1997 uses non-traded sector wages as a proxy for consumption. Period: 1968-2019. Number of observations: 2,550. Panel B displays the risk-sharing parameters implied by Panel A, e.g. $\hat{\beta}^M = \frac{\hat{\gamma}^M}{\hat{\gamma}^m}$, with standard errors computed from the variance-covariance matrix of the stacked coefficients. Panel C displays the risk-sharing coefficients from regressing the various dependent variables on the log change of output using OLS, as in ASY. The OLS regressions use the same sample as in Panel A, including state fixed effects and time fixed effects.

The world economy consists of two countries that form a currency union. The focus is on a limiting case where the size of the foreign economy (denoted by an asterisk) goes towards infinity and the home economy becomes a small open economy.[14] Several model ingredients are standard in this class of models, such as sticky wages, Armington trade and a production process that requires capital and labor. As in the simple model, we focus on external demand shocks as a driving force.

We complement this baseline model with four features to capture the distinct risk-sharing channels emphasized in the empirical section: First, we allow for labor mobility across countries as in House, Proebsting and Tesar (2023) to model the migration channel. Second, households' stock holdings are potentially biased towards domestic firms to capture imperfect risk sharing through

[14]The exposition concentrates on the home economy, but equivalent equations hold for the foreign economy as well.

capital markets. Third, the fiscal transfer channel is modeled through a progressive (federal) income tax as in Heathcote, Storesletten and Violante (2017). Finally, households can trade in non-contingent bonds subject to adjustment costs to capture the credit market channel.

4.1 Households, Population and Migration

The small open economy is populated by capital owners and workers. Capital owners are immobile. Their number is given by $\mathbb{N}^k$. Workers are mobile and can move between the small open economy and the rest of the world if they find it optimal to do so. The number of workers in the small open economy at time t is given by $\mathbb{N}_t^w$ and the total population is

$$\mathbb{N}_t = \mathbb{N}^k + \mathbb{N}_t^w. \tag{4.1}$$

Variables are indicated in per capita terms. For example, C_t^w is the consumption of a single worker while the total consumption for workers is $\mathbb{N}_t^w C_t^w$.

4.1.1 Capital owners

Capital owners receive income from dividends and interest. They own a (fixed) share κ of domestic firms that pay out dividends $P_t D_t$, where P_t is the price index of the final good and D_t are real dividends per capital owner. They also own a share $1-\kappa$ of foreign firms that pay out D^* (the foreign price index is normalized to 1). The parameter κ indicates the home bias in capital owners' firm ownership. As κ goes to zero, capital owners are more and more diversified. *We later calibrate κ to match the empirically observed strength of the capital market channel.*

Capital owners can also trade in bonds, but adjusting one's bond position entails some cost, given by $\frac{\iota}{2} P_t \left(B_t^k\right)^2$, where B_t^k is the number of bonds purchased by capital owners in t, and $\iota \geq 0$ disciplines the strength of the adjustment costs.[15] The higher ι, the more costly it is for capital owners to use bonds to smooth income shocks. *We later calibrate ι to match the empirically observed strength of the credit market channel.* Bonds mature in one period and pay a constant nominal interest rate i.[16] The face value of bonds purchased in $t-1$ is given by $P_{t-1} B_{t-1}^k$. Taken together, nominal capital income is

$$P_t Y_t^k = \kappa P_t D_t + (1-\kappa) D^* + P_{t-1} B_{t-1}^k i - \frac{\iota}{2} P_t \left(B_t^k\right)^2.$$

[15] These adjustment costs are common in the literature on open-economy models because they induce stationarity (Schmitt-Grohé and Uribe, 2003). A very small value of ι is enough to induce stationarity. More recently, calibrations with larger values of ι have been proposed to model segmentation in international financial markets, see e.g. Maggiori (2022).

[16] The fact that the nominal interest rate on bonds is fixed results from our SOE assumption.

Income is taxed according to a log-linear tax and transfer function (Heathcote, Storesletten and Violante, 2017): For any pre-government income y, the disposable income is given by $\mu y^{1-\tau}$. The parameter μ captures the level of taxation, while the parameter $\tau < 1$ measures tax progressivity. Positive values of τ indicate a progressive tax system, while negative values indicate a regressive tax system. If $\tau = 0$, the taxes paid are proportional to income and the tax rate is given by $1 - \mu$. In that case, a 1 percent increase in pre-tax income raises post-tax income by 1 percent and the fiscal system does not contribute to consumption smoothing. More generally, a 1 percent increase in pre-tax income leads to a $1 - \tau$ percent rise in post-tax income. The larger τ, the stronger the fiscal transfer channel. *We later calibrate τ to match the empirically observed strength of the fiscal transfer channel.*

Capital owners use their income to purchase consumption goods at price P_t and to save or borrow by adjusting their bond position, $P_t B_t^k - P_{t-1} B_{t-1}^k$. Taken together, at any point in time t, capital owners choose consumption, C_t^k, and bond holdings, B_t^k, to maximize the expected discounted sum of utility

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \ln\left(C_{t+j}^k\right)$$

subject to the nominal budget constraint

$$P_t C_t^k + P_t B_t^k - P_{t-1} B_{t-1}^k = \mu \left(P_t Y_t^k\right)^{1-\tau}. \tag{4.2}$$

The optimal demand for bonds satisfies the following Euler equation

$$1 = \beta \mathbb{E}_t \left\{ \frac{C_t^k}{\pi_{t+1} C_{t+1}^k} \frac{1 + \mu(1-\tau)\left(P_{t+1} Y_{t+1}^k\right)^{-\tau} i}{1 + \iota\mu(1-\tau)\left(P_{t+1} Y_{t+1}^k\right)^{-\tau} B_t^k} \right\}, \tag{4.3}$$

where $\pi_{t+1} = \frac{P_{t+1}}{P_t}$ is inflation. Notice that absent income taxation and adjustment costs, this collapses to the standard Euler equation, $1 = \beta(1+i)\mathbb{E}_t \left\{ \frac{C_t^k}{\pi_{t+1} C_{t+1}^k} \right\}$.

4.1.2 Workers

Workers are mobile and earn labor income. Workers choose their location (Home or Foreign) at the beginning of each period and consume and work in their new location within the same period. Each period a share n_t of workers migrate from Home to Foreign, whereas a share n_t^* of workers migrate from Foreign to Home. The number of workers available to work in Home in period t is therefore:

$$\mathbb{N}_t^w = (1 - n_t)\mathbb{N}_{t-1}^w + n_t^* \mathbb{N}^*.$$

Workers residing in Home receive nominal labor income $W_t L_t$, where wages W_t and labor L_t are set by a labor union (see below).

Similar to capital owners, we allow workers to trade bonds. But to keep the model tractable, we assume perfect risk sharing within the group of workers in each location (see e.g. Lucas Jr, 1990; Bilbiie, 2024, for a similar approach). As workers move across countries, they take their bond holdings with them. The average worker leaving Home in t therefore leaves with B^w_{t-1} and the average worker arriving in Home in t arrives with zero bonds (since bond holdings in the large Foreign economy are equal to their steady-state level, which we set to zero). This makes borrowing particularly cheap for the group of workers because they do not need to repay any debt incurred by emigrants. To compensate for this incentive to overborrow, we assume that they face a higher interest rate, which is set to $\frac{i}{1-\bar{n}}$. Total income for workers in Home is therefore

$$\mathbb{N}^w_t P_t Y^w_t = \mathbb{N}^w_t W_t L_t + \mathbb{N}^w_{t-1} \frac{1-n_t}{1-\bar{n}} P_{t-1} B^w_{t-1} i - \mathbb{N}^w_t \frac{\iota}{2} P_t \left(B^w_t\right)^2.$$

Similar to capital workers, workers are subject to taxation. Then, the budget constraint is

$$\mathbb{N}^w_t P_t C^w_t + \mathbb{N}^w_t P_t B^w_t - \mathbb{N}^w_{t-1} P_{t-1} B^w_{t-1} = \mu \mathbb{N}^w_t \left(P_t Y^w_t\right)^{1-\tau}. \tag{4.4}$$

When living abroad, workers' consumption is fixed and given by $C^{w,*}$. The "household head" of workers in Home chooses workers' consumption and bond holdings to maximize the expected discounted sum of utility of the average worker subject to (4.4), taking the migration decisions as given. The optimal demand for bonds then satisfies the following Euler equation:

$$1 = \beta \mathbb{E}_t \left\{ \frac{C^w_t}{\pi_{t+1} C^w_{t+1}} \frac{1 + \mu(1-\tau)\left(P_{t+1} Y^w_{t+1}\right)^{-\tau} \frac{1-n_{t+1}}{1-\bar{n}} i}{1 + \iota\mu(1-\tau)\left(P_{t+1} Y^w_{t+1}\right)^{-\tau} B^w_t} \right\}. \tag{4.5}$$

We next discuss workers' migration decision and the determination of labor supply.

Labor mobility At the beginning of each period, workers choose to migrate or remain in their current country. Here, we focus on the location decision faced by workers currently living in the home country, but equivalent conditions hold for workers living in the foreign country.

A worker moving from Home to Foreign incurs a migration cost v. Each period, workers receive idiosyncratic (i.e., worker-specific) shocks for each of the two countries, denoted by ϵ_t and ϵ^*_t. Define $v_t\left(\epsilon_t, \epsilon^*_t\right)$ as the value of a worker currently living in Home conditional on the aggregate

state and the worker's vector of idiosyncratic shocks. The value of living in Home at time t is

$$v_t(\epsilon_t, \epsilon_t^*) = \max\left\{U\left(C_t^w\right) + \frac{1}{\gamma}\epsilon_t + \beta\mathbb{E}_t\left(V_{t+1}\right), U\left(C^{w,*}\right) + \frac{1}{\gamma}\epsilon_t^* - v + \beta\mathbb{E}_t\left(V^*\right)\right\}$$

The value V_t is the expected value of $v_t\left(\epsilon_t, \epsilon_t^*\right)$ prior to the realization of the idiosyncratic shocks and thus V_t is the average expected utility of any worker in Home at the beginning of time t. Parameter γ governs how strongly idiosyncratic location shocks affect migration decisions.

We follow Artuç, Chaudhuri and McLaren (2010) and assume that the idiosyncratic shocks are i.i.d. over time and across individuals and are distributed according to a Gumbel distribution with zero mean. Given these assumptions, the average utility of a worker in Home at the beginning of time t is

$$V_t = \frac{1}{\gamma}\ln\left(\exp\left\{\gamma\left[U\left(C_t^w\right) + \beta\mathbb{E}_t\left(V_{t+1}\right)\right]\right\} + \exp\left\{\gamma\left[U\left(C^{w,*}\right) - v + \beta\mathbb{E}_t\left(V^*\right)\right]\right\}\right). \tag{4.6}$$

Migration decisions depend on this average utility. The share of workers who relocate from Home to Foreign, denoted by n_t, is

$$n_t = \frac{\exp\left\{\gamma\left[U\left(C^{w,*}\right) - v + \beta\mathbb{E}_t\left(V^*\right)\right]\right\}}{\exp\left\{\gamma V_t\right\}}. \tag{4.7}$$

Naturally, a decrease in consumption in Home, C_t^w, leads to out-migration, especially if γ is large.[17] Rewriting this expression in differences from the non-stochastic steady yields the number of emigrants as

$$n_t\mathbb{N}_{t-1}^w = \exp\left\{-\gamma\Delta V_t\right\}\bar{n}\mathbb{N}_{t-1}^w.$$

where $\Delta V_t = V_t - \bar{V}$ and $\bar{X}$ is the non-stochastic steady state of any variable X_t. In the Appendix, we derive a corresponding expression describing the number of immigrants:

$$n_t^*\mathbb{N}^{w,*} = \exp\left\{\gamma\left(\Delta U_t + \beta\mathbb{E}_t\Delta V_{t+1}\right)\right\}n^*\mathbb{N}^{w,*}.$$

Then, the number of workers living in Home at time t consists of the number of workers last period, plus net immigration:

$$\mathbb{N}_t^w = \mathbb{N}_{t-1}^w + \exp\left\{\gamma\left(\Delta U_t + \beta\mathbb{E}_t\Delta V_{t+1}\right)\right\}n^*\mathbb{N}^{w,*} - \exp\left\{-\gamma\Delta V_t\right\}n\mathbb{N}_{t-1}^w. \tag{4.8}$$

Equations (4.6) and (4.8) fully characterize the dynamics of the workers' population. The strength of

[17] Note that in the limiting case of a small open economy with $\mathbb{N}^{w,*} \to \infty$, the utility terms in Foreign, $U(C^{w,*})$ and V^*, do not respond to fluctuations in the business cycle in the small open economy.

the response of migration to consumption differentials is driven by two parameters: the migration cost υ and the inverse of the standard deviation of the idiosyncratic shocks, γ. A higher value of γ implies a lower variance for the idiosyncratic shock and migration decisions are more driven by consumption differentials rather than idiosyncratic preferences. This makes the migration response larger. The migration cost parameter pins down the share of migrants, n, in steady state. *Given υ, we calibrate γ to match the empirically observed strength of the migration channel.*

Labor supply per worker Workers supply fixed amounts of labor, normalized to 1, in their country of residence. As in Erceg, Henderson and Levin (2000), workers are randomly assigned a job $\iota \in [0,1]$ that pays a nominal wage $W_t(\iota)$. For each job, there is a labor union with market power that acts in the interest of its members in setting a wage rate.

Job-specific labor $L_t(\iota)$ is employed by competitive labor-aggregating firms who sell aggregate effective labor to goods-producing firms at wage W_t. Labor-aggregating firms choose a combination of jobs $L_t(\iota)$ to maximize their profits

$$W_t L_t - \int_0^1 W_t(\iota) L_t(\iota) d\iota.$$

subject to the definition of effective labor L_t:

$$L_t = \zeta + \left(\int_0^1 \left(L_t(\iota) - \zeta \right)^{\frac{\psi_w - 1}{\psi_w}} d\iota \right)^{\frac{\psi_w}{\psi_w - 1}},$$

where $\psi_w > 1$ is the elasticity of substitution across jobs.The parameter $\zeta > 0$ in this aggregator ensures a positive equilibrium wage in the presence of inelastic labor supply (**?**). The demand for each labor type then satisfies

$$L_t(\iota) = \zeta + \left(\frac{W_t(\iota)}{W_t} \right)^{-\psi_w} (L_t - \zeta) \tag{4.9}$$

Each period, labor unions can reset wages for individual jobs, $W_t(\iota)$, with probability $1 - \theta_w$. If they can adjust the wage, labor unions set it to maximize the expected net present value of total labor income for their workers taking the demand curve for each job (4.9) as given. This optimization problem then gives rise to the following labor supply curve:[18]

$$\tilde{\pi}_t^w = \frac{(1 - \theta_w \beta)(1 - \theta_w)}{\theta_w} \tilde{L}_t + \beta \mathbb{E}_t \left[\tilde{\pi}_{t+1}^w \right], \tag{4.10}$$

where a tilde denotes log deviations from the non-stochastic steady state and π_t^w is wage inflation at time t. Notice that if wages were fully flexible ($\theta_w \to 0$), then hours per worker would be constant.

[18] See the appendix for more details.

4.2 Firms, Production and Trade

Production takes place in a two-stage process. A first set of competitive firms produce intermediate goods and sell these either domestically to final-good firms or to firms abroad. Final-good firms, in turn, combine the domestically produced intermediates with imported intermediates to produce a final good that is used for consumption and investment.

4.2.1 Intermediate goods

Intermediate good producers own the capital stock and pay dividends to households. Nominal dividends are nominal sales less labor costs:

$$\mathbb{N}^k P_t D_t = \mathbb{N}_t P_t^Y Y_t - \mathbb{N}_t^w W_t L_t \tag{4.11}$$

Here, P_t^Y is the price of the intermediate good and Y_t is the per-capita volume of intermediate goods produced.

The production of intermediate goods requires overhead labor, $\hat{L}$, and variable labor $L_t - \hat{L}$. The introduction of overhead labor allows us to replicate the countercyclical labor share observed in the empirical analysis. Output of intermediates is

$$\mathbb{N}_t Y_t = \left(\mathbb{N}_t^w L_t - \bar{\mathbb{N}}^w \hat{L}\right)^{1-\alpha}. \tag{4.12}$$

Each period t, intermediate good firms choose dividends, D_t, and labor, L_t, to maximize the expected discounted sum of their dividends, $\mathbb{E}_t \sum_{s=0}^{\infty} \beta^{t+s} Q_{t,t+s} D_{t+s}$ subject to (4.11) and (4.12). Here, $\beta^{t+s} Q_{t,t+s}$ denotes the stochastic discount factor between t and $t+s$ and that reflects that a fraction κ of firms are owned by domestic capital owners whose marginal utility of consumption fluctuates and a fraction $1-\kappa$ is owned by foreign households with constant consumption:

$$\beta^{t+s} Q_{t,t+s} = \beta^{t+s} \left(\kappa \left(\frac{C_{t+s}^k}{C_t^k}\right)^{-1} + 1 - \kappa\right).$$

Optimal labor demand implies

$$W_t \left(\mathbb{N}_t^w L_t - \bar{\mathbb{N}}^w \hat{L}\right) = (1-\alpha) P_t^Y \mathbb{N}_t Y_t \tag{4.13}$$

4.2.2 Final-good producers

A second set of producers combine domestically produced intermediates, H_t, and imported intermediates, M_t, to produce a final consumption good C_t. We normalize the nominal price of imported

intermediate goods to 1. Final-good producers choose inputs to maximize profits

$$\max\left\{P_t C_t - P_t^Y H_t - M_t\right\}$$

subject to the CES production function

$$C_t = \left((\omega)^{\frac{1}{\psi}}(H_t)^{\frac{\psi-1}{\psi}} + (1-\omega)^{\frac{1}{\psi}}(M_t)^{\frac{\psi-1}{\psi}}\right)^{\frac{\psi}{\psi-1}}. \tag{4.14}$$

Here, the parameter ψ describes the elasticity of substitution between domestic goods and imports, and ω describes the home bias. The optimal demand for domestic intermediate goods and imported intermediate goods is given by

$$H_t = \omega C_t \left(\frac{P_t^Y}{P_t}\right)^{-\psi} \tag{4.15}$$

$$M_t = (1-\omega) C_t \left(\frac{1}{P_t}\right)^{-\psi}. \tag{4.16}$$

In a similar fashion, the Foreign country demands some of the Home country's intermediate goods. We assume that foreign demand takes the same form as (4.16). Adjusting for fluctuations in the home country's population, this demand schedule is

$$\mathbb{N}_t X_t = \mathbb{N} X Y_t^* \left(P_t^Y\right)^{-\psi}, \tag{4.17}$$

where X is the per capita amount of exported intermediate goods in the non-stochastic steady state and Y_t^* is the external demand shifter that follows an AR(1) with persistence ρ:

$$\ln Y_t^* = \rho \ln Y_{t-1}^* + \epsilon_t. \tag{4.18}$$

4.3 Market clearing

Market clearing of the intermediate good requires that the total production of tradable intermediates equals total demand:

$$Y_t = X_t + H_t. \tag{4.19}$$

Final consumption equals consumption by capital owners and by workers:

$$\mathbb{N}_t C_t = \mathbb{N}^k C_t^k + \mathbb{N}_t^w C_t^w. \tag{4.20}$$

A similar definition holds for aggregate bond holdings.

5 Quantitative analysis

We solve the model using a first-order approximation around a zero inflation steady state. We calibrate the model to U.S. states. The model is expressed at a quarterly frequency.

We partition the parameters into a set of calibrated parameters and a set of estimated parameters. Parameters that have commonly accepted values used in the international business cycle literature or parameters that have direct analogues in the data (e.g., trade shares, migration shares, etc.) are calibrated accordingly. Taking the calibrated parameters as given, we estimate the remaining four parameters pertaining to the risk-sharing channels.

5.1 Calibration

Table 2 lists the calibrated parameter values. The discount factor is set to $\beta = 0.99$, which implies an annual real interest rate of 4 percent. We set $\alpha = 0.40$ to match a labor income share of 0.60 observed for the United States (Karabarbounis and Neiman, 2013). The wage rigidity parameter is set to $\theta_w = 0.85$, consistent with recent estimates by Grigsby, Hurst and Yildirmaz (2021). Nekarda and Ramey (2020) estimate the share of overhead workers in the total workforce to be 18%. We therefore set $\hat{L} = 0.18\bar{L}$.

The migration cost parameter, v, enters the system of log-linearized equations only through its effect on the migration share, $\bar{n}$. Instead of explicitly solving for the migration cost parameter, we directly condition on the observed migration share. House, Proebsting and Tesar (2023) report that a bit more than 3 percent of residents move across state border lines each year in the U.S. We therefore set $\bar{n} = 0.8\%$.

An important set of parameters for the open economy multiplier pertains to the trade block. Drawing on U.S. Commodity Flow Survey data, Nakamura and Steinsson (2014) estimate an import share of 31% for the average U.S. state, leading us to set the home bias parameter to $\omega = 0.69$. We calibrate the elasticity of substitution between domestic and foreign goods, ψ, at 1.5, consistent with the international business cycle literature (Chari, Kehoe and McGrattan, 2002). Although this value exceeds the cross-country 2-year elasticity estimate of 1 (Boehm, Levchenko and Pandalai-Nayar, 2020), it can be explained by the fact that trade within U.S. states typically encounters lower barriers and enjoys greater market integration than international trade.

5.2 Estimation

Given the set of calibrated parameters, we employ a simulated method of moments to estimate the risk-sharing parameters—the migration elasticity (γ), the equity home bias (κ), the tax progressivity

Table 2: CALIBRATION

Description	Parameter	Value	Source / Target
Preferences			
Discount factor	β	0.99	Standard value
Production and labor			
Labor income share	$1-\alpha$	0.60	Karabarbounis and Neiman (2013)
Wage stickiness	θ_w	0.85	Grigsby, Hurst and Yildirmaz (2021)
Share overhead labor	$\hat{L}/\bar{L}$	0.18	Nekarda and Ramey (2020)
Share of migrants	$\bar{n}$	0.8%	House, Proebsting and Tesar (2023)
Trade			
Trade openness	$1-\omega$	0.31	Nakamura and Steinsson (2014)
Armington elasticity	ψ	1.5	Chari, Kehoe and McGrattan (2002)
Shock process			
Shock persistence	ρ	0.93	Nakamura and Steinsson (2014)

(τ) and the adjustment cost on bond holdings (ι).[19] Our parameter estimates are chosen to match the risk-sharing betas of the four channels: β^M, β^K, β^F and β^C. We calculate these objects by simulating quarterly data from the quantitative model, time-aggregating it up to an annual frequency, and running the regression (3.2) on this data. To run this regression, we need to define the model counterparts of the various income concepts in the data: *Aggregate nominal GDP* is

$$GDP_t^{nom} = \mathbb{N}_t P_t^Y Y_t.$$

Personal income is labor income of workers plus dividends received by capital owners as well as net interest:

$$PI_t^{nom} = \mathbb{N}_t^w W_t L_t + \mathbb{N}^k \left(\kappa P_t D_t + (1-\kappa) D^*\right) + \mathbb{N}_t P_{t-1} B_{t-1} i.$$

It differs from GDP by excluding retained earnings (which, in the model, only refer to investment since corporate saving is zero), but including dividends and interest income earned abroad. Hence, the capital market channel reflects movements in retained earnings / investment and differences in

[19] More detail on the simulation procedure and model regressions is provided in Appendix XXX.

capital market returns.[20] *Disposable income* is

$$DI_t^{nom} = \mu \left(PI_t^{nom}\right)^{1-\tau}$$

And *per-capita nominal consumption* is given by

$$P_t C_t.$$

Esimated parameters Table 3 displays the estimated parameters together with the targeted moments. Since we are exactly identified, our model perfectly matches the empirical moments. Our estimate for the equity home bias is 0.56, implying that a \$1 increase in capital income generated in state i raises received capital income in that state by \$0.56. To see whether this number is sensible, we classify capital income sources by whether they are linked to a household's local economy. Over our sample period, about 40% of capital income is likely to be generated locally (proprietors' income including taxes, rental income)—close to the home bias estimate of 56%, whereas the remaining 60% are more likely to be diversified (corporate profits including taxes and net interest).[21] Our estimate of tax progressivity is $\tau = 0.22$ is somewhat higher than the estimate found by Heathcote, Storesletten and Violante (2017). They use household-level data on income and disposable income from the Panel Study of Income Dynamics combined with the NBER's TAXSIM program to estimate $\tau = 0.184$.

Panel B of Table 3 displays some non-targeted moments. First, the model generates a countercyclical labor income share. In the data, a \$1 increase in income is split into a 45 cents increase in labor income and a 55 cents increase in capital income. With labor income accounting for 60% of total income, this implies that the labor income share falls, i.e. it is countercyclical. In the model, the split is 47 - 53. The model can generate this countercylical labor share through the presence of overhead labor. The higher the share of labor that is fixed, the smaller the share of labor that responds to fluctuations in demand and therefore, the smaller the increase in labor income.

The next row indicates the output multiplier, which is 1.43 in the data and 1.48 in the model.

[20] In particular, the difference is

$$GDP_t^{nom} - PI_t^{nom} = \mathbb{N}^k \left(P_t I_t - (1-\kappa)\left(D^* - P_t D_t\right)\right) - \mathbb{N}_t P_{t-1} B_{t-1} i.$$

[21] Over our sample period, about a quarter of capital income stems from proprietors' income, generally income generated by self-employed (e.g. restaurant owners) that, in practice, is linked to the household's home state. Another 5% stems from rental income, which might also be biased towards a household's home state. About 45% of capital income consists of corporate profits and interest rate income, where geographical diversification is easier. And another quarter relates to taxes on production and imports less subsidies. Distributing those taxes proportionately to proprietors' income and corporate profits suggests that about 40% of capital income is accounted for by proprietors' income including taxes and rental income.

Table 3: ESTIMATION

Moment	Data	Model	Parameter	Value
Panel A: Targeted moments and estimated parameters				
Migration channel (β^M)	0.12	0.12	Migration elasticity (γ)	16.48
Capital market channel (β^K)	0.23	0.23	Equity home bias (κ)	0.56
Fiscal transfer channel (β^F)	0.14	0.14	Tax progressivity (τ)	0.22
Credit market channel (β^C)	0.04	0.04	Bond adj. cost (ι)	0.23
Panel B: Non-targeted moments				
Marg. labor income share	0.45	0.47		
Output multiplier	1.43	1.48		
Consumption multiplier	0.68	0.70		

Notes: Panel A displays the estimated parameters as well as the targeted moments. The targeted moments correspond to those reported in Panel B of Table 1. To get the counterparts to these moments in the data, we simulate our model by feeding in random sequences of external demand shocks and then run the same regressions as we do in the data (see equation (3.2)). Panel B displays non-targeted moments as estimated from the empirical data and the simulated data.

Since the output multiplier is slightly larger in the model than in the data, the consumption multiplier is also slightly larger in the model (0.70) than in the data (0.68).

6 Role of risk sharing in reducing consumption volatility

To what extent does risk sharing across U.S. states lower states' consumption volatility? We answer this question by first comparing our benchmark economy to a counterfactual economy where all risk-sharing channels are shut down, and then by analyzing the role of each risk-sharing channel separately.

6.1 Role of all risk-sharing channels combined

In our benchmark economy, the output multiplier is 1.48 and the consumption multiplier is 0.70 (see Figure 3). The consumption multiplier is lower because risk sharing smooths out 47% of all income fluctuations. Removing all risk sharing raises the consumption multiplier. If income were not affected by risk sharing—in other words, if risk sharing did not affect the output multiplier—raising the share of unsmoothed income fluctuations from $\beta^U = 0.53$ to $\beta^U = 1$ would raise the consumption multiplier to the value of the output multiplier, i.e. 1.48. This effect represents the direct effect of risk sharing.

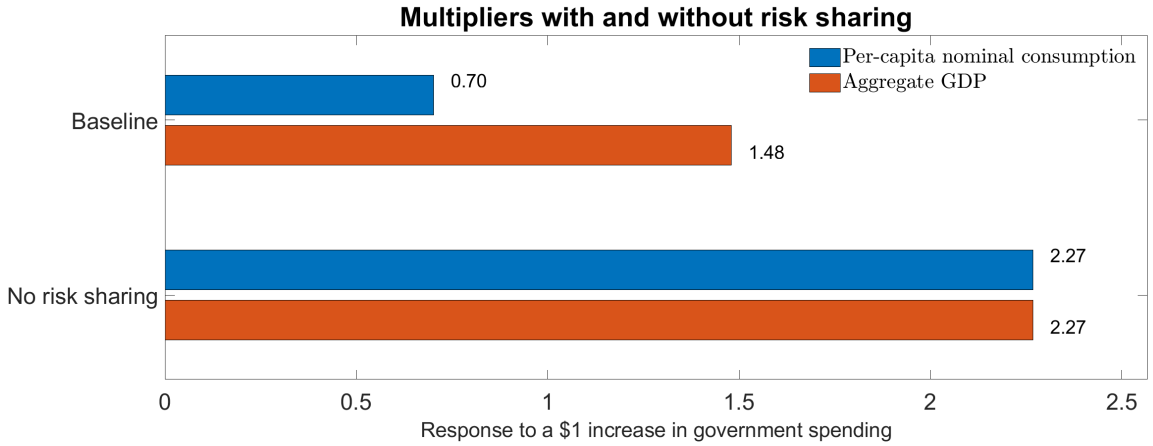


Figure 3: MULTIPLIERS WITH AND WITHOUT RISK SHARING

Notes: The figure displays multipliers for aggregate nominal GDP and per-capita nominal consumption. The estimates are based on model data, simulated from the benchmark calibration and a counterfactual calibration that shuts down all risk-sharing channels ($\gamma \to 0$, $\kappa = 1$, $\tau = 0$ and $\iota \to \infty$).

However, the full impact on consumption volatility is even larger because the absence of risk sharing also increases the output multiplier itself. In other words, a reduction in risk sharing not only makes β^U larger but also amplifies output fluctuations, thereby further increasing consumption volatility through general equilibrium effects. As shown in Figure 3, the output multiplier increases to 2.27 in the model without risk sharing. Therefore, taking into account that removing risk sharing also amplifies output fluctuations, the consumption multiplier rises from 0.70 in the benchmark to 2.27 in the counterfactual. This additional increase in the consumption multiplier (from 1.48 to 2.27) reflects the indirect benefits of risk sharing and accounts for about 40% of the overall increase in the multiplier.[22]

6.2 Role of individual risk-sharing channels

The previous section studied the role of all risk-sharing channels combined in smoothing out consumption volatility, highlighting the feedback mechanism of risk sharing on income volatility. In this section, we study the role of individual risk-sharing channels by shutting down only one channel while keeping the remaining parameters at their baseline value. This helps identify the key risk-sharing channels in the United States.

[22]Based on equation (2.1), the log change in the consumption multiplier can be written as

$$\Delta log\left(m_t^C\right) = \Delta log\left(\beta_t^U\right) + \Delta log\left(m_t\right),$$

where the first term reflects the direct effect of risk-sharing via the unsmoothed income share and the second term reflects the indirect effects via the output multiplier.

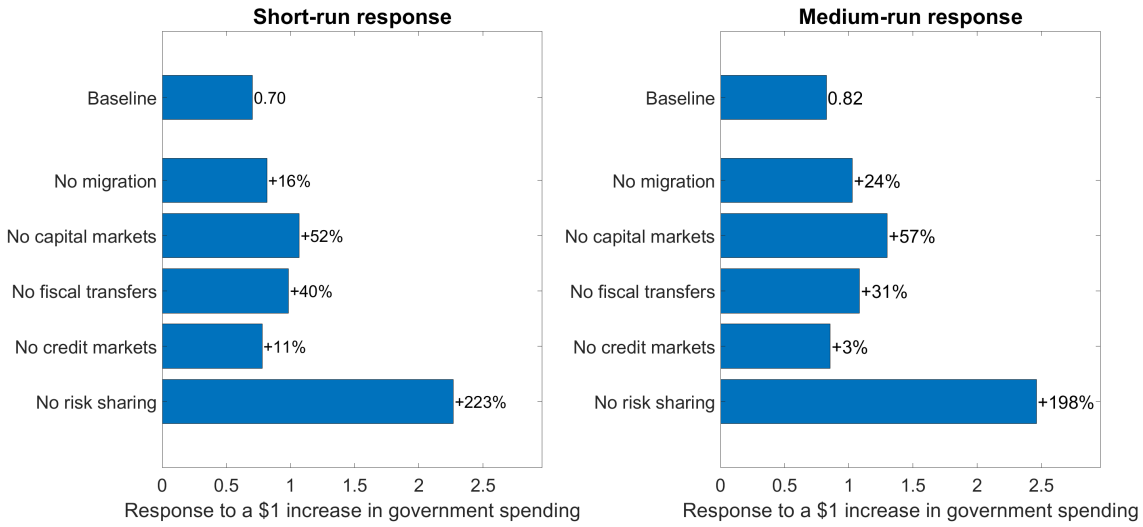


Figure 4: CONSUMPTION MULTIPLIERS FOR COUNTERFACTUALS

Notes: The first horizontal bar in each subplot refers to the response of per-capita nominal consumption to a $1 increase in income. The short-run response denotes the equivalent response to the empirical specification (2-year differences). The medium-run responses denotes the change in per-capita nominal consumption, 4 years after the short-run response.

One might be tempted by deriving the importance of a risk-sharing channel from its β coefficient. But this would ignore that households optimally adjust their behavior in response to policy changes. For instance, lowering fiscal transfers across states raises households' incentives to migrate in response to reduced personal income. Our model allows us to answer these policy counterfactuals by accounting for these interdependencies across risk-sharing channels.

Figure 4 displays consumption multipliers for various counterfactuals where only one risk-sharing channel is shut down at a time. Multipliers are expressed in percent changes compared to the baseline multiplier. The short-run response (left panel) displays the impact multipliers (estimated along the lines of regression (3.2)), whereas the medium-run multipliers (right panel) show the response at a horizon of four years obtained from regressing consumption changes between $t-2$ and $t+4$ on changes in government spending between $t-2$ and t.

The figure reveals that risk-sharing channels act as substitutes: the overall increase in the consumption multiplier when all risk-sharing channels are eliminated (+223%) is more than double the combined increase from removing one channel at a time (+119%). Removing only one channel triggers compensatory increases in other channels that partly mitigate the loss of risk sharing. The figure therefore underscores why a structural model is necessary to capture the changes in household behavior and to quantify these interdependencies between risk-sharing channels.

Looking more specifically at the risk-sharing channels reveals that each channel's estimated β

coefficients provide only a partial picture of each channel's smoothing role. For instance, although the β's for the migration and fiscal transfer channels are similar (0.12 vs. 0.14), removing fiscal transfers raises the consumption multiplier by 40% upon impact, whereas removing migration only increases it by 16%, indicating that fiscal transfers are more than twice as important as migration in reducing consumption fluctuations. Similarly, shutting down credit markets increases consumption by 11%, even though the corresponding β is only one-third that of migration (0.04 vs. 0.12). Shutting down capital markets produces the largest effect, raising consumption by 52%.

The medium-run response (right panel of Figure 4) further demonstrates that the relevance of different channels evolves over time. In the short run, fiscal transfers and credit markets play a more prominent role, whereas migration becomes increasingly important in the medium run.[23]

Figure 5 provides additional information on these counterfactuals by showing for each counterfactual the response of the other risk-sharing channels, together with that of output and consumption. In general, eliminating one channel tends to raise the remaining channels' β's, partially compensating for the reduction in overall risk sharing. For instance, in the absence of fiscal transfers, households respond to state-specific shocks by migrating across U.S. state borders. This dynamic is reflected in the migration channel's counterfactual risk-sharing coefficient: Whereas a 1% rise in output increases population by 0.12% in the baseline, this figure rises to 0.14% in the absence of fiscal transfers. Despite these compensatory effects, the part of income fluctuations that remain unsmoothed rises when fiscal transfers are eliminated: In the baseline, 14% of income fluctuations are smoothed out through fiscal transfers and 47% of income fluctuations remain unsmoothed. Reducing the fiscal channel to zero raises the part that remains unsmoothed by eleven percentage points to 55%. While these compensatory responses of the remaining risk-sharing channels partly undo the loss in risk sharing, the overall effect on consumption is generally larger than predicted by the direct effect.

7 Conclusion

We argue that risk-sharing β's underestimate the benefits of risk sharing because they ignore the general equilibrium effects on income volatility. Quantitatively, these indirect benefits can be large. Calibrating a DSGE model of a small open economy to match these empirical findings we show that the observed level of risk sharing between U.S. states cuts state-level consumption volatility by a factor of 3.5, with 40% of the reduction being due to indirect, general equilibrium effects.

While we do not perform a quantitative welfare analysis, our finding that risk sharing affects

[23]This growing role of migration as an adjustment mechanism is consistent with findings in, e.g., Foschi et al. (2022).

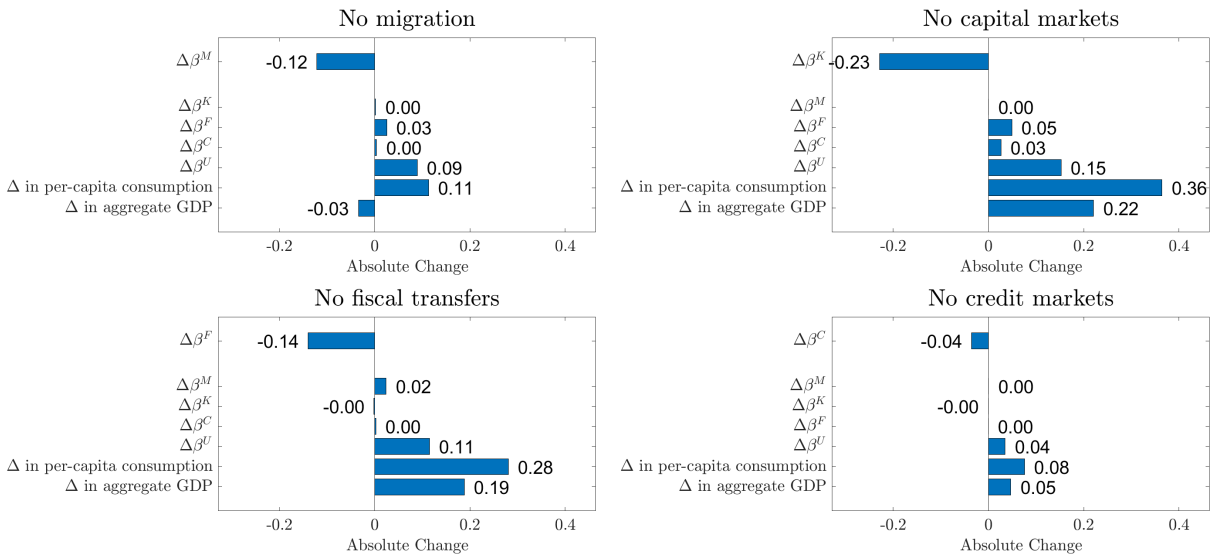


Figure 5: RISK-SHARING β'S FOR COUNTERFACTUALS

Notes: Risk-sharing β's and per-capita nominal consumption multipliers are reported for alternative model calibrations. The notes to Figure 3 explain how the various β's and the consumption multiplier are estimated from the simulated data. In the counterfactuals, eliminating migration implies $\gamma \to 0$, eliminating capital markets implies $\kappa = 1$, eliminating fiscal transfers implies $\tau = 0$, eliminating credit markets implies $\iota \to \infty$, and frictionless credit markets imply $\iota \to 0$.

consumption volatility indirectly through its effect on output volatility is important for such an exercise. First, consumption volatility directly shows up in households' welfare. Finding that the indirect effects of risk sharing amplify its effect on consumption volatility should therefore lead to larger welfare gains from risk sharing. Second, higher output volatility is typically associated with lower welfare due to first-order effects, especially if it goes along with higher inflation volatility. The New Keynesian literature emphasizes that higher inflation volatility reduces average output and hence, consumption, leading to first-order welfare effects (Woodford, 2003; Galí, 2008). More recently, several papers have proposed models in which output volatility directly reduces average output, e.g. through downward nominal wage rigidity (Schmitt-Grohé and Uribe, 2016; Dupraz, Nakamura and Steinsson, 2023). Taking these first-order effects into account would suggest that the potential gains from risk sharing are substantially higher than what is implied by a reduction of consumption volatility.

While our empirical and quantitative analysis focus on U.S. states, our conclusions are relevant for other currency unions as well, in particular the euro area. As shown in the introduction, the euro area does not only feature a higher consumption volatility, but also higher output volatility, in line with our model predictions. Disentangling whether this higher output volatility is due to less risk

sharing or a higher volatility of the exogenous shock processes is an important question for future research. Our paper provides a relevant framework to analyze this question.

References

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales. 2019. "Shift-share designs: Theory and inference." *The Quarterly Journal of Economics*, 134(4): 1949–2010.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro. 2020. "On the informativeness of descriptive statistics for structural estimates." *Econometrica*, 88(6): 2231–2258.

Angrist, Joshua D, and Peter Hull. 2023. "Instrumental variables methods reconcile intention-to-screen effects across pragmatic cancer screening trials." *Proceedings of the National Academy of Sciences*, 120(51): e2311556120.

Angrist, Joshua, Peter Hull, and Christopher Walters. 2023. "Methods for measuring school effectiveness." *Handbook of the Economics of Education*, 7: 1–60.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. "New trade models, same old gains?" *American Economic Review*, 102(1): 94–130.

Armington, Paul S. 1969. "A Theory of Demand for Products Distinguished by Place of Production (Une théorie de la demande de produits différenciés d'après leur origine)(Una teoría de la demanda de productos distinguiéndolos según el lugar de producción)." *Staff Papers-International Monetary Fund*, 159–178.

Artuç, Erhan, Shubham Chaudhuri, and John McLaren. 2010. "Trade shocks and labor adjustment: A structural empirical approach." *American Economic Review*, 100(3): 1008–45.

Asdrubali, Pierfederico, Bent E Sørensen, and Oved Yosha. 1996. "Channels of interstate risk sharing: United States 1963–1990." *Quarterly Journal of Economics*, 111(4): 1081–1110.

Backus, David K., and Gregor W Smith. 1993. "Consumption and Real Exchange Rates in Dynamic Economies with Non-Traded Goods." *Journal of International Economics*, 35(3-4): 297–316.

Balli, Faruk, Bent E Sørensen, et al. 2006. "The impact of the EMU on channels of risk sharing between member countries." 399–429.

Becker, Sascha O, and Mathias Hoffmann. 2006. "Intra-and international risk-sharing in the short run and the long run." *European Economic Review*, 50(3): 777–806.

Beraja, Martin. 2022. "Does US Fiscal Integration Stabilize Regional Business Cycles?" Working paper, Dept. Econ., Mass. Inst. Tech.

Bilbiie, Florin O. 2024. "Monetary policy and heterogeneity: An analytical framework." *Review of Economic Studies*, rdae066.

Boehm, Christoph E, Andrei A Levchenko, and Nitya Pandalai-Nayar. 2020. "The long and short (run) of trade elasticities." National Bureau of Economic Research.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-experimental shift-share research designs." *The Review of economic studies*, 89(1): 181–213.

Briganti, Edoardo, and Victor Sellemi. 2023. "Why Does GDP Move Before Government Spending? It's all in the Measurement." *UCSD Manuscript*.

Broda, Christian, Joshua Greenfield, and David Weinstein. 2006. "From Groundnuts to Globalization: A Structural Estimate of Trade and Growth." NBER Working paper No. 12512.

Brunet, Gillian. 2022. "When Does Government Spending Matter? Evidence from a New Measure of Spending." 41.

Chari, V. V, Patrick J. Kehoe, and Ellen R. McGrattan. 2002. "Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?" *The Review of Economic Studies*, 69(3): 533–563.

Chodorow-Reich, Gabriel. 2019. "Geographic cross-sectional fiscal spending multipliers: What have we learned?" *American Economic Journal: Economic Policy*, 11(2): 1–34.

Cochrane, John H. 1991. "A simple test of consumption insurance." *Journal of political economy*, 99(5): 957–976.

Devereux, Michael B, and James Yetman. 2014. "Capital controls, global liquidity traps, and the international policy trilemma." *The Scandinavian Journal of Economics*, 116(1): 158–189.

Dupor, Bill, and Rodrigo Guerrero. 2017. "Local and aggregate fiscal policy multipliers." *Journal of Monetary Economics*, 92: 16–30.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson. 2023. "A plucking model of business cycles." National Bureau of Economic Research.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, geography, and trade." *Econometrica*, 70(5): 1741–1779.

Erceg, Christopher J, Dale W Henderson, and Andrew T Levin. 2000. "Optimal monetary policy with staggered wage and price contracts." *Journal of Monetary Economics*, 46(2): 281–313.

Erickson, Timothy, and Toni M Whited. 2002. "Two-step GMM estimation of the errors-in-variables model using high-order moments." *Econometric Theory*, 18(3): 776–799.

Evers, Michael P. 2015. "Fiscal federalism and monetary unions: A quantitative assessment." *Journal of International Economics*, 97(1): 59–75.

Flood, Robert P, Nancy P Marion, and Akito Matsumoto. 2012. "International risk sharing during the globalization era." *Canadian Journal of Economics/Revue canadienne d'économique*, 45(2): 394–416.

Foschi, Andrea, Christopher L. House, Christian Proebsting, and Linda L. Tesar. 2022. "Moving to opportunity within and across borders."

Friedman, Milton. 1953. "The case for flexible exchange rates." *Essays in Positive Economics*, 157–203. University of Chicago Press.

Galí, Jordi. 2008. *Monetary policy, inflation, and the business cycle: An introduction to the New Keynesian framework.* Princeton University Press (Princeton, New Jersey: 2008).

Gali, Jordi, and Tommaso Monacelli. 2005. "Monetary policy and exchange rate volatility in a small open economy." *The Review of Economic Studies*, 72(3): 707–734.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik instruments: What, when, why, and how." *American Economic Review*, 110(8): 2586–2624.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz. 2021. "Aggregate nominal wage adjustments: New evidence from administrative payroll data." *American Economic Review*, 111(2): 428–71.

Guren, Adam, Alisdair McKay, Emi Nakamura, and Jón Steinsson. 2021. "What do we learn from cross-regional empirical estimates in macroeconomics?" *NBER Macroeconomics Annual*, 35(1): 175–223.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2017. "Optimal tax progressivity: An analytical framework." *The Quarterly Journal of Economics*, 132(4): 1693–1754.

Hoffmann, Mathias. 2008. "The lack of international consumption risk sharing: can inflation differentials and trading costs help explain the puzzle?" *Open Economies Review*, 19: 183–201.

Hoffmann, Mathias, Egor Maslov, Bent E Sørensen, and Iryna Stewen. 2019. "Channels of risk sharing in the eurozone: What can banking and capital market union achieve?" *IMF Economic Review*, 67(3): 443–495.

House, Christopher L, Christian Proebsting, and Linda L Tesar. 2023. "Quantifying the benefits of labor mobility in a currency union." National Bureau of Economic Research.

Ingram, James C. 1962. "A proposal for financial integration in the Atlantic community." 118.

Kalemli-Ozcan, Sebnem, Bent E Sørensen, and Oved Yosha. 2003. "Risk sharing and industrial specialization: Regional and international evidence." *American Economic Review*, 93(3): 903–918.

Karabarbounis, Loukas, and Brent Neiman. 2013. "The global decline of the labor share." *Quarterly Journal of Economics*, 129(1): 61–103.

Kenen, Peter. 1969. "The Theory of Optimum Currency Areas: An Eclectic View." *Monetary Problems of the International Economy*, , ed. R.A. Mundell and A.K. Swoboda, 41–60. University of Chicago Press.

Kohler, Wilhelm, Gernot J Müller, and Susanne Wellmann. 2023. "Risk sharing in currency unions: The migration channel." *European Economic Review*, 158: 104540.

Lucas Jr, Robert E. 1990. "Liquidity and interest rates." *Journal of economic theory*, 50(2): 237–264.

Mace, Barbara J. 1991. "Full insurance in the presence of aggregate uncertainty." *Journal of Political Economy*, 99(5): 928–956.

Maggiori, Matteo. 2022. "International macroeconomics with imperfect financial markets." In *Handbook of international economics*. Vol. 6, 199–236. Elsevier.

McKinnon, Ronald I. 1963. "Optimum Currency Areas." *The American Economic Review*, 717–725.

Melitz, Marc J, and Gianmarco IP Ottaviano. 2008. "Market size, trade, and productivity." *Review of Economic Studies*, 75(1): 295–316.

Montiel Olea, José Luis, and Carolin Pflueger. 2013. "A robust test for weak instruments." *Journal of Business & Economic Statistics*, 31(3): 358–369.

Mundell, Robert A. 1961. "A theory of optimum currency areas." *American Economic Review*, 51(4): 657–665.

Mundell, Robert A. 1973. "Uncommon arguments for common currencies." *The economics of common currencies*, 114–132.

Nakamura, Emi, and Jon Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from US Regions." *The American Economic Review*, 104(3): 753–792.

Nekarda, Christopher J, and Valerie A Ramey. 2020. "The cyclical behavior of the price-cost markup." *Journal of Money, Credit and Banking*, 52(S2): 319–353.

Pennings, Steven. 2016. "Cross-regions transfers in a monetary union: evidence from the US and implications for Europe." *Unpublished, World Bank*, 5.

Schmitt-Grohé, Stephanie, and Martın Uribe. 2003. "Closing small open economy models." *Journal of International Economics*, 61(1): 163–185.

Schmitt-Grohé, Stephanie, and Martin Uribe. 2016. "Downward Nominal Wage Rigidity, Currency Pegs, and Involuntary Unemployment." *Journal of Political Economy*, 124(5): 1466–1514.

Sørensen, Bent E, and Oved Yosha. 1998. "International risk sharing and European monetary unification." *Journal of International Economics*, 45(2): 211–238.

Townsend, Robert M. 1994. "Risk and insurance in village India." *Econometrica: journal of the Econometric Society*, 539–591.

Wolf, Christian K. 2023. "The missing intercept: A demand equivalence approach." *American Economic Review*, 113(8): 2232–2269.

Woodford, Michael. 2003. *Interest and prices: Foundations of a theory of monetary policy.* Cambridge University Press.