Journal of Monetary Economics 147 (2024) 103579

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

A HANK^2 model of monetary unions☆

Christian Bayer [a,b,c,d], Alexander Kriwoluzky [e,f], Gernot J. Müller [g,b,c,*], Fabian Seyrich [e,f]

[a] University of Bonn, Germany
[b] CEPR, UK
[c] CESifo, Germany
[d] IZA, Germany
[e] Freie Universität Berlin, Germany
[f] DIW Berlin, Germany
[g] University of Tübingen, Germany

ARTICLE INFO

JEL classification:
F45
E52
D31

Keywords:
OCA theory
Two-country model
Monetary union
Monetary policy
Household heterogeneity
Inequality

ABSTRACT

How does a monetary union alter the impact of business cycle shocks at the household level? We develop a Heterogeneous Agent New Keynesian model of two countries (HANK^2) and show in closed form that a monetary union shifts the adjustment to a shock horizontally across countries, within the brackets of the union-wide wealth distribution, rather than vertically, that is, across the brackets of the union-wide wealth distribution. Calibrating the model to the euro area reveals that a monetary union alters the impact of shocks most strongly in the tails of the wealth distribution but leaves the middle class almost unaffected.

1. Introduction

Following the seminal work of Mundell (1961), Optimum Currency Area theory analyzes the costs and benefits of monetary unions at the level of regions or countries. Likewise, the policy debate is framed in such terms, as the 20-year plus history of the euro illustrates: discussions of whether specific *countries* would have been better off without the euro abound. Heterogeneity is at the heart of the issue: if countries differ, say, because of country-specific shocks, one (monetary) policy does not fit all. However, heterogeneity across households–in terms of income, wealth and shocks–dwarfs the heterogeneity across countries. Hence, we offer a change of perspective. We focus on households rather than countries and ask: How does a monetary union alter the impact of business cycle shocks at the household level?

To answer this question, we propose a Heterogeneous Agent, New Keynesian model of two countries: HANK^2. The model features incomplete markets, idiosyncratic risk, and self-insurance in a standard New Keynesian two-country setup. It is therefore able to

☆ We thank an anonymous referee, our discussants Tommaso Monacelli and Vincent Sterk as well as Ralph Luetticke for insightful comments. Luis Huxel provided valuable research assistance. This research has received financial support by the Leibniz Association, Germany through the project "Distributional effects of macroeconomic policies in Europe". Christian Bayer gratefully acknowledges support by the DFG, Germany through CRC-TR224 EPoS, subarea C05. The usual disclaimer applies.

* Corresponding author at: University of Tübingen, Germany.
E-mail addresses: christian.bayer@uni-bonn.de (C. Bayer), akriwoluzky@diw.de (A. Kriwoluzky), gernot.mueller@uni-tuebingen.de (G.J. Müller), fabian.seyrich@gmail.com (F. Seyrich).

https://doi.org/10.1016/j.jmoneco.2024.103579
Received 1 March 2024; Received in revised form 23 March 2024; Accepted 25 March 2024
Available online 27 March 2024

capture key features of the business cycle and the wealth distribution; and we may use it to analyze how a monetary union alters–at three levels of aggregation–the effects of country-specific business-cycle shocks relative to a scenario where independent monetary policies are in place. First, we show that when aggregated across both countries and households, macroeconomic dynamics are independent of whether there is a monetary union or not. Second, the overall effect of the shock for households in specific brackets of the wealth (and income) distribution, aggregated across countries of residence, does not depend on whether monetary union is in place or not. In other words, the monetary union itself does not shift the impact of the shock *vertically* across wealth classes. Instead, and this is our third and main result, it shifts the impact of the shock *horizontally* across borders within the brackets of the wealth distribution. Quantitatively, this shift is strongest in the tails of the distribution.

We first study a small-scale version of the model similar to Auclert et al. (2021b) where we make a number of simplifying assumptions. These assumptions allow us to obtain strong analytical knife-edge results. First, we restrict the two countries in the model to be symmetric, except for the occurrence of country-specific shocks. What is more, we abstract from capital accumulation and allow for trade in liquid, one-period debt only. Further, we abstract from price stickiness and maintain wage rigidities as the only nominal friction. A key difference to Auclert et al. (2021b) is that we consider a two-country model rather than a small-open economy. Relying on a sequence-space representation as in Auclert et al. (forthcoming) or McKay and Wolf (2022), we then show that it is possible to cast the union-wide dynamics of the simplified model into the canonical form which is familiar from the textbook-version of the Representative Agent New Keynesian (RANK) model (Galí, 2015).

To assess the quantitative relevance of the analytical results, we consider a richer version of the model. Specifically, we introduce capital formation, portfolio choice, and price rigidities building on the medium-scale HANK model in Bayer et al. (forthcoming). We calibrate this version of the model to capture key features of (and asymmetries between) the Italian economy and the German one at the household level. At this level, Italy and Germany are very different, notably in terms of wealth inequality. According to a number of indicators, wealth inequality is significantly higher in Germany. We show that the model is able to account for these structural differences. At the macro level, we maintain the assumption of identical frictions. These have been the subject of earlier research, which we review below. Still, we verify that the model is able to capture key aspects of the business cycle, including its co-movement across Italy and Germany. Importantly, we find that the results that we establish for the simplified model approximately hold for the medium-scale HANK model, even though household-level heterogeneity differs considerably across countries.

In more detail, we present our main results as we put forward three propositions. First, we show that whether countries operate a monetary union or independent monetary policies makes no difference for how country-specific shocks play out at the union level: Monetary union is irrelevant for union-wide dynamics. This result holds exactly for a first-order approximation in the aggregate states and under the assumption that countries are perfectly symmetric (except for the incidence of shocks). Moreover, we require that the monetary union is designed in such a way that the (implied) monetary policy rule for the union-wide interest rate does not differ from the case of independent monetary policies. This holds if we assume–in line with actual practice in the EA–that the common monetary policy adjusts the policy rate to the average inflation rate (and possibly the output gap) in both countries, to which we refer as "Home" and "Foreign". Against this background the irrelevance result is intuitive. Under a monetary union, monetary policy does not fit all: relative to a benchmark with independent monetary policies, the common policy responds too much in one country and too little in the other. It follows that macro dynamics at the country level do very much depend on whether countries operate a monetary union or not. But when countries are symmetric the changes induced by the monetary union in both countries offset each other such that union-wide dynamics do not change with monetary union.

Our second proposition concerns the household level. Taking a union-wide perspective and aggregating households across countries of residence within the brackets of the wealth distribution, we find that the impact of a business cycle shock for specific wealth classes does not depend on whether there is a monetary union or not. Put differently, just like with union-wide aggregate dynamics, monetary union is also irrelevant for the impact of shocks along the union-wide wealth distribution. It does not, say, shift the adjustment vertically from the rich to the poor or vice versa. Intuitively, how saving and consumption at the household level change in response to shocks within one country does depend on whether there is a monetary union in place or not because these depend on the price adjustments within that country. But given that a monetary union is irrelevant for union-wide price paths, it follows that the changes that a monetary union induces for the response of a generic household in Home are perfectly offset by the changes of its "twin" in Foreign—where a twin is defined in terms of its location in the income and wealth space. Aggregating across countries of residence, we thus find the overall adjustment unchanged within the brackets of the wealth distribution.

Our third proposition is implicit in the argument above. It establishes that monetary union is potentially very relevant for the impact of country-specific shocks along the wealth distribution within a country. More specifically, comparing the outcome under a monetary union to the outcome under independent monetary policies, we observe that the union shifts the impact of shocks horizontally across borders within the brackets of the wealth distribution. Put differently, the impact of shocks changes for specific households at the expense of their twins in the other country: in the face of specific shocks, the poor (rich) in one country benefit from union membership at the expense of the poor (rich) in the other country. Hence, monetary union makes a difference for how shocks impact the rich and the poor within a country.

We simulate the calibrated model and verify that our main results hold approximately once we allow for asymmetries in terms of household-level heterogeneity across countries. In particular, as we study the adjustment to country-specific shocks we find–consistent with Proposition 1–that union-wide aggregate dynamics are basically independent of whether there is a monetary union in place or not. In contrast, country-level dynamics change fundamentally due to the monetary union. In this regard, our model simulation confirms the classic notion that one size does not fit all. The response of Gini coefficients in our simulations suggests that Proposition 2 also holds approximately in the asymmetric model.

Last, we perform a quantitative analysis that relates to Proposition 3. We compute the consumption equivalent welfare variation of a shock as a comprehensive (ex-post) measure of its impact and find that a monetary union induces strong changes in this measure in the tails of the wealth distribution, both in Home and Foreign. These changes can be traced back to how a monetary union changes the interest-rate dynamics to which households in the tails of the wealth distribution–rich and poor–are more exposed than the middle class which neither borrows nor saves much (in excess of what it implicitly owes through government debt). We find accordingly, that monetary union does not change the impact of shocks for the middle class. This result offers a fresh perspective on the euro. During its 20-year-plus history, the euro area witnessed various political movements in several of its member states that campaigned against the euro, yet their appeal to the electorate turned out to be limited. Our analysis offers an explanation for why this occurs.

Related literature. Our analysis builds on two earlier generations of OCA theory. The first generation stresses that countries should be sufficiently homogeneous to qualify as an OCA. The original contribution of Mundell (1961) emphasizes that economic regions as opposed to nation-states or countries are the relevant category when it comes to operating a common currency. We thus follow Mundell's lead as we attempt to shift the focus away from countries (and towards households). Other contributions to the first generation of OCA theory stress the role of trade openness and the asymmetry of shocks (McKinnon, 1963; Kenen, 1981; Bayoumi and Eichengreen, 1992; Krugman, 1993). Lastly, influential work has emphasized the potential endogeneity of the OCA criteria (Frankel and Rose, 1998; Rose, 2000).

The second generation of OCA theory zooms in on specific aspects, notably on the trade-offs faced by monetary and fiscal policy in monetary unions as well as on the conduct of optimal policy, relying on explicit welfare criteria (Beetsma and Uhlig, 1999; Alesina and Barro, 2002). These criteria are typically micro-founded within New Keynesian models featuring representative agents (see, for instance, Benigno, 2004; Kollmann, 2004; Benigno and López-Salido, 2006; Beetsma and Jensen, 2005; Corsetti, 2008; Gali and Monacelli, 2008; Galí and Monacelli, 2016; Farhi and Werning, 2017; Hettig and Müller, 2018; Groll and Monacelli, 2020).

The present paper belongs to a new set of studies that explicitly accounts for within-country heterogeneity when revisiting open-economy issues. In particular, several studies rely on small open-economy HANK models to reassess the merits of alternative exchange-rate policies. de Ferra et al. (2021) find that household heterogeneity rationalizes "fear of floating" in the face of sudden stops. Auclert et al. (2021b), in turn, stress that household heterogeneity can amplify the real income channel of exchange rates, potentially giving rise to contractionary depreciations. Guo et al. (2020) find that fixing the exchange rate leads to larger spillovers of foreign shocks but dampens their distributional impact, in contrast to what we find for HANK^2. Oskolkov (2023) and Zhou (2021) also study the distributional impact of foreign shocks and exchange-rate policies in small open-economy HANK models. Aggarwal et al. (2023) study the implications of fiscal deficits through the lens of a multi-country HANK model. Bellifemine et al. (2023) develop a HANK model of a monetary union composed of small open economies. What sets our paper apart is the two-country structure of HANK^2: it allows us to study how a monetary union alters the impact of shocks along the wealth distribution—both, vertically and horizontally across borders. Bayer et al. (2022) and Chen et al. (2023) also develop a two-country HANK models and calibrate them to the EA. They focus on fiscal frameworks rather than on monetary union as such. In Bayer et al. (2022), in particular, we develop the notion that "attitudes" towards fiscal policy may be traced back to how differences in income and wealth interact with different social security systems.

2. The model

We develop a two-country New Keynesian model with incomplete markets, idiosyncratic risk, and heterogeneous agents (HANK^2). In this section, we first introduce a smaller model, a one-asset-HANK^2 model, for which we are able to establish a number of closed-form results in Section 3. We extend the model in Section 4 to a two-asset, medium-scale-HANK^2 model and calibrate it to data for the EA in order to assess the quantitative relevance of our results.

We borrow our two-country framework from Corsetti et al. (2012), while the specification of the household problem follows the small-open economy setup of Auclert et al. (2021b). Their setup, in turn, extends Galí and Monacelli (2005) by allowing for household heterogeneity. Countries are isomorphic and our exposition focuses on the domestic economy or "Home". "Foreign" looks the same. Countries differ only in terms of shocks and in terms of size: We normalize the total population to unity, a fraction n of which resides in Home. In what follows, we denote foreign variables with the superscript * and use subscripts H and F to distinguish between domestic and foreign variables within a country. To benchmark the case of a monetary union against a scenario of independent monetary policies, we allow Home and Foreign to operate different currencies. In case there is a monetary union there will be an irrevocable conversion rate. We further assume that households and firms have perfect foresight and focus on a first-order approximation around the stationary equilibrium.

2.1. Households

There is a continuum of households, each of which faces idiosyncratic income risk. This, in turn, is due to idiosyncratic productivity, $e_{i,t}$, which is determined exogenously by a first-order Markov chain with mean $\mathbb{E}e_{i,t} = 1$. Households save via a riskless bond which is denominated in domestic currency and issued by a mutual fund which, in turn, holds government debt as well as foreign-currency bonds. This yields the familiar UIP condition but is otherwise inconsequential for the household savings decisions given perfect foresight.

Household labor supply, N_t, is determined by a labor union as described below and we assume that the labor union allocates hours worked uniformly across households. In what follows we state the household problem recursively, using time subscripts only

for aggregate variables exogenous to the individual decision problem (and the value function because it is time-dependent in the face of aggregate shocks). At time t, a generic household with bond holdings a and productivity level e chooses consumption, c, and savings a', by solving the dynamic program

$$V_t(a,e) = \max_{c,a'} u(c,N_t) + \xi_t \beta E_t[V_{t+1}(a',e')] \tag{1}$$

$$\text{s.t. } c + a' = (1+r_t^b)a + e\frac{W_t}{P_t}N_t - \tilde{\tau}_t e \tag{2}$$

$$a' \geq \underline{a},$$

P_t is the consumption price index specified below, r_t^b is the return on the bond, W_t the nominal wage, $0 < \beta < 1$ is the time discount factor, E_t the expectation operator, $\underline{a}$ an exogenous borrowing limit, and $\tilde{\tau}_t$ is a non-distortionary tax on households (it depends on e which is exogenous).[1] In addition, ξ_t is an impatience shock which we use to showcase how country-specific demand shocks will enter the IS relation later.

For now, we assume the functional form

$$u(c,N_t) = \frac{c^{1-\gamma}}{1-\gamma} - \psi\frac{N_t^{1+\varphi}}{1+\varphi},$$

where $\gamma, \varphi > 0$ are the inverse of the intertemporal elasticity of substitution and the inverse of the Frisch elasticity of labor supply respectively.

We state the solution to the household's problem in sequence-space form as in Auclert et al. (2021a), which we will use later to characterize the dynamics of the model. In particular, the solution to household's i consumption-savings problem described by (1) and (2) maps the time paths of wages, written in boldface to indicate vectors, $\boldsymbol{W}$, hours worked $\boldsymbol{N}$, real returns $\boldsymbol{r^b}$, taxes $\tilde{\boldsymbol{\tau}}$, prices $\boldsymbol{P}$, and shocks $\boldsymbol{\xi}$ to that of consumption of household i:

$$\boldsymbol{c}_i = C_i\left(\boldsymbol{W}/\boldsymbol{P}, \boldsymbol{N}, \boldsymbol{r^b}, \tilde{\boldsymbol{\tau}}, \boldsymbol{\xi}\right). \tag{3}$$

Aggregating across all domestic households, we obtain an aggregate domestic consumption function $C(\cdot)$, similar as in Auclert et al. (forthcoming) or McKay and Wolf (2022):

$$\boldsymbol{c} = C\left(\boldsymbol{W}/\boldsymbol{P}, \boldsymbol{N}, \boldsymbol{r^b}, \tilde{\boldsymbol{\tau}}, \boldsymbol{\xi}\right). \tag{4}$$

In each period, households allocate their consumption expenditures, c, across a domestically produced good c_H and an imported good c_F so as to enjoy overall consumption level

$$c = \left\{ [1-(1-n)\alpha_H]^{\frac{1}{\sigma}} c_H^{\frac{\sigma-1}{\sigma}} + [(1-n)\alpha_H]^{\frac{1}{\sigma}} c_F^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{\sigma}{1-\sigma}}. \tag{5}$$

Here $\alpha_H \in [0,1]$ indicates a home bias in consumption: The weight of the domestic good in total consumption is larger than what the size of the domestic economy would imply. If $\alpha_H = 1$ there is no home bias. σ is the elasticity of substitution between the domestic and the imported good. Letting $P_{H,t}$ and $P_{F,t}$ denote the price of these goods, both expressed in domestic currency, expenditure minimization implies for the consumer price index:

$$P_t = \left\{ [1-(1-n)\alpha_H] P_{H,t}^{1-\sigma} + [(1-n)\alpha_H] P_{F,t}^{1-\sigma} \right\}^{\frac{1}{1-\sigma}}. \tag{6}$$

The optimal intratemporal allocation of expenditures implies the demand functions:

$$c_H = (1-(1-n)\alpha_H)\left(\frac{P_{H,t}}{P_t}\right)^{-\sigma} c, \quad c_F = (1-n)\alpha_H\left(\frac{P_{F,t}}{P_t}\right)^{-\sigma} c.$$

Let $\mathcal{E}_t$ denote the nominal exchange rate, that is, the price of foreign currency expressed in terms of the domestic currency. We assume that the law of one price holds, that is, the foreign currency price of the domestically produced good is given by $P_{H,t}^* = \mathcal{E}_t P_{H,t}$ and likewise for the foreign-currency price of the imported good. For future reference, it is also useful to define the terms of trade as the relative price of foreign goods to domestic goods $s_t = P_{F,t}/P_{H,t}$ and the real exchange rate $Q_t = P_t \mathcal{E}_t / P_t^*$.

2.2. Production

The production function is linear in labor:

$$Y_t = N_t, \tag{7}$$

[1] Non-distortionary taxes simplify our arguments below but they are not necessary: our result also holds if taxes are standard distortionary labor taxes.

where N_t is the aggregate labor input. For now, we assume perfect competition in the domestic goods market such that the price of domestic goods is equal to marginal costs given by the nominal wage: $P_{Ht} = W_t$. It is convenient to rewrite the real wage as a function of the terms of trade:

$$w_t = \frac{W_t}{P_t} = \frac{P_{Ht}}{P_t} = [(1-(1-n)\alpha_H) + ((1-n)\alpha_H)s_t^{1-\sigma}]^{-\frac{1}{1-\sigma}}. \tag{8}$$

Aggregate labor is composed of differentiated types:

$$N_t = \left(\int_k N_{k,t}^{\frac{\epsilon_t - 1}{\epsilon_t}} dk \right)^{\frac{\epsilon_t}{\epsilon_t - 1}}, \tag{9}$$

where ϵ_t is the elasticity of substitution between labor types and may vary over time. We use this "cost-push shock" to showcase how country-specific shocks to the Philips curve affect the dynamics of the model. Labor types, in turn, are efficiency units of work: $N_{k,t} = \int e_{i,t} n_{i,k,t} di$, where i indexes a household, as before, and $k \in [0,1]$ indexes the labor type. As in the recent literature, we assume that the number of hours a household works as type k, $n_{k,t}$, is determined by a union, which also determines the wage for each type $W_{k,t}$ (Erceg et al., 2000; Auclert et al., forthcoming; McKay and Wolf, 2022). A union can reset the wage with a constant probability θ.

The solution to the union problem yields a standard linearized open-economy Philips curve:

$$\hat{\pi}_{H,t} = \kappa((1-n)\alpha_H \hat{s}_t + \varphi \hat{Y}_t + \gamma \hat{C}_t) + \beta \hat{\pi}_{H,t+1} + \psi \hat{\epsilon}_t, \tag{10}$$

where $\pi_{H,t} := \frac{p_{H,t}}{p_{H,t-1}}$ is gross domestic producer price inflation, $\kappa \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta}$, C_t is aggregate consumption, and $\psi \equiv -\frac{\kappa}{(\bar{\epsilon}-1)}$; and a "$\hat{}$" denote the log deviation of a variable from its steady state value.[2]

2.3. Financial markets

There are two bonds, a home bond and a foreign bond, each denoted in the country's own currency (which is identical in case there is monetary union). In the absence of arbitrage, the expected returns on both bonds are equal which implies the standard uncovered interest parity (UIP) condition:

$$1 + i_t = (1 + i_t^*)\frac{\mathcal{E}_{t+1}}{\mathcal{E}_t}, \tag{11}$$

with i_t being the nominal interest rate. In order to abstract from potentially heterogeneous household portfolios we assume that bond trading takes place via the mutual fund. Generally, up to the first order, the fund's portfolio is indeterminate. We assume that, in the steady state, the fund only holds domestic-currency debt. In this way, we rule out valuation effects that may arise in response to shocks under flexible exchange rates. Off steady-state, without loss of generality, we assume that cross-border trade is restricted to the foreign bond.

The domestic real interest rate is then pinned down by the Fisher equation and given by:

$$1 + r_t = \frac{1 + i_t}{1 + \pi_{t+1}}, \tag{12}$$

where $\pi_{t+1} := \frac{P_t}{P_{t-1}}$ is domestic CPI inflation. Given our assumption that the fund only holds domestic-currency debt in steady state, $r_t^b = r_t$ holds up to first-order.

2.4. Monetary and fiscal policy

In case there is a monetary union the common central bank adjusts interest rates based on the following simple rule:

$$i_t = \theta_\pi \left(n\pi_{Ht} + (1-n)\pi_{Ft}^* \right). \tag{13}$$

where π_{Ht} is producer price inflation at Home while π_{Ft}^* is producer price inflation in Foreign. The coefficient $\theta_\pi \geq 0$ governs the extent to which the central bank adjusts the policy rate in response to average inflation in the monetary union. In the expression above, we assume that the rate is adjusted in response to producer price inflation but in our setup "targeting" the CPI inflation is equivalent. Our results below also extend to the case where interest rates are adjusted to the (average) output gap. Note also that UIP (11) implies that $i_t = i_t^*$ once the nominal exchange rate is irrevocably fixed.

Alternatively, we consider a case with flexible exchange rates, assuming the following rules for monetary policy in Home:

$$i_t = \theta_\pi \pi_{Ht}, \tag{14}$$

and symmetrically for Foreign:

$$i_t^* = \theta_\pi \pi_{Ft}^*. \tag{15}$$

[2] We assume that the union neglects the impatience shock of households when setting wages in order to study the role of distinct demand and supply shocks.

The conduct of fiscal policy is independent of whether there is a monetary union or not. It is set at the national level. Each government issues government bonds B_t to finance deficits and sets the tax rate. The budget constraint of the national fiscal policy reads as follows:

$$\frac{1+i_t}{\pi_t}B_t = B_{t+1} + \tau_t w_t N_t, \tag{16}$$

with $\tau_t = \frac{\bar{\tau}_t}{w_t N_t}$. We assume that tax rates adjust to stabilize the level of government debt:

$$\frac{\tau_t}{\bar{\tau}} = \left(\frac{B_{t+1}}{\bar{B}}\right)^{\gamma_B^\tau}, \tag{17}$$

where γ_B^τ governs the speed with which debt returns to its target value $\bar{B}$.

2.5. Market clearing

Bond market clearing requires:

$$A_{t+1} = B_{t+1} + \frac{B_{F,t+1}}{Q_t}, \tag{18}$$

that is, the total amount of domestic savings, $A_{t+1} \equiv \int_0^n a_{i,t+1}di$, equals the domestic bonds plus the net foreign asset position, $B_{F,t+1}$, which is held in foreign bonds. Analogously bond market clearing requires for Foreign:

$$A^*_{t+1} = B^*_{t+1} - \frac{n}{1-n}B_{F,t+1}. \tag{19}$$

Aggregating over the domestic households' budget constraints gives the net amount of domestic holdings of foreign bonds, $B_{F,t}$:

$$w_t Y_t - T_t + (1+r_t)B_t + \frac{(1+r^*_t)}{Q_t}B_{F,t} = C_t + B_{t+1} + \frac{B_{F,t+1}}{Q_t}. \tag{20}$$

Finally, goods markets clearing requires:

$$Y_t = \left(p_{Ht}\right)^{-\sigma}\left[(1-(1-n)\alpha_H)C_t + (1-n)\alpha_H Q_t^{-\sigma}C^*_t\right] \tag{21}$$

$$Y^*_t = \left(p^*_{Ft}\right)^{-\sigma}\left[n\alpha_H Q_t^{\sigma}C_t + (1-n\alpha_H)C^*_t\right]. \tag{22}$$

In Online Appendix A.1 we provide a formal definition of a linearized perfect-foresight equilibrium for which we derive results in the following section.

3. Closed-form results

In this section, we derive our main results in closed form. In particular, we show that a monetary union shifts the impact of country-specific shocks at the household level horizontally, that is, across borders within the brackets of the wealth distribution. We abstract from union-wide shocks because, given our assumptions on symmetry, the one-size-fits-all problem arises only in the face of country-specific shocks.

To set the stage, we first derive two propositions that show that a monetary union makes no difference for union-wide outcomes, both in terms of how aggregate variables respond to country-specific shocks and how the impact of the shock spreads vertically across the brackets of the union-wide wealth distribution. Put differently, it is irrelevant to union-wide outcomes whether countries form a monetary union or not. At the same time, a monetary union alters the adjustment to country-specific shocks across borders—both, at the aggregate level and at the household level.

For what follows, we define union-wide variables as a weighted average of the realizations in Home and Foreign, $X_t^W = nX_t + (1-n)X^*_t$, and write the canonical form for union-wide dynamics using the sequence-space representation (see Online Appendix A.2.1 for details).[3] As with the textbook representative agent version of the New Keynesian model, the canonical form is sufficient to describe the aggregate dynamics of the economy. Specifically, we summarize inflation dynamics with a union-wide New Keynesian Phillips curve:

$$\widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} = \kappa\widehat{\boldsymbol{y}^{\boldsymbol{W}}} + \beta\widehat{\boldsymbol{\pi}^{\boldsymbol{W}}_{+\mathbf{1}}} + \psi\widehat{\boldsymbol{\eta}^{\boldsymbol{W}}}, \tag{23}$$

where η^W is a sequence of cost shocks. The union-wide IS relation, in turn, is given by:

$$\widehat{\boldsymbol{y}^{\boldsymbol{W}}} = \tilde{C}_y\widehat{\boldsymbol{y}^{\boldsymbol{W}}} + \tilde{C}_i\widehat{\boldsymbol{i}^{\boldsymbol{W}}} + \tilde{C}_\pi\widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} + \tilde{C}_\xi\widehat{\boldsymbol{\xi}^{\boldsymbol{W}}}. \tag{24}$$

Importantly, (23) and (24), hold independently of whether there is a monetary union or not. To close the model, we need to specify a rule that pins down the union-wide interest rate $\widehat{\boldsymbol{i}^{\boldsymbol{W}}}$. This is where the monetary union comes into play. However, given that

[3] For lack of a better term, we also refer to these variables as "union-wide" variables even if the two countries operate independent monetary policies and let the exchange rate float.

the interest-rate rule for the union has the same functional form as the rules under monetary independence–except that it targets weighted average inflation–it is irrelevant for the dynamics of the union-wide interest rate whether or not the two countries operate a monetary union.

Proposition 1. *The union-wide aggregate dynamics are characterized by (23) and (24) and a mapping from aggregate union-wide inflation to aggregate union-wide policy rates. Because under the assumptions in Section 2.4 above this mapping is the same in a monetary union and with independent monetary policies, so are aggregate dynamics.*

Proof. With a monetary union, we have $i_t = i_t^* = i_t^W$ and given Eq. (13), we have:

$$i_t^W = \theta_\pi(n\pi_{H,t} + (1-n)\pi_{F,t}^*) = \theta_\pi \pi_t^W .$$

With independent monetary policies, given by Eqs. (14) and (15), the union-wide interest rate is:

$$i_t^W = n i_t + (1-n) i_t^* = n(\theta_\pi \pi_{H,t}) + (1-n)(\theta_\pi \pi_{F,t}^*) = \theta_\pi \pi_t^W \tag{25}$$

and, hence, *exactly the same as with a monetary union.* □

Proposition 1 implies that if a monetary union experiences country-specific shocks, union-wide aggregates like output, consumption, and inflation behave exactly the same independently of whether countries form a monetary union or not. To see why, consider a shock originating in Foreign. In a monetary union, the response of monetary policy is a response to the weighted average of the dynamics in both countries. This implies, for instance, that monetary policy reacts "too much" from the perspective of Home and "too little" from the perspective of Foreign, compared to what would happen under independent policies. But given that the countries are isomorphic–in particular given $(\kappa, C, \theta_\pi) = (\kappa^*, C^*, \theta_\pi^*)$–"too little" and "too much" means the same in absolute value and, thus, the contribution of each country to union-wide dynamics exactly offsets each other. Note that this holds even if the countries are not of the same size, in which case the size-weighted absolute value would be the same.

From a union-wide perspective, monetary union is also irrelevant to the impact of shocks along the wealth distribution. To see this, consider a generic household j in Home. Given symmetry, there are $\frac{n}{1-n}$ times identical households in Foreign, that is, households with the same idiosyncratic productivity and the same wealth. We label these twin households j^* and define $c_J = n c_j + (1-n) c_{j^*}$ as aggregate consumption of household j and its twins. Note that in linearized form, we have

$$\hat{c}_j = C_{w,j}\hat{\boldsymbol{w}} + C_{N,j}\hat{\boldsymbol{N}} + C_{i,j}\hat{\boldsymbol{i}} + C_{\pi,j}\hat{\boldsymbol{\pi}} + C_{\tau,j}\hat{\boldsymbol{\tau}} + C_{\xi,j}\hat{\boldsymbol{\xi}} \tag{26}$$

$$\widehat{c^*_{j^*}} = C^*_{w^*,j^*}\widehat{\boldsymbol{w}^*} + C^*_{N^*,j^*}\widehat{\boldsymbol{N}^*} + C^*_{i^*,j^*}\widehat{\boldsymbol{i}^*} + C^*_{\pi^*,j^*}\widehat{\boldsymbol{\pi}^*} + C^*_{\tau^*,j^*}\widehat{\boldsymbol{\tau}^*} + C^*_{\xi^*,j^*}\widehat{\boldsymbol{\xi}^*}. \tag{27}$$

Given symmetry, $C_{x,j} = C^*_{x^*,j^*}$ and, thus:

$$\hat{c}_J = C_{w,j}\widehat{\boldsymbol{w}^W} + C_{N,j}\widehat{\boldsymbol{N}^W} + C_{i,j}\widehat{\boldsymbol{i}^W} + C_{j,\pi}\widehat{\boldsymbol{\pi}^W} + C_{j,\tau}\widehat{\boldsymbol{\tau}^W} + C_{\xi,j}\widehat{\boldsymbol{\xi}^W}. \tag{28}$$

The same logic applies to all policy functions of the household. Given Proposition 1, the inputs of the aggregate policy functions of the twin households do not depend on whether there is a monetary union or not. Hence, the weighted average (or union-wide aggregate) of a choice variable of household j in Home and its $\frac{n}{1-n}$ twins in Foreign does therefore not depend on whether there is a monetary union or not. Our next irrelevance result follows directly:

Proposition 2. *The impact of country-specific shocks along the union-wide wealth and income distribution is independent of whether two countries form a monetary union or not. In other words, the monetary union does not alter the impact of the shock vertically.*

Against this background, the next proposition follows directly. It summarizes our main result.

Proposition 3. *Monetary union shifts the distributional impact of country-specific shocks horizontally across borders within the brackets of the wealth distribution.*

To see what drives this result, recall from our arguments above that a monetary union alters the dynamics of *country-specific* variables relative to what would be observed under independent monetary policies. This means that the arguments that feature in the consumption function of individual households in Home and Foreign, (26) and (27), generally differ compared to what would be the case with independent monetary policies. Also, the consumption choice of a household with a given wealth and productivity state in Home will generally differ from that of its twin in Foreign. Yet, as established in Proposition 2, how the *union-wide* wealth distribution changes in response to country-specific shocks does not depend on the monetary union (because it does not shift the impact of the shock vertically). Assuming countries are of the same size, this then requires that monetary union changes the effect of a shock on a household's consumption choice in Home in exactly the opposite way as it does for its Foreign twin. When countries differ in size, the differential impact of monetary union on the choice of a generic household in Home is of the opposite sign as that of its Foreign twin, weighted by the number of twins that a Home household has in Foreign. It follows that the distributional effect of monetary union operates horizontally across borders: It shifts the distributional impact of shocks (compared to a scenario of independent monetary policy) between households in Home and Foreign *with the same individual states* or within the same bracket of the wealth distribution. For instance, if consumption of the poor at Home is higher with a monetary union in place than with independent monetary policies after a given shock, consumption of the poor in Foreign must be lower by the (weighted) same amount.

4. Quantitative analysis

Our analysis has established that a monetary union alters the impact of business cycle shocks at the household level. It does so by shifting the adjustment horizontally across countries within the brackets of the wealth distribution. We now perform a quantitative analysis in order to assess how strongly this effect plays out for different types of households. We perform the quantitative analysis in a version of the model that is extended along a number of dimensions and, importantly, it no longer restricts Home and Foreign to be symmetric at the household level. Specifically, we calibrate the model to two countries of the euro area that represent polar cases in terms of the wealth distribution: Germany and Italy. For this version of the model, we also show that the results established by Propositions 1 and 2, which rely on symmetry, are still approximately satisfied.

4.1. Medium-sized HANK^2 model

Since our question at hand is a quantitative one, we enrich our model laid out in Section 2 by features that are frequently used in medium-sized business cycle models. In particular, we use a two-country version of the model developed in Bayer et al. (forthcoming) which has been shown to be able to generate business-cycle dynamics that conform well with the data. We calibrate this medium-sized HANK^2 model to capture key aspects when it comes to asset holdings and wealth distributions in Germany and Italy. In what follows, we briefly sketch the main extensions of the model and delegate a full description of the extended model to Online Appendix A.3. As before, the structural features are the same in Home and Foreign. Yet by assigning different parameter values below we make sure that Home and Foreign differ—in accordance with the data.

Households. We modify the household side in three ways in order to be better able to match the wealth distribution in the data. First, we assume that a group of households is employed by firms while others are self-employed entrepreneurs. The former group receives only labor income while entrepreneurs earn firm profits that arise due to monopolistic competition in the goods market (see below). Yet, households may move from one group (or employment state) to the other according to some exogenous probability. Both labor income and profit income are subject to a proportional income tax. Second, we assume that households can hold two different types of assets, liquid government bonds, and illiquid capital. Capital holdings are illiquid because we assume that only a random share of households can trade capital in a given period. Third, we assume that in Foreign, which will be calibrated to Germany, households will receive a minimum income benefit which we model as a targeted transfer which those households receive whose income is below a certain threshold. As Bayer et al. (2022) show in detail, large differences in minimum income benefits across Germany and Italy can explain a large part of the differences in the wealth distribution,[4] and, as a result, requires large differences in government debt (high in Italy and low in Germany) in order to obtain the same real interest in both countries in steady state.

Firm sector. We also extend the firm sector by assuming that not only wages but also prices are adjusted infrequently. To this end, we assume a multi-layered production structure. Intermediate goods producers operate under perfect competition using both domestic capital and labor which we assume are immobile across countries. We also assume that production is subject to country-specific total factor productivity (TFP). Final good producers, in turn, differentiate domestic intermediate goods under monopolistic competition and are subject to Calvo (1983)-type price setting frictions in Home and Foreign. Domestically and imported goods are then bundled into consumer goods as in Section 2. Capital producers also use intermediate goods and face quadratic investment adjustment costs.

Fiscal policy. Lastly, we also consider a somewhat richer set of fiscal policies. First, the government in Foreign has to fund the minimum income benefits. Second, we now also consider government spending. This will allow us to analyze how a government spending shock plays out, both under flexible exchange rates and in the monetary union.

Shocks. In what follows we focus on TFP shocks and government spending shocks which may originate either in Home or Foreign. We assume each of these four shocks follows an exogenous AR(1)-process.

4.2. Symmetric calibration

Compared to the stylized model in Section 2, the medium-sized HANK^2 model outlined here features a richer structure in order to better capture key aspects of the data, both at the micro and the macro level. Before we calibrate the model to the EA, we therefore verify that the results established by Propositions 1–3 still hold exactly once we simulate a perfectly *symmetric* version of the model.

For this purpose, we pick parameter values for both Home and Foreign in line with the "Italy calibration" below. We provide detailed results in Online Appendix A.5. We study, in particular, the transmission of TFP shocks at the country level and at the union level. We find that monetary union alters the effects of a country-specific shock at the country level very much. Yet, while there is "too little" adjustment in one country and "too much" in the other these effects also offset each other completely in the larger model—in line with Proposition 1. The response of union-wide prices and quantities to country-specific shocks is independent of whether the countries operate a monetary union or not. Consequently, as argued above, the terms that enter the consumption function of union-wide twins are the same (Proposition 2), implying that monetary union does not shift the impact of shocks vertically across the wealth and income distribution. Instead, it shifts the impact of shocks horizontally across countries, as established in Proposition 3. To the extent that our results below differ from those stated in Propositions 1–3, this thus reflects the asymmetric calibration which in turn captures the differences in household-level heterogeneity in Germany and Italy.

[4] Pham-Dào (2016) shows that this is the case more generally across euro area countries. However, she uses a single-asset incomplete markets model.

Table 1
Calibrated model v Data.

			Model		Data	
			H	F	ITA	GER
Steady state	Assets	Debt (% of output)	132	71	132	71
		Capital–Output-Ratio	3.3	3.2	3.3	3.2
(targeted)	Distribution	Wealth gini	0.60	0.72	0.61	0.73
		Top-10% wealth share	0.43	0.55	0.44	0.52
		Bottom-50% wealth share	0.10	0.01	0.09	0.02
		Borrowers	0.08	0.18	0.08	0.18
Business Cycle	Volatility	Std(Y)*100 (targeted)	3.78	2.90	3.78	2.74
		Std(C)/Std(Y)	0.95	0.86	0.95	0.90
		Std(I)/Std(Y)	2.52	3.00	1.82	1.60
		Std(π)/Std(Y)	0.63	0.67	0.33	0.40
	Co-Movement	Corr(Y, Y^*) (targeted)	0.80		0.80	
		Corr(C, C^*)	0.95		0.79	
		Corr(I, I^*)	0.89		0.33	
		Corr(π, π^*)	0.97		0.77	

Notes: Model predictions based on baseline calibration, see Appendix A.4 for details. Micro data based on the 2017 wave of the Household Finance and Consumption survey of the ECB. Macro data from Eurostat and Worldbank (Inflation). Quantities are measured in real per capita terms, yoy changes; sample: 1999Q1–2022Q2.

4.3. Asymmetric calibration to the euro area

We outline how we calibrate the model to the EA and refer readers to Online Appendix A.4 for more details. Importantly, we now allow countries to differ not only in terms of shocks but also in terms of heterogeneity at the household level, in line with the data for Germany and Italy. For this purpose, we set parameters to target the wealth distributions and asset holdings in both countries.

For most parameters, we use standard values as listed in Online Appendix A.4. We specify the parameters that determine the income process at the household level to match micro-level estimates established for German and Italian data. In particular, we set the persistence of idiosyncratic income shocks to a standard value found for the euro area, see for example Pham-Dào (2016); and set the respective standard deviations as to match income inequality in Italy and in Germany. Moreover, we assume that in Foreign there are minimum income benefits which, overall, amount to 1% of GDP, in line with data for Germany. There are no minimum income benefits in Home which represents Italy.[5]

We then use six parameters to target key features of the wealth distributions and asset holdings in Germany and Italy. In particular, we use the discount factor, the portfolio adjustment probability, the probability which governs the transition of households to become entrepreneurs, and the borrowing penalty to match the level of government debt, the capital-to-output ratio, the wealth Gini, the top-10% wealth share, the bottom-50% wealth share, and the mass of borrowers.

Table 1 reports key moments as predicted by the model under the baseline calibration and contrasts them with their empirical counterparts. The top panels show values for the steady state where the empirical moments have been used as calibration targets. Note that the model is able to generate the observed large asymmetry between both countries: The wealth distribution is much more unequal in Germany compared to Italy, while government debt is considerably higher in Italy. As explained in detail in Bayer et al. (2022), from the perspective of the model these two aspects are interrelated and can be explained by the stronger need for self-insurance when minimum income benefits are not available. It should be noted that the degree of asymmetry between the two countries in these areas is one of the largest in the euro area (see also Pham-Dào, 2016; Kindermann and Kohls, 2017).

At the same time, the model is able to capture key features of the business cycle. For this purpose, we set parameters which capture macro frictions and policies in line with estimates of Bayer et al. (forthcoming). Note that for the baseline, we assume a common monetary policy (monetary union) and assume that the exchange rate is permanently fixed. Monetary policy is described by an interest rate feedback rule with interest rate smoothing. For the open economy parameters, we rely on standard parameter values in the literature, see again Online Appendix A.4. We follow Enders et al. (2013) in specifying country-specific government spending shocks and TFP shocks. We also add a common TFP shock since, otherwise, the model cannot account for the high degree of business-cycle comovement across both countries. We specify TFP shocks and the common component as we target output volatility in both countries and the co-movement of output across countries. The lower part of Table 1 shows that our model then does a fairly good job in matching other key statistics of the Euro area business cycle. In particular, the relative volatility of investment, consumption, and inflation is in the right ballpark, as is the co-movement across countries.

[5] The discretization of the income process effectively also establishes a non-zero lower income bound for the Home country, that is however much lower than the Foreign country's minimum income benefits.

4.4. Macroeconomic adjustment to country-specific shocks

We use the calibrated model to analyze the macroeconomic adjustments to country-specific shocks. In particular, we consider a TFP shock originating in Foreign and a government spending shock originating in Home. Even though the countries are no longer symmetric, a very similar pattern emerges for shocks that originate in the other country. Hence we do not report results for this case to economize on space.

Fig. 1 shows the responses of country aggregates in Home and Foreign as well as the union-wide aggregates to a contractionary TFP shock that originates in Foreign. Throughout, we contrast results for the monetary union (blue solid line) with those for independent monetary policies (red dashed line) in Home (left column), in Foreign (middle column), and in the entire union by displaying the aggregate responses (right column). Recall that in case of independent monetary policies, the interest rate feedback rule is the same as in the monetary-union case, except that monetary policy in each country responds to country-level rather than union-wide inflation rates. In each panel, the horizontal axis measures time in quarters, the vertical axis measures the percentage (or percentage point) deviation from steady state.

The top panels show the shock process which is independent of whether there is a monetary union or not: TFP in Foreign contracts; it is unchanged in Home. The second row shows the adjustment of output. Here monetary union makes a fundamental difference. Output in Foreign *increases* on impact with a monetary union in place, but *decreases* under independent policies. Likewise, output in Home also responds very differently across monetary regimes: it increases much more under independent monetary policies.

To rationalize these differences, it is instructive to study the adjustment of the policy rate, shown in the fourth row: in the monetary union it responds in the same way in both countries, while with independent monetary policies, we observe an increase in the short rate in Foreign and a decline in Home. This reflects, in turn, the differential impact of the shock on inflation in Foreign and Home which is shown in the third row: the contractionary TFP shock is strongly inflationary in Foreign, thus necessitating a monetary contraction. The policy response is much weaker in Foreign once it operates in a monetary union. For Home, it is the opposite: operating inside the monetary union implies a restrictive instead of an expansionary monetary policy. This is the one-size-doesn't-fit-all issue that is at the heart of the policy discussion in monetary unions. In our context, it even induces a change in the sign of the response of Foreign output in response to a TFP shock.

The right column of Fig. 1 shows the aggregate response of Home and Foreign under monetary union and with independent monetary policies. For all variables, the two impulse responses lie almost perfectly on top of each other. This shows that even though the two countries now differ substantially in terms of household-level heterogeneity, the result of Proposition 1 still holds approximately: monetary union does change the adjustments to country-specific shocks at the country level but it does so by shifting the adjustment between countries. The overall effect of a monetary union on the adjustment of union-wide aggregates turns out to be negligible.

Against this background, the bottom panels of the figure show the response of the consumption Gini. Consumption inequality increases after the shock in both countries and across both exchange rate regimes for reasons which become clear below. At this point it is important to point out that the response of aggregate, union-wide consumption inequality (shown in the right column) is basically independent of the exchange rate regime—consistent with Proposition 2. Yet consumption inequality increases less at Foreign and more at Home in case of monetary-union, compared to what we observe under independent monetary policies.

Turning to the effects of a government spending shock, shown in Online Appendix A.6, we observe that monetary union also alters the dynamics within countries profoundly. In particular, a monetary union amplifies the output response in Foreign, but dampens it in Home. This again reflects the common monetary stance in the union. However, the overall effect of monetary union on the adjustment of union-wide aggregates is close to zero—just like in the case of TFP shocks. Hence, the result of Proposition 1 also holds approximately for government spending shocks, even if we allow countries to have asymmetries.

4.5. Adjustment at the household level

We are finally in a position to address the main question of the paper: How does monetary union alter the impact of country-specific shocks at the household level? Proposition 3 above establishes for the symmetric case that monetary union indeed shifts the adjustment to country-specific shocks across borders. Specifically, monetary union shifts the adjustment at the household level horizontally across borders within the brackets of the wealth distribution. For the calibrated version of the model, we have shown that the results established in Propositions 1 and 2 approximately hold in the larger model even as it is calibrated to capture cross-country heterogeneity at the household level.

We now use this version of the model to quantify how monetary union alters the impact of shocks at the household level. Specifically, we compute the welfare impact of a shock for each household across the wealth distribution and contrast results for the monetary union with those for independent monetary policies. We measure the welfare impact using the consumption equivalent variation, which is the permanent consumption change that would make an individual household equally well off as the shock under consideration. We stress upfront that we take an ex-post perspective, evaluating welfare based on specific shocks (that is, one-sided welfare), rather than providing an ex-ante welfare analysis based on a second-order approximation of the utility function.

To synthesize results, we compute how the consumption equivalent variation due the shock under consideration changes for each decile of the wealth distribution as countries move from independent monetary policy to monetary union. Fig. 2 shows the results. The left (right) panels show the welfare differences in the adjustment due to monetary union for an adverse TFP shock that originates in Foreign (Home). The upper (lower) panel depicts the effect along the deciles of the wealth distribution in Home

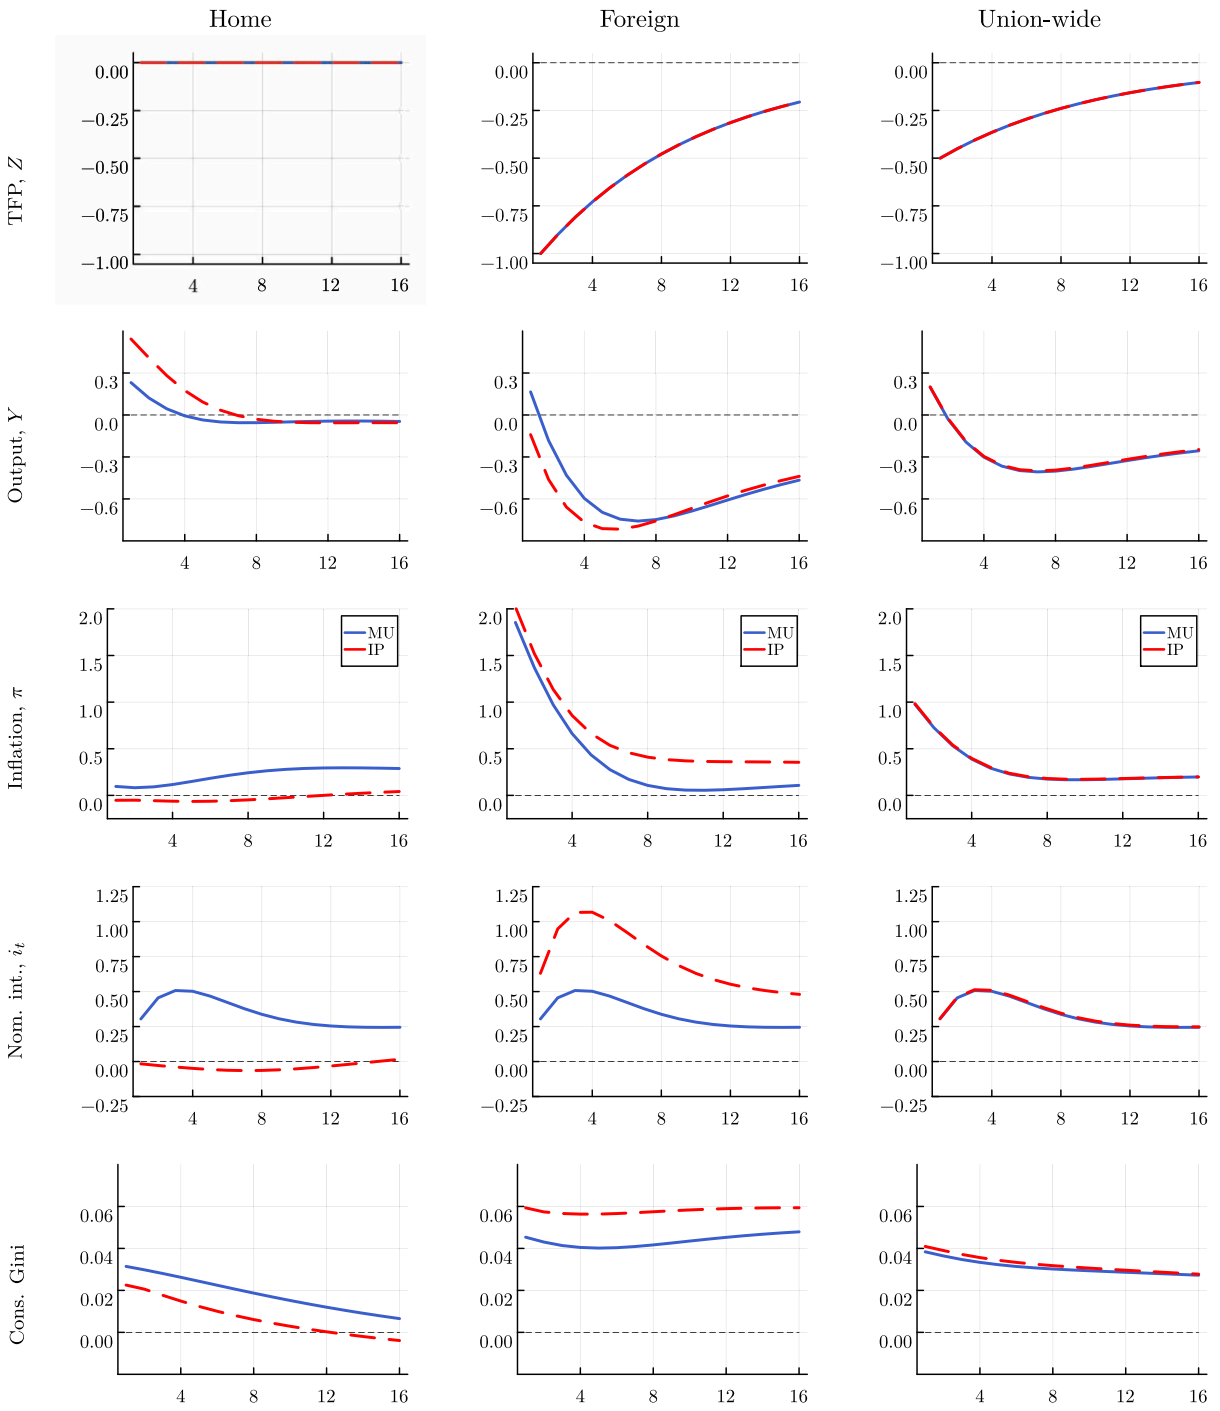

Fig. 1. Adjustment to adverse TFP shock originating in Foreign.
Notes: monetary union v independent monetary policies in Home (left), in Foreign (middle), and aggregate of Home and Foreign (right). Y-axis: Percentage deviation from steady state and percentage points in case of interest rates. X-axis: Quarters.

(Foreign). The emerging pattern is clear-cut and warrants three observations. First, the pattern is consistent with the result of Proposition 3 according to which monetary union shifts the impact of shocks across borders within the brackets of the wealth distribution. For the calibrated model, this does not hold exactly because we relax the assumption of country symmetry. And yet, we find that the adjustment in Home and Foreign is still fairly symmetric across the wealth distribution—the adjustment in Foreign

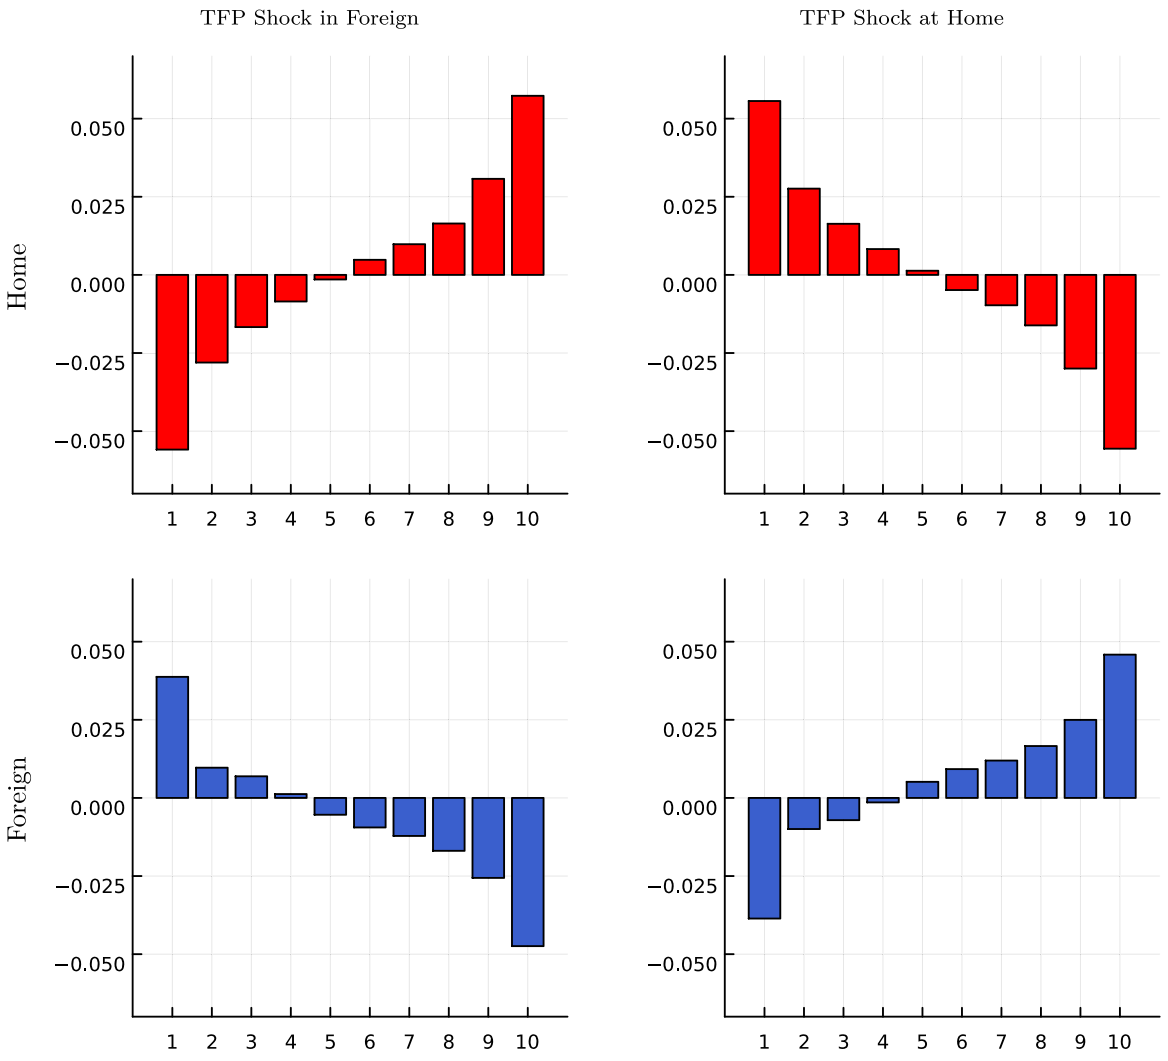

Fig. 2. How monetary union alters the welfare impact of shocks.
Notes: Difference of welfare impact of a Foreign adverse TFP shock (left) and a Home adverse TFP shock (right) between monetary union and independent monetary policies in Home (upper panel) and Foreign (lower panel). Y-axis: Difference in terms of consumption equivalent compensating variations. X-axis: Wealth deciles.

mirrors those in Home. Interestingly, we observe that there is now some vertical shift in the impact of the shock: the overall effect for each wealth bracket in Home and Foreign is not exactly zero.[6]

Second, we find that the change in the impact of shocks on welfare due to monetary union is much more concentrated in the tails of the wealth distribution. The middle class is much less affected. Put differently, whether there is a monetary union in place or not matters for the shock's welfare impact, but only for the poor and the rich. This is consistent with Proposition 3, and holds true for both Home and Foreign.

Third, these patterns do not depend on the origin of the shock. This becomes clear when comparing the left and the right column of the figure. If the shock originates in Home instead of Foreign, the exact same patterns emerge—just with a flipped sign. Likewise, the pattern emerges for positive and negative shocks equally; and it is not specific to TFP shocks. A similar pattern also emerges for government spending shocks (see Figure A.4 in the Online Appendix).

To see how monetary union alters the welfare impact of shocks along the wealth distribution, we decompose the overall effect for each decile into partial equilibrium effects. For this purpose, we exploit the fact that how the welfare of a given household is affected by a TFP shock depends on the arguments that enter its policy functions and thus the choice of variables which directly impact its welfare. By changing the adjustment of these at the country level, a monetary union changes the effect of a shock on households' welfare.

[6] Figure A.2 in the Online Appendix reports results for the symmetric calibration: in this case, the patterns in both countries are also perfectly symmetric.

To economize on space, we focus on an adverse TFP shock originating in Foreign and show results in Fig. 3. Based on our decomposition, each row shows the contribution of a specific variable to the overall effect: the real rate on liquid bonds, capital income, labor and profit income, and taxes. The blue bars represent the consumption equivalent variation of the shock under a monetary union, while the red bars represent the counterpart for the case of independent monetary policies. The left (right) panels report results for the deciles of the wealth distribution in Home (Foreign). We observe that monetary union has a strong bearing on the welfare impact of the shock by changing the way the real interest rate (top panel) and taxes (bottom panel) respond to the shock, and more so than for labor and capital income (middle panels). This is intuitive because–as discussed above–monetary union changes the interest rate response to a country-specific TFP shock. Changes in the interest rate then impact governments' budgets by altering the interest rate burden on the outstanding debt which, ultimately, results in an adjustment of the tax rate.

Importantly, the way in which the different adjustments of the real interest rate and the taxes affect households' welfare is highly heterogeneous along the wealth distribution. And we observe that the changes due to monetary union are largest at the tails of the wealth distribution—reflecting a different interest rate exposure. High-wealth households are directly exposed to interest rate changes through their assets, while low-wealth households are exposed through the tax response. As shown in Fig. 1 above, a negative TPF shock in Foreign raises interest rates and hence the return on the liquid asset in Home. At the same time, wages fall and labor income taxes rise. This benefits the asset-rich and harms the asset-poor. The monetary union changes the size of these price responses. As a result, the welfare impact of the shock on the poor and the rich depends on whether there is a monetary union in place or not.

Fig. 3 also illustrates why the middle class is largely unaffected by a monetary union. In this case monetary union also changes the welfare impact of different aspects of the shock: for example, the middle class in Home benefits from the higher real rates in response to an adverse TFP shock in Foreign and it suffers from the increased tax response. But these two effects roughly cancel each other out. This is intuitive as the middle class roughly holds an average amount of wealth and pays taxes on roughly an average income.

These results offer a new perspective on some of the policy debates surrounding the euro. In fact, they may explain why the European Monetary Union did not break up in the face of sizeable asymmetric shocks during its 20-year-plus history. In every country, those that benefit from the (union-wide) monetary response can always form a sufficiently large coalition with the middle class to support the union, as long as there is a small (and here non-modeled) cost of breaking up the union. Focusing on how monetary union alters the welfare impact of business cycle shocks thus offers new insights into the political economy of monetary unions, an issue that calls for further research.

5. Conclusion

Asymmetric shocks are a classic theme of OCA theory. They bring to the fore the one-size-doesn't-fit-all problem from which monetary unions are bound to suffer at times. We revisit the issue through the lens of a Heterogeneous Agent New Keynesian model with two countries: HANK^2. It belongs to a class of models that breaks with the representative agent paradigm and offers new perspectives. In particular, in contrast to earlier generations of OCA theory, we are no longer confined to analyzing what membership in a monetary union means for countries or regions as such.

Instead, we can investigate how monetary union alters the impact of shocks for individual households. In particular, the HANK^2 structure allows us to distinguish how monetary union alters the impact of shocks horizontally across borders within the brackets of the wealth distribution and vertically across the union-wide wealth brackets. A key result of our analysis is that a monetary union shifts the adjustment to shocks horizontally and not so much vertically. We show in closed form that a monetary union neither changes the union-wide dynamics after a country-specific shock nor its vertical impact across the brackets of the union-wide wealth distribution. Instead, it shifts the impact horizontally across borders within the brackets of the wealth distribution, from the poor in one country, for instance, to the poor in the other country.

Our quantitative analysis shows that this effect is particularly strong for the tails of the wealth distribution and weaker for the middle class. This brings to the fore questions about the political economy of currency unions which we take up in a companion paper (Bayer et al., 2022). Here we just note that our results may provide a rationale for why the EA did not break up during its 20-year-plus history despite several severe crises and calls for an exit of individual countries: In the face of a specific shock (or crisis), whether a country operates inside a monetary union or not does not matter so much for a large fraction of the population. That being said, we find that it can matter a lot for a smaller fraction of the population. This part of the population, however, lacks political majorities.

Transparency declaration

The publication of the supplement is sponsored by BANQUE DE FRANCE.

Data availability

Data will be made available on request.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jmoneco.2024.103579.

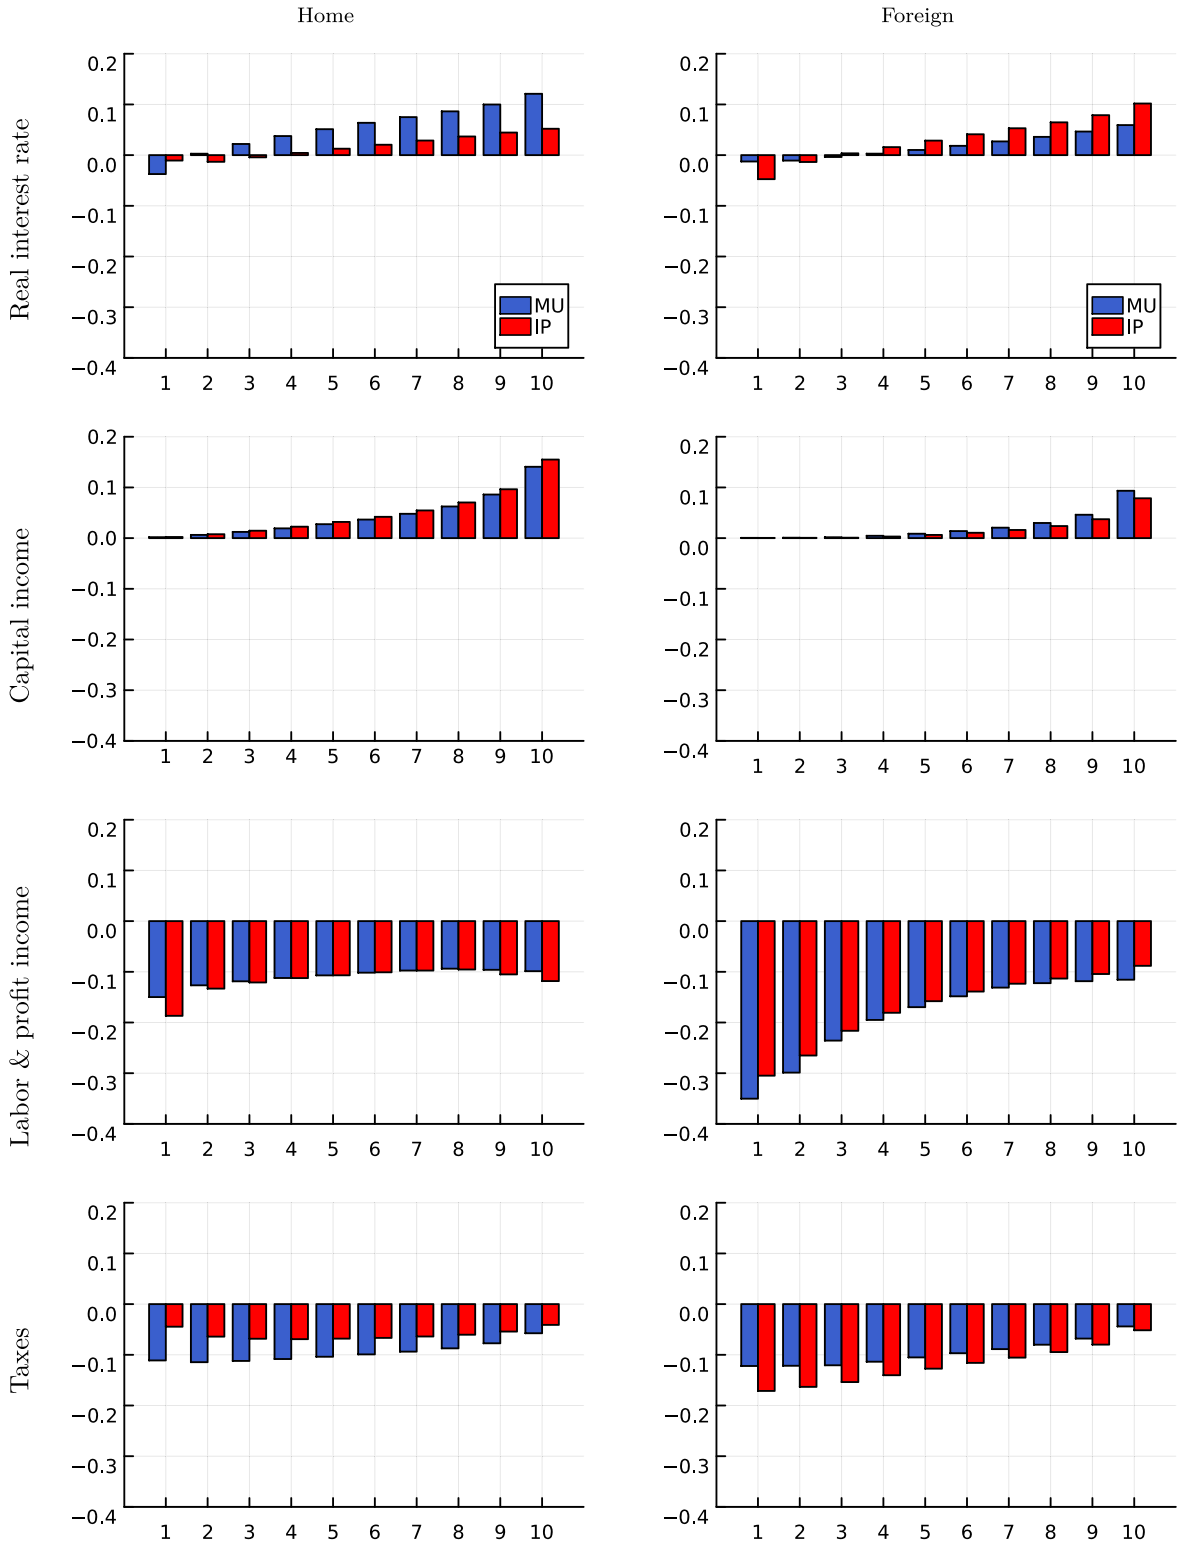

Fig. 3. Decomposition of welfare effect of TFP shock in Foreign.
Notes: Decomposition of welfare effects of a contractionary TFP shock in Foreign across the deciles of the wealth distribution. Monetary union (blue) vs Independent monetary policies (red). Y-axis: Difference in terms of consumption equivalent compensating variations. X-axis: Wealth deciles.

References

Aggarwal, R., Auclert, A., Rognlie, M., Straub, L., 2023. Excess savings and twin deficits: The transmission of fiscal stimulus in open economies. NBER Macroecon. Ann. 37 (1), 325–412.

Alesina, A., Barro, R.J., 2002. Currency unions. Q. J. Econ. 117 (2), 409–436.

Auclert, A., Bardóczy, B., Rognlie, M., Straub, L., 2021a. Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models. Econometrica 89 (5), 2375–2408.

Auclert, A., Rognlie, M., Souchier, M., Straub, L., 2021b. Exchange rates and monetary policy with heterogeneous agents: Sizing up the real income channel. mimeo.

Auclert, A., Rognlie, M., Straub, L., 2024. The intertemporal keynesian cross. J. Polit. Econ..

Bayer, C., Born, B., Luetticke, R., 2024. Shocks, frictions, and inequality in US business cycles.

Bayer, C., Kriwoluzky, A., Müller, G.J., Seyrich, F., 2022. Das Sein bestimmt das Bewußtsein: Attitudes towards fiscal policy according to HANK^2. mimeo.

Bayoumi, T., Eichengreen, B., 1992. Shocking Aspects of European Monetary Unification. NBER Working Papers No. 3949.

Beetsma, R., Jensen, H., 2005. Monetary and fiscal policy interactions in a micro-founded model of a monetary union. J. Int. Econ. 67 (2), 320–352.

Beetsma, R., Uhlig, H., 1999. An analysis of the stability and growth pact. Econom. J. 109 (458), 546–571.

Bellifemine, M., Couturier, A., Jamilov, R., 2023. The Regional Keynesian Cross. mimeo.

Benigno, P., 2004. Optimal monetary policy in a currency area. J. Int. Econ. 63 (2), 293–320.

Benigno, P., López-Salido, J.D., 2006. Inflation persistence and optimal monetary policy in the euro area. J. Money Credit Bank. 38 (3), 587–614.

Calvo, G.A., 1983. Staggered prices in a utility-maximizing framework. J. Monetary Econ. 12 (3), 383–398.

Chen, X., Lazarakis, S., Varthalitis, P., 2023. Debt targets and fiscal consolidation in a two-country HANK model for Eurozone. mimeo.

Corsetti, G., 2008. A Modern Reconsideration of the Theory of Optimal Currency Areas. European Economy- Economic Papers 2008 - 2015 No. 308.

Corsetti, G., Meier, A., Müller, G.J., 2012. Fiscal Stimulus with Spending Reversals. Rev. Econ. Stat. 94 (4), 878–895.

de Ferra, S., Mitman, K., Romei, F., 2021. Why does capital flow from equal to unequal countries? CEPR Discussion paper 15647.

Enders, Z., Jung, P., Müller, G.J., 2013. Has the euro changed the business cycle? Eur. Econ. Rev. 59, 189–211.

Erceg, C.J., Henderson, D.W., Levin, A.T., 2000. Optimal monetary policy with staggered wage and price contracts. J. Monetary Econ. 46 (2), 281–313.

Farhi, E., Werning, I., 2017. Fiscal unions. Amer. Econ. Rev. 107 (12), 3788–3834.

Frankel, J.A., Rose, A.K., 1998. The endogeneity of the Optimum Currency Area criteria. Econ. J. 108 (449), 1009–1025.

Galí, J., 2015. Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications. Princeton University Press.

Galí, J., Monacelli, T., 2005. Monetary policy and exchange rate volatility in a small open economy. Rev. Econom. Stud. 72, 707–734.

Gali, J., Monacelli, T., 2008. Optimal monetary and fiscal policy in a currency union. J. Int. Econ. 76 (1), 116–132.

Galí, J., Monacelli, T., 2016. Understanding the gains from wage flexibility: The exchange rate connection. Amer. Econ. Rev. 106 (12), 3829–3868.

Groll, D., Monacelli, T., 2020. The inherent benefit of monetary unions. J. Monetary Econ. 111 (C), 63–79.

Guo, X., Ottonello, P., Perez, D.J., 2020. Monetary policy and redistribution in open economies. mimeo, National Bureau of Economic Research.

Hettig, T., Müller, G.J., 2018. Fiscal policy coordination in currency unions at the effective lower bound. J. Int. Econ. 115 (C), 80–98.

Kenen, P.B., 1981. The theory of Optimum Currency Areas: An eclectic view. In: Essays in International Economics. Princeton University Press, Princeton, pp. 163–182.

Kindermann, F., Kohls, S., 2017. Rental Markets and wealth inequality across Europe. mimeo.

Kollmann, R., 2004. Welfare effects of a monetary union: The role of trade openness. J. Eur. Econom. Assoc. 2 (2/3), 289–301.

Krugman, P., 1993. Lessons of massachusetts for EMU. In: Giavazzi, F., Torres, F. (Eds.), Adjustment and Growth in the European Monetary Union. Cambridge University Press, pp. 241–261.

McKay, A., Wolf, C., 2022. Optimal policy rules in HANK. mimeo.

McKinnon, R.I., 1963. Optimum Currency Areas. Am. Econ. Rev. 53 (4), 717–725.

Mundell, R.A., 1961. A theory of Optimum Currency Areas. Am. Econ. Rev. 51 (4), 657–665.

Oskolkov, A., 2023. Exchange rate policy and heterogeneity in small open economies. J. Int. Econ. 142, 103750.

Pham-Dào, L.A., 2016. Income Risk, Self-Insurance and Public Policies: Consequences for Aggregate Fluctuations and Wealth Inequality (Ph.D. thesis). Universitäts-und Landesbibliothek Bonn.

Rose, A.K., 2000. One money, one market: the effect of common currencies on trade. Econ. Policy 15 (30), 08–45.

Zhou, H., 2021. Open economy, redistribution, and the aggregate impact of external shocks. mimeo.