International Risk Sharing in the European Monetary Union

Alessandro Ferrari†

Anna Rogantini Picco†

Abstract

This paper aims at empirically assessing the effect of the adoption of the euro on the ability of euro area member states to smooth consumption and share risk. With the objective of evaluating the economic performance of euro area countries in the scenario where the euro had not been adopted, we construct a counterfactual dataset of macroeconomic variables via the Synthetic Control Method. In order to get some preliminary measures of risk sharing, we compute correlations between consumption and GDP within a country, bilateral consumption correlations, and Brandt-Cochrane-Santa Clara indices across euro area member states. We then decompose risk sharing in different channels by means of the Asdrubali, Sorensen and Yosha (1996) output variance decomposition. Our difference in difference estimates show that the euro has not affected the level of international risk sharing across euro area countries, but has partially reduced the ability of member states to smooth consumption. We attribute this change to the higher GDP growth generated by the adoption of the euro, which has been accompanied by a greater output volatility. We also report differential effects for euro area core and periphery countries, showing that the former have not suffered any negative effect from the adoption of the euro in terms of risk sharing whereas the latter are now less able to smooth consumption.

Keywords: Consumption Smoothing, Risk Sharing Channels, European Monetary Union, Synthetic Control Method

JEL codes: F32, F36, F41

†European University Institute, Villa La Fonte, Via delle Fontanelle, 18, 50014 San Domenico di Fiesole, Italy; Email: alessandro.ferrari@eui.eu and anna.rogantini@eui.eu

Acknowledgments

_____

1   Introduction

The adoption of a common currency by euro area member states has guaranteed a high degree of price stability. At the same time, however, member states have lost the possibility to use monetary policy as a leeway to respond to idiosyncratic shocks. While this drawback is intrinsic to being part of a currency union, what makes the euro area different from other existing monetary unions is that fiscal policy is still conducted at national level. The coexistence of a common monetary policy and a decentralised fiscal policy generates remarkable tensions in the euro area when it comes to absorbing national idiosyncratic shocks. On the one hand, being part of a currency union, member states cannot absorb idiosyncratic shocks through monetary policy as opposed to countries which still keep monetary policy at national level; on the other hand, not being part of a fiscal federation, they do not receive fiscal transfers from a central budget as happens, for example, to US states for state-level shocks or to German federal regions for regional shocks. In the absence of some common fiscal capacity, shocks are mainly absorbed through the issuance of national non-contingent debt. In spite of being a powerful instrument to smooth consumption, debt is not intrinsically an international risk sharing mechanism, unless it is traded across the borders. That is why the recent financial and economic crises have endeavoured, on the one side, the creation of new mechanisms to generate public international risk sharing[1], and, on the other side, the consolidation of newly founded institutions such as the Banking Union to boost private international risk sharing. The lack of public risk sharing mechanisms is a major rationale behind the proposal included in the Five Presidents' Report (2015) regarding the creation of a euro area fiscal capacity able to absorb asymmetric shocks.

Given this ambitious project at euro area level, it is crucial to quantify the level of risk sharing across euro area members and to assess whether the adoption of the euro has brought any change in the ability to share risk. Estimates of risk sharing at euro area level already exist in the literature – see van Beers, Bijlsma, and Zwart (2014), Kalemli-Ozcan, Luttini, and Sørensen (2014), and Furceri and Zdzienicka (2015). However, to the best of our knowledge, no one has yet tried to evaluate whether the adoption of the euro has had any impact on the level of risk sharing across euro area member states and this is where the contribution of our paper lies. In order to evaluate any possible effect of the adoption of the euro on the level of risk sharing across euro area member states we proceed in three steps. First, we generate a counterfactual dataset for the scenario in which euro area countries had not adopted the common currency. Then, we compute some preliminary measures of risk sharing both with the actual and the counterfactual data in order to assess whether there has been a change in risk sharing due to the adoption of the euro. Thirdly, we attempt to decompose risk sharing into several different channels to evaluate how risk sharing has changed with the euro. We now proceed to explain more

---

[1]For specific proposals on public international risk sharing institutions, see Abraham, Carceles-Poveda, Liu, and Marimon (2016), Poghosyan, Senhadji, and Cottarelli (2016), Carnot, Evans, Fatica, and Mourre (2015).

in detail these three steps.

To begin with, there is a major obstacle to evaluating the effect of the adoption of a common currency by euro area member states: the lack of an appropriate set of countries to be used as a counterfactual pool for the scenario in which the member states had not adopted the euro. We tackle this problem by using the so called Synthetic Control Method (SCM). The main benefit of this method is that it allows to build synthetic time series that can be used in the absence of a natural counterfactual. Precisely for this reason, SCM is exploited by the literature to estimate the effect of policy interventions when it is not possible to have a real counterfactual.

After generating a synthetic dataset of time series as a counterfactual in the scenario of no adoption of the common currency, our second step is to compute some preliminary measures of risk sharing across euro area countries both with the actual dataset of euro area countries' variables and with their synthetic counterparts. Our aim is to evaluate whether the adoption of the euro has had any impact on the level of risk sharing across euro area countries. We do so by computing difference in difference estimates of our measures of risk sharing calculated using the actual and the synthetic series pre and post adoption of the euro. Our measures of risk sharing are based on the following theoretical prescription: the more risk is shared across countries, the less each country's consumption and GDP are correlated and the more the consumption of two countries sharing risk are correlated. Consequently, we estimate correlation between consumption and GDP of each of our countries, and the correlation of consumption of all pairs of countries in our sample. The third risk sharing measure that we adopt is the so called Brandt-Cochrane-Santa Clara (BCS) Index, an indicator of bilateral risk sharing, which relies on the similarity of stochastic discount factors across countries. If the adoption of the euro has had any impact on risk sharing we should find that the difference in difference estimates of our risk sharing measures are significant.

Once computed difference in difference estimates for correlations and BCS indices and assessed the effect of the adoption of the common currency on risk sharing across euro area countries, our third step is to decompose risk sharing into different channels. In particular, our final goal is to assess which of these channels have changed after the adoption of the euro as opposed to the counterfactual scenario in which euro area countries had not adopted the common currency. In order to better understand the channels through which risk sharing is accomplished, we adopt the GDP variance decomposition introduced by Asdrubali, Sorensen, and Yosha (1996). This method allows us to distinghuish four possible channels of risk sharing across countries: private risk sharing through private cross-border investments, public risk sharing through government taxes and transfers, private savings, and public savings. It is worth noticing that *international* risk sharing happens only through the fist two channels, i.e. cross-border investments and government transfers.

Our main result, which is robust to our different risk sharing measures and

specifications, is that *international* risk sharing across euro area member states through both private and public channels has not changed due to the adoption of the common currency. Nonetheless, we find evidence of a decrease in consumption smoothing across euro area countries in the period after the introduction of the euro. Difference in difference estimates of the correlation between the GDP and consumption of each country are positive, indicating a reduction in consumption smoothing. Similarly, difference in difference estimates of bilateral consumption correlations and BCS indices are negative, again pointing at a reduction in consumption smoothing. At the same time, we report that the adoption of the common currency has had a positive effect on GDP growth, which has been accompanied by an increase in output volatility. We interpret our result on the lower level of consumption smoothing after the adoption of the common currency as follows: we attribute the lower consumption smoothing to the inability of agents to insure against larger shocks to GDP compared to the pre euro period. Furthermore, we provide evidence of heterogeneous effects by splitting the sample of euro area member states into core and periphery countries. We show that the euro has not affected significantly consumption smoothing of core countries, whereas the aggregate negative effect that we find for the whole sample of countries is due to a reduction in consumption smoothing for periphery member states. Finally, we show that our results are robust to changes in the matching strategy, exclusion of potentially affected units from the group of non euro area countries, and changes of the year of euro adoption.

Related literature The SCM is first introduced by Abadie and Gardeazabal (2003) to test for the impact of the outbreak of terrorism in the Basque Country in the late 60s. Building on that seminal paper, it is further employed by Abadie, Diamond, and Hainmueller (2010) to estimate the effect of a large-scale tobacco control programme that California implemented in 1988. In addition, Billmeier and Nannicini (2012) use SCM to investigate the impact of economic liberalisation on real GDP per capita in a worldwide sample of countries. Closer to our focus, Campos, Coricelli, and Moretti (2014) make use of SCM to evaluate the benefits from being part of the European Union, while Saia (2016) employs SCM to estimate counterfactual trade flows between the UK and Europe if the UK had joined the euro.

Among the risk sharing measures that we use, there is the BCS index, which is proposed by Brandt, Cochrane, and Santa-Clara (2006). This is an indicator of bilateral risk sharing which relies on the similarity of pricing kernels. The BCS index is computed also by Rungcharoenkitkul (2011) to assess risk sharing among some Asian countries in the first decade of the 2000s. In order to understand the channels through which risk sharing is accomplished, we use the GDP variance decomposition presented by Asdrubali et al. (1996). The latter paper introduces this methodology to identify risk sharing channels in the US over the period 1963-1990. Later on, the GDP variance decomposition is used by Furceri and Zdzienicka (2015) to analyze and compare risk sharing among euro area countries with risk sharing across the US states. The analysis of Furceri and Zdzienicka (2015) is updated with more recent data by van Beers et al. (2014), who try to assess the functioning

of insurance mechanisms in the euro area, and by Kalemli-Ozcan et al. (2014) who separately consider countries hit by the sovereign debt crisis in 2010.

2   Methodology

2.1   The Synthetic Control Method

The main purpose of this paper is to assess whether the introduction of the common currency has had any effects on the level of risk sharing between member states. In order to meaningfully address this question, one would need to estimate risk sharing between euro area member states under the alternative scenario in which the currency area had not been established. Since it is not possible to have a real counterfactual for this situation, we use the SCM by Abadie and Gardeazabal (2003) to generate a synthetic counterpart. This method is a data driven procedure that has been used to estimate the effect of policy interventions in the absence of a natural counterfactual.

Our first step is to generate the synthetic counterpart of the following macroeconomic variables in per capita term: gross domestic product (GDP), household final consumption (C), government expenditure (G), national income (NI), and disposable national income (DNI). We are going to use these variables to compute the measures of risk sharing across countries discussed in sections 2.2 and 2.3. In order to compute the synthetic counterpart of our macroeconomic variables of interest, we proceed as in Abadie and Gardeazabal (2003). Let N be the number of countries in the potential counterfactual pool, and let $W = (w_i)_{i=1}^N$ an $N \times 1$ vector of country weights such that $w_i \geq 0$ and $\sum_i w_i = 1$ for $i = 1, ..., N$. Moreover, let X_1 be the $K \times 1$ vector of our variables of interest for *euro area* member states before the introduction of the euro. Similarly, let X_0 be the $K \times N$ matrix values of the same K variables of interest for all N *non euro area* countries in our counterfactual pool before the introduction of the euro. In addition, let V be a $K \times K$ diagonal matrix with non negative components representing the relevance of our variables of interest in determining the macroeconomic outcome variables. As discussed in Abadie and Gardeazabal (2003), while the choice of the matrix V could be arbitrarily based on economic considerations, here it is computed through a factor model. Then, the algorithm of Abadie and Gardeazabal (2003) looks for the vector W^* of weights that minimises $(X_1 - X_0 W)' V (X_1 - X_0 W)$, subject to $w_i \geq 0$ and $\sum_i w_i = 1$ for $i = 1, ..., N$. The vector W^* determines the linear combination of macroeconomic variables for non euro area countries, which best reproduces each variable of interest for the euro area countries in the period before the introduction of the euro. Therefore, let Y_1 and Y_0 be the outcome variables for respectively the euro area and the non euro area countries in our group of non euro area countries. Then, the method uses $Y_1^* = Y_0 W^*$ as counterfactual for the outcome variables of euro area countries after the introduction of the euro. The choice of the covariates in matrix X_0 is such that it maximises the ability of the synthetic series to reproduce the behaviour of the series of the euro area countries in the period before the introduction of the euro. The baseline matching function always takes the past

value of the variable we investigate. This means that if we are evaluating what Portuguese consumption C would have been without Portugal being part of the euro, we always start by matching on the consumption of Portugal in every year before the introduction of the euro. To continue on the example, in order to generate the counterfactual series of Portuguese C for the scenario in which Portugal had not adopted the euro, the method uses the variables GDP, C, G, NI, DNI of the non euro area countries in our sample and it chooses the vector of weights W so as to minimise the distance between Portuguese C and the combination of the macroeconomic variables we have at our disposal, in the subsample before the introduction of the euro. Once we have a synthetic series of Portuguese C which mimics the actual series in the matching period before the euro, we can use that series as a counterfactual for Portuguese C in the scenario where Portugal had not joined the euro in the period after the introduction of the euro.

This method relies on two identification assumptions: 1) the choice of the covariates on which the matching is carried out before the introduction of the euro should be such that the variables that are able to mimic the pre euro path are included, but do not rely on observables that anticipate the effect of the introduction of the euro itself; 2) the variables concerning the group of non euro area countries in our counterfactual pool should not be affected by the adoption of the euro. For the latter reason, the matching is carried out for one euro area country at the time, meaning that we iteratively drop all but one euro area member state, so that the procedure always involves one euro area country and N non-euro area countries.

A relevant assumption for the correct use of the SCM is that the non euro area group is unaffected by the adoption of the euro. This assumption can be troublesome since, given the magnitude of the potential effect of the euro, one might indeed think that the introduction of the common currency indirectly affected all countries in the world. This could be particularly true for the countries in our non euro area group, which is made of OECD countries with strong trade and financial linkages with our euro area sample. This concern is legitimate if we look at the effect of the introduction of the euro itself. However, one can think of the total effect of the euro for member states as being made of two components: i) the effect of the mere existence of the euro; ii) the effect of having adopted the euro and being a member of the currency union. Under this decomposition, even though all countries in the world may be subject to the first effect, only euro area member states are subject to the second one. Hence, the effect that we want to analyse should be interpreted as being the membership of the euro, conditional on the existence of the euro.

The intuition behind this method is that one can use the best linear combination of synthetic series in terms of matching the behaviour of actual series as a counterfactual for the national account aggregates of the euro area countries after the adoption of the euro. It is worth mentioning that the evaluation of the robustness of these estimates has been discussed in the literature, but no analytical result is available to compute the standard deviation of these

estimates, namely because the estimated component is the weighing vector. Robustness checks can then be carried out in three possible ways: i) performing bootstrap, by randomly resampling the donor pool of non euro area countries (see Saia (2016)); ii) estimating a difference in difference regression and testing whether the outcome is significantly different from zero (see Campos et al. (2014)); iii) running placebo studies on units in the donor pool in order to assess whether the method delivers spurious effect of the adoption of the euro. We will use all these techniques in order to check the robustness of our results.

2.2 Bilateral Correlations and the BCS Index

Economic theory predicts that, under the assumption of no arbitrage and complete markets, countries fully share risk. This has two major implications. Firstly, idiosyncratic shocks to income in a country are not transmitted to the consumption in the same country as they are fully insulated. In other words, even if a country's income is volatile, its consumption level should remain unchanged. Secondly, the stochastic discount factors (henceforth SDFs) of two countries which fully share risk are equalised – see for example Cochrane (2001). Put it differently, let $M_{i,t} = \beta u'(c_{i,t+1})/u'(c_{i,t})$ and $M_{j,t} = \beta u'(c_{j,t+1})/u'(c_{j,t})$ be the SDF of country i and country j. Under complete markets, it has to hold that $M_{i,t} = M_{j,t}$. In this situation, the growth of marginal utility is perfectly correlated across individuals. More specifically, if preferences u and discount factors β are assumed to be the same across countries, the growth rate of consumption is identical. Whenever the assumption of complete markets is violated, there are two consequences. On the one hand, part of the income volatility of a country is transmitted to consumption, so full insulation of consumption from variations in income no longer applies. On the other hand, SDFs between countries are not equalised and part of the risk remains untraded. What individuals seek to do in this case is to use all available assets to share the biggest possible portion of risk. Put differently, by trading the available assets they seek to get the SDF of the two countries as close as possible – see Svensson (1988). With this theoretical background in mind, we start our analysis by computing three potential measures of risk sharing across euro area countries. We are going to express the computed risk sharing measures in a difference in difference framework. To be more specific, we are going to compute the three risk sharing measures both with actual data and with synthetic data, across two sample periods: 1990-1998, and 1999-2011. We then take the difference of the risk sharing measures obtained with the actual and the synthetic data both in the first and the second subsample. We finally compute the difference of the two subsample differences in order to get the difference in difference estimate. The obtained estimates should tell us whether the adoption of the euro has had any impact on the risk sharing measures that we are computing. In order to test the significance of the estimate that we get, we compute t-statistics using bootstrap techniques. In other words, we randomly drop one country from the donor pool needed to generate the synthetic series and we compute the synthetic series. We repeat the exercise for 100 iterations. We can then compute the sample distribution of the synthetic estimate and its t-statistics.

We now turn to describe which are the three risk sharing measures that we are computing.

First of all, we compute correlations between consumption and GDP of each country. If risk sharing across euro area countries has increased, we expect to find a decrease in the correlation between each country's consumption and GDP. Viceversa, a decrease in risk sharing should result into an increase in the same correlation as consumption would be less insulated from variations in GDP. Secondly, we calculate bilateral consumption correlations across euro area members. As noted before, economic theory suggests that a higher level of risk sharing should increase consumption correlation between countries even when their GDP correlation is low.[2] If the adoption of the euro has had an impact (either positive or negative) on risk sharing, we should find that our difference in difference estimate of pairwise consumption correlations is significant. The third measure that we compute is the bilateral risk sharing indicator proposed by Brandt, Cochrane, and Santa-Clara (2006). This indicator, referred to as BCS index, captures the level of risk sharing between country i and country j and takes the following form:

$$\text{BCS}_{i,j} = 1 - \frac{\text{var}(\log M_{i,t+1} - \log M_{j,t+1})}{\text{var}(\log M_{i,t+1}) + \text{var}(\log M_{j,t+1})} \tag{1}$$

The numerator measures how far apart the SDFs of the two countries are from one another, i.e. what portion of risk is not shared. The denominator quantifies the volatility of SDF in the two countries, i.e. what is the total portion of risk to be shared. This metric ranges between minus one and one with a higher number meaning a higher degree of risk sharing. As noted in Brandt et al. (2006) this index differs from correlation. Indeed, like a correlation, it is equal to one when the two SDFs are the same, it is zero when they are uncorrelated, and it is minus one if $M_{i,t+1} = -M_{j,t+1}$. However, differently from a correlation, it detects violations of scale in the growth rate of marginal utilities. In fact, risk sharing requires the two countries' SDFs to be equal, not just perfectly correlated. Nevertheless, both the BCS index and the correlation of SDFs are statistical descriptions of how far we are from perfect risk sharing. In terms of computation, we assume that households in the two countries have the same preferences and, in particular, CRRA utilities with risk aversion $\sigma = 2$ and discount factor $\beta = 0.95$. Given this, their SDFs

---

[2]We are aware that there is contrasting evidence of this theory in the data. For example, Baxter and Crucini (1995) find that GDP correlation is much higher than consumption correlation even across countries which are known to share risk. More recently, instead, Krueger and Perri (2006) show that income volatility in the US over the period 1972-1998 was not accompanied by a corresponding rise in consumption volatility and they attribute this to the development of credit markets, which played a crucial role in isolating consumption against higher income risk. In any case, what we are mainly interested in is not much the level of correlation itself, but the difference in the correlation obtained from the actual data and the one obtained from the synthetic data.

look as follows:

$$M_{i,t} = \beta \left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{-\sigma}$$

$$M_{j,t} = \beta \left(\frac{C_{j,t+1}}{C_{j,t}} \right)^{-\sigma}.$$

2.3   GDP Decomposition

Following our preliminary investigation on consumption correlations and BCS indices, our next step is to identify the different channels through which risk sharing has changed. We carry out this analysis through a methodology proposed by Asdrubali et al. (1996). The idea of this analysis is to check which of the potential risk sharing channels absorb output shocks. In particular, this is implemented by decomposing GDP into the following national account aggregates: Gross Domestic Product (GDP), Net National Income (NI), Disposable National Income (DNI), and Private and Government Consumption (C+G). According to this decomposition, each country's GDP can be disaggregated as this accounting identity:

$$\text{GDP} \equiv \frac{\text{GDP}}{\text{NI}} \frac{\text{NI}}{\text{DNI}} \frac{\text{DNI}}{\text{DNI+G}} \frac{\text{DNI+G}}{\text{C+G}} (\text{C+G}) \tag{2}$$

Because of the differences in the national account aggregates, the ratios on the right-hand side can be interpreted as specific channels through which risk is absorbed. The first ratio, $\frac{\text{GDP}}{\text{NI}}$, accounts for income insurance stemming from internationally diversified investment portfolios. This is because NI measures the income (net of depreciation) earned by residents of a country, whether generated on the domestic territory or abroad, while GDP refers to the income generated by production activities on the economic territory of the country. Therefore, the ratio $\frac{\text{GDP}}{\text{NI}}$ captures the *private international* insurance channel due to private cross-border investments or, as Kalemli-Ozcan et al. (2014) refer to, holding of claims against the output of other regions. The ratio $\frac{\text{NI}}{\text{DNI}}$, instead, can be interpreted as the *public international* insurance channel due to government taxes and transfers as DNI is the income that households are left with after subtracting taxes and adding transfers. Finally, the ratios $\frac{\text{DNI}}{\text{DNI+G}}$ and $\frac{\text{DNI+G}}{\text{C+G}}$ account for smoothing through respectively *public* and *private* saving channels.

In order to measure how much of the variations in output is absorbed by each channel, we proceed as in Asdrubali et al. (1996). We first take logs of equation 2, we difference the series, we multiply by the change of log GDP, and we take expectations to get:

$$
\begin{aligned}
\text{Var}(\Delta \log \text{GDP}_{i,t}) \equiv\ & \text{Cov}(\Delta \log \text{GDP}_{i,t}, \Delta \log \text{GDP}_{i,t} - \Delta \log \text{NI}_{i,t}) \\
& + \text{Cov}(\Delta \log \text{GDP}_{i,t}, \Delta \log \text{NI}_{i,t} - \Delta \log \text{DNI}_{i,t}) \\
& + \text{Cov}(\Delta \log \text{GDP}_{i,t}, \Delta \log \text{DNI}_{i,t} - \Delta \log(\text{DNI}_{i,t} + \text{G}_{i,t})) \\
& + \text{Cov}(\Delta \log \text{GDP}_{i,t}, \Delta \log(\text{DNI}_{i,t} + \text{G}_{i,t}) - \Delta \log(\text{C}_{i,t} + \text{G}_{i,t})) \\
& + \text{Cov}(\Delta \log \text{GDP}_{i,t}, \Delta \log(\text{C}_{i,t} + \text{G}_{i,t}))
\end{aligned}
$$

9

Dividing both sides by $\text{Var}(\Delta \log \text{GDP}_{i,t})$ we get the following identity:

$$1 \equiv \beta^m + \beta^g + \beta^p + \beta^s + \beta^u$$

where we have defined

$$\beta^m \equiv \frac{\text{Cov}(\Delta \log \text{GDP}_{i,t}, \Delta \log \text{GDP}_{i,t} - \Delta \log \text{NI}_{i,t})}{\text{Var}(\Delta \log \text{GDP}_{i,t})}$$

$$\beta^g \equiv \frac{\text{Cov}(\Delta \log \text{GDP}_{i,t}, \Delta \log \text{NI}_{i,t} - \Delta \log \text{DNI}_{i,t})}{\text{Var}(\Delta \log \text{GDP}_{i,t})}$$

$$\beta^p \equiv \frac{\text{Cov}(\Delta \log \text{GDP}_{i,t}, \Delta \log \text{DNI}_{i,t} - \Delta \log(\text{DNI}_{i,t} + \text{G}_{i,t}))}{\text{Var}(\Delta \log \text{GDP}_{i,t})}$$

$$\beta^s \equiv \frac{\text{Cov}(\Delta \log \text{GDP}_{i,t}, \Delta \log(\text{DNI}_{i,t} + \text{G}_{i,t}) - \Delta \log(\text{C}_{i,t} + \text{G}_{i,t}))}{\text{Var}(\Delta \log \text{GDP}_{i,t})}$$

$$\beta^u \equiv \frac{\text{Cov}(\Delta \log \text{GDP}_{i,t}, \Delta \log(\text{C}_{i,t} + \text{G}_{i,t}))}{\text{Var}(\Delta \log \text{GDP}_{i,t})}$$

All β coefficients can be estimated thorugh the system of equations proposed by Asdrubali et al. (1996):

$$\Delta \log \text{GDP}_{i,t} - \Delta \log \text{NI}_{i,t} = \beta^m \Delta \log \text{GDP}_{i,t} + \epsilon_{i,t}^m \tag{3}$$

$$\Delta \log \text{NI}_{i,t} - \Delta \log \text{DNI}_{i,t} = \beta^g \Delta \log \text{GDP}_{i,t} + \epsilon_{i,t}^g \tag{4}$$

$$\Delta \log \text{DNI}_{i,t} - \Delta \log(\text{DNI}_{i,t} + \text{G}_{i,t}) = \beta^p \Delta \log \text{GDP}_{i,t} + \epsilon_{i,t}^p \tag{5}$$

$$\Delta \log(\text{DNI}_{i,t} + \text{G}_{i,t}) - \Delta \log(\text{C}_{i,t} + \text{G}_{i,t}) = \beta^s \Delta \log \text{GDP}_{i,t} + \epsilon_{i,t}^s \tag{6}$$

$$\Delta \log(\text{C}_{i,t} + \text{G}_{i,t}) = \beta^u \Delta \log \text{GDP}_{i,t} + \epsilon_{i,t}^u \tag{7}$$

where each β coefficient represents the share of income shocks smoothed by a given channel. In particular, β^m accounts for the share of GDP shocks smoothed by capital markets, β^g by fiscal transfers, β^p by public savings, β^s by private savings. What is left, β^u, is the unsmoothed part of the GDP shock. A zero β^u coefficient in the regression of total (private and public) consumption on GDP, i.e. equation 7, means that a shock to GDP is fully absorbed through capital markets, fiscal transfers, public and private savings, thus leaving consumption unchanged. Instead, a high β^u coefficient in the same regression means that only a minor part of the shock is absorbed through risk sharing, while a significant part stays unsmoothed.

In estimating the coefficients in the above system we do not control for country size. In other words, the national account series of each country are not adjusted by population or size of the economy. The estimation of coefficients is carried out using the following methods: OLS with time fixed effects and clustered standard errors, OLS with time fixed effects and panel correlated standard errors, generalized method of moments, and seemingly unrelated regressions. The inclusion of time fixed effects is important as it allows us to take out euro area business cycle fluctuations. In this way, we make sure that the effects that we find are deviations from the euro area business cycle and

not, rather, fluctuations of the euro area business cycle itself. In the rest of the paper our baseline estimation for the analysis of risk sharing channels will be an OLS estimation with time fixed effects. We also perform OLS with panel correlated standard errors, seemingly unrelated regression, and GMM. In particular, in GMM we separately estimate the above described relations using up to three lags of GDP growth as an instrument. The estimation procedure, which follows Arellano and Bond – see Roodman (2009) – automatically includes past values of the dependent variable as instruments.[3]

We show the results of these estimation strategies as computed in a difference in difference model, which is equivalent to separate estimations. Namely, we stack together our actual and synthetic samples and include the independent variable interacted with the four possible combinations of actual/synthetic and euro/no euro. In particular, the regressions that we are going to estimate are:

$$y_{i,t} = \beta_0 + \beta_1 x_{i,t} + \beta_2 \text{Tr}_i x_{i,t} + \beta_3 \text{Eur}_t x_{i,t} + \beta_4 (\text{Tr}_i * \text{Eur}_t) x_{i,t} + \nu_t + \epsilon_{i,t}$$

where x_i is $\Delta \log \text{GDP}_{i,t}$, y_i are the dependant variables in Equations 3 - 7, Tr_i is a dummy variable taking the value of 1 if the series comes from the actual dataset and 0 if it comes from the synthetic dataset, Eur_t is a dummy taking the value of 1 after the adoption of the euro in 1999 and 0 otherwise, and ν_t are time fixed effects. β_2 represents the share of GDP variation smoothed by a given channel for our actual data before the introduction of the euro in deviations from its synthetic counterpart β_1. If our matching is successful, we should find that β_2 is not significantly different from zero. For our euro period, i.e. for the period in which $\text{Eur}_t = 1$, we should compare β_4 with β_3. If the euro has had an effect on the analysed risk sharing channel channel, β_4 should be significantly different from zero.

Provided that we have a good match for the pre euro period, we can be relatively sure that the underlying difference in difference assumption of common trend is fulfilled. We provide an example of this in Figure 2 which shows the last dependent variable, $\Delta \log(\text{C} + \text{G})$ (the one that delivers us the coefficient of the unsmoothed component) for both the actual and the synthetic group over the whole sample period.

2.4   Data

The data that we use for our analysis comes from the OECD National Account Statistics. In particular, we use household final consumption expenditure for C, general government expenditure for G, gross domestic product computed following the output approach for GDP, net national income for NI, and net disposable income for DNI.

Our initial dataset covers 31 countries, has a yearly frequency and spans

---

[3]In the interest of space, we only report OLS with clustered standard errors and OLS with panel correlated standard errors. Results are robust to the other specifications and are available upon request.

from 1960 to 2014. However, as SCM requires the data to display no missing values, in order to keep in our sample the greatest number of countries, we limit our matching window to the period 1990-1998. This limitation leaves us with 21 countries having a complete set of data for the variables we need. Out of these 21 countries, 11 are euro area member states, while 10 are OECD countries that are not in the currency area.[4]

With the aim of increasing the number of non euro area countries in the donor pool for the synthetic control matching, we repeat the exercise with data series from the World Bank. Again, since SCM requires to have a complete dataset with no missing values, World Bank data only allows to increase the donor pool at expenses of the number of euro area countries. Therefore, in our baseline analysis we work with OECD data, while we leave World Bank data to run robustness checks.

From the actual data the SCM procedure allows to generate synthetic series for the euro area group of countries. In particular, the SCM algorithm produces the vector $W = (w_1, ..., w_N)$ of weights that maximise the matching between the actual series and the linear combinations of non euro area countries in our sample to produce the synthetic series. For the sake of clarity, Tables 1 and 2 display the optimal weights to generate the synthetic series of GDP using OECD and World Bank data respectively. For example, the synthetic GDP of Finland using OECD data is made of the Mexican, Swedish, and British GDP in the percentages of 11.3, 40.2, and 48.5 respectively. This is the linear combination of GDP series of the non euro area countries in our sample, which best matches the GDP series of Finland. For each national account series of each euro area country the algorithm finds the linear combination of national account aggregates of the non euro area countries that maximises the matching.

3 Results

3.1 Matching with the SCM

As discussed in Section 2.1, our first step is to generate the synthetic series of national account aggregates to be used as counterfactual series for the euro area countries over the period after the introduction of the euro. A sample of our match is shown in Figure 1. The figure exhibits the actual and synthetic series of household and government consumption expenditure for Finland. The two series are very close in the matching period spanning from 1990 to 1998, and then start to diverge over the post euro period. Even if only household and government consumption series for Finland are displayed here, actual and synthetic series for the national account aggregates of all euro area countries in our sample look similar to those reported. Given that the aim

---

[4]Euro area countries: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain; Non euro area countries in in our sample: Australia, Canada, Japan, Korea, Mexico, New Zealand, Sweden, Switzerland, UK, US.

of SCM is to get synthetic series, which are as close as possible to the actual ones over the matching window, our matching proves to be successful.[5]

3.2   Bilateral Correlations and the BCS Index

The results for the exercise described in Section 2.2 are shown in Table 3. As anticipated in Section 2.2, we compute difference in difference estimates of three indicators of risk sharing: $\text{Corr}(C_i, C_j)$, $\text{Corr}(C_i, Y_j)$, and $\text{BCS}_{i,j}$ for any pair of countries (i, j). The numbers shown in the table are the average of pairwise correlations or BCS indices across our panel of countries. The first column shows the estimates for the full sample, the second column excludes the crisis period, while the last two columns are estimates for two subsamples, only core countries and only periphery countries respectively. All our difference in difference estimates point to the same direction: a decrease in risk sharing due to the adoption of the euro. In particular, the first row of the table displays a significant increase in the difference in difference estimate for $\text{Corr}(C_i, Y_j)$. This indicates a reduction in risk sharing, given that with the adoption of the euro consumption and output within countries appear more synchronised. The sample split results show that the reduction in risk sharing is even sharper in periphery countries than in core countries. By excluding the crisis period we get a dampening in the fall in risk sharing, that is a reduction in $\text{Corr}(C_i, Y_j)$. The second row of the table shows the average values for the difference in difference estimates of $\text{Corr}(C_i, C_j)$. All the values are negative (though not significant), meaning that the pairwise consumption correlation has decreased with the euro. The last row reports the average values for the difference in difference estimates of pairwise BCS indices. The estimates are negative and significant for the sample which excludes the crisis period and for the sample of only periphery countries, again implying a reduction in risk sharing due to the adoption of the euro.

3.3   Risk Sharing Channels

3.3.1   Full sample of countries

The actual and synthetic series of national accounts are used to estimate Equations 3-7. Table 4 and 5 display the results of our estimations for the full sample period, i.e. 1990-2011. Each table shows both the estimations for the actual and the synthetic series in the period before (pre euro) and after (post euro) the introduction of the euro. In particular, each *Synthetic* row reports the coefficient estimates computed with synthetic data, while each *Actual* row reports the difference between the coefficients computed with the actual and those computed with the synthetic data. Table 4 shows the OLS estimates with clustered standard errors, while Table 5 displays the OLS estimates with panel correlated standard errors. In both tables the pre euro period coefficients of the euro and non euro group are never significantly different from each other (i.e. the coefficients in the row *Pre euro - Actual* are

---

[5]In order to more rigorously assess the goodness of our matching, we will discuss the non-significance of the pre-euro coefficients for synthetic series in the difference in difference estimation in Section 3.3.1 and some placebo tests in Section 4.4.

not significant), implying that the quality of our match is good. In the pre euro subsample (1990-1998), international risk sharing significantly happens only through the channel of international transfers, which absorb 4% of the shocks to GDP. Most of the shock absorption happens through consumption smoothing via public and private savings, which absorb respectively 14% and 35% of GDP shocks. The unsmoothed portion of GDP fluctuations is 50%. The estimates of the effect of the introduction of the euro are displayed in the row *Post euro - Actual*. With clustered standard errors, none of the coefficients for the smoothing channels is significant. To the contrary, the coefficient for the unsmoothed component is significant. In particular, we find that the introduction of the euro increases the unsmoothed component of the shock by 18%. With panel correlated standard errors also the private saving channel is significant. As a result of the reduction of smoothed variations in the private savings channel, the overall ability to absorb income changes is reduced by about 18% (coefficient of the unsmoothed component) compared to the scenario of no euro adoption.

A potentially surprising result is that we find no evidence of an effect of euro membership on pure *international* risk sharing, meaning through capital markets and international transfers. This suggests that the elimination of exchange rate risk has not generated an increase in the component of output variation smoothed through cross border lending and foreign direct investment. What might have reduced the unsmoothed component of the shock is private savings, which is not intratemporal cross border risk sharing, but rather intertemporal consumption smoothing. The *Post euro - Actual* coefficient for the private risk saving channel in Table 5 is negative and significant at 10% (we will also show later in Table 8 that the same coefficient is even more negative and significant for the periphery countries only). One possible explanation for this result could be that the decrease in consumption smoothing is the consequence of an increase in GDP growth and volatility due to the adoption of the euro. It could be argued that the common currency has triggered a boost in GDP growth for the countries which adopted it, as it has eliminated exchange rate risk and increased cross-member trade – at least before the outburst of the 2008 financial crisis. Figure 3 displays how much the actual GDP has increased compared to its synthetic counterpart. In the left-hand side panel the series of GDP for actual (blue line) and synthetic (red line) countries show that, with the adoption of the euro, countries display on average higher GDP. The right-hand side panel exhibits the cross-country average of the percentage difference between the actual and the synthetic series of GDP and it witnesses an increase of GDP with the adoption of the euro. Figures 4 and 5 exhibit measures of the increase in GDP volatility due to the adoption of the euro. In particular, Figure 4 displays the variance of GDP: the left-hand side chart exhibits the actual and synthetic data variances, while the right-hand panel shows the percentage difference in volatility between the actual and the synthetic series of detrended GDP with a linear quadratic trend. In Figure 5 the same graphs are shown for the coefficient of variation of detrended GDP. The left-hand side chart shows actual and synthetic sample statistics, while the right-hand side chart por-

trays the percentage difference of a coefficient of variation of detrended GDP computed from the actual and the synthetic series. The coefficient is obtained as the volatility of detrended GDP scaled by each subsample average GDP. These charts provide preliminary evidence that the currency union member states saw an increase in GDP growth and volatility after the adoption of the euro, which were higher than the ones they would have observed had they not adopted the common currency. We proceed by econometrically testing this claim by means of a difference in difference estimator. Using the same metric used in the analysis of risk sharing channels, we regress our outcomes of interest, namely GDP growth, the variance of GDP and the coefficient of variation of GDP, on a set of dummies spanning the possible combinations of pre euro/post euro and euro/no euro. Table 6 displays the results of this simple estimation. We find that the adoption of the euro had a positive and significant effect on GDP growth, but also on measures of volatility, thereby confirming the intuition provided by the graphs and providing a rationale for the effect we found on the ability to smooth consumption.

A general and legitimate concern regarding our main results on risk sharing channels is that they may be prone to measurement error driven bias. This may be particularly worrisome given that we are estimating our parameters on data we may have generated with error. As it is well known, random measurement error generates attenuation bias, which would bring our risk sharing channel for the counterfactual data closer to zero than the true parameter. Because of the GDP variance decomposition, however, this cannot be the case for all the parameters given the identity nature of our problem. In particular, assuming that we generate our series with random error, we can only have that the first four parameters suffer from attenuation bias, while the last one in fact can be computed as a residual. If the first four parameters are closer to zero than their true counterpart, this implies that the unsmoothed share must be higher than the true value. Since we consistently find that the unsmoothed parameter is lower in the counterfactual experiment than in the actual data and we have no reason to believe that the actual data is subject to the same measurement error, then our estimated difference in smoothed income variation can only be a lower bound to the actual value. By the same token, our estimated changes in the risk sharing channels can be viewed as lower bounds since we consistently find that the channels would be more effective in the counterfactual and, given the potential attenuation bias, we may be underestimating this change. This argument applies to the change in the private savings channel. By adding up the relevant estimates we find that the channel displays higher coefficient for the synthetic than for the actual data, implying that if we were to measure the coefficient without bias, the two would be further apart. In particular, the difference between actual and synthetic estimate for the private savings channel, which is already statistically different from zero, would be even larger.

3.3.2 Sample split in core and periphery countries

In order to evaluate potential heterogeneous effects of the adoption of the common currency we perform the same analysis on two subsamples of euro

area countries, namely core and periphery. The core euro area countries include Austria, Belgium, Finland, France, Germany, and Netherlands, while the periphery countries consist of Greece, Ireland, Italy, Portugal, Spain. Results are shown in Tables 7 and 8 for core and periphery countries respectively. The first interesting piece of evidence is that countries in the subgroups adopted the euro with different levels of risk sharing. In particular we find that core countries were able to smooth a larger share of output variations than the periphery counterpart - this can be seen from the coefficient of the unsmoothed component of the shock in the row *Pre - Synthetic*. This difference is mostly explained by the higher ability of public and private savings channels to smooth consumption. Regarding the effect of the adoption of the euro we find that for core, no channel is significant, meaning that the adoption of the euro did not affect either international risk sharing or consumption smoothing in the core countries. For the periphery countries, we find that the adoption of the euro has raised consumption smoothing through private dissavings. In line with our main results, the overall effect is an increase in the unsmoothed component of the shock by 34%.

4   Robustness Checks

4.1   World Bank Data

In order to assess the robustness of our findings, we carry out a series of robustness checks. The first one consists of using World Bank data instead of OECD data. This allows us to enlarge our non euro area group considerably, though one could argue that the newly added countries are probably not a good group of countries for our euro area sample. The lack of Disposable National Income in the World Bank dataset forces us to compute the measure from the raw series of international transfers from abroad, which is often missing also for developed countries. For this reason our euro area sample reduces to 8 countries. With this World Bank sample of countries, we implement the SCM to generate new synthetic series. For explanatory purposes, Table 2 displays the matrix of country weights to generate the synthetic GDP series. By running our GDP decomposition analysis on this synthetic dataset we find no change in the results.

4.2   Match SDF on growth rates

In our previous discussion regarding SDFs, we first matched over the consumption series of euro area countries and we then computed the SDFs with actual and synthetic data. As a robustness check, we take another route to evaluate the SDF pre and post euro. The alternative strategy is to compute the SDF on actual data and, only after, to generate a synthetic counterpart. This competing procedure is potentially more convoluted because, while usually the matching is carried out on levels, the SDF is a function of gross growth rates of consumption. To exemplify why this difference may be troublesome, consider that one wants to check the SDF of Germany under two policy regimes, pre and post euro. Then, matching on consumption levels would optimally

put weight on countries with similar levels of per capita consumption. In particular, it is likely that the counterfactual is a linear combination of developed countries. We may, instead, directly match on SDF, which ultimately results in matching on consumption growth. This could result in the counterfactual being composed of countries with completely different fundamentals, which happen to display similar dynamic behaviour as pre euro Germany. Ultimately, although both strategies are econometrically correct, their outcomes may somewhat vary, and, thus, one may have different preferences on the two competing procedures. We do not take a stand on which of the two is more advisable, though it is worth mentioning that, precisely because of what explained above, they may deliver quite different results.

Both methodologies produce a very good match on the pre euro period, even though the approach that directly matches on consumption does not perform as well as the one that matches on growth rates once we compute the ensuing SDF. This happens because with the former approach the synthetic series is generated to closely resemble the level and not the growth rate of consumption. On the other hand, with the latter approach we are able to reproduce relatively well the dynamics of the SDF as we directly match on growth rates of consumption and from that we compute SDF. Figure 6 displays the actual and synthetic SDF for Greece over the period 1990-2011 with the matching window stretching from 1990 to 1998.

4.3   Match on first differences

The main results in this paper are carried out by means of a difference in difference estimation. Among the assumptions of this method, the hardest to fulfill is normally the assumption of parallel trend between the actual and the synthetic series in the pre euro period. This problem can be partially dismissed by the use of the SCM since, if the matching is successful and the synthetic series closely mimic the actual ones, the common trend assumption is implied. However this assumption, which ensures that the dynamic behaviour of the actual and the synthetic series before the euro is close enough to attribute post euro differences to the introduction of the euro itself, has to hold for the dependent variable of the regression that is carried out.

In the results discussed above our analysis is applied to first differenced data, whereas the matching is produced via covariates in per capita levels. In fact, even though our matching on levels is such that the dynamics of the synthetic series are very close to the ones of the actual series, this is not enough to ensure that the first difference data will have the same trend. To address this potential issue we replicate the matching by using already first differenced covariates and outcomes, while still maintaining pre euro averages in levels. In other words, when using covariates X_0, we actually match on $\{\Delta X_{0,t}\}_{t=0}^{T^*-1}$ and $\overline{X}_0$, where the bar variables stand for pre euro period averages.

The reason for this matching strategy is that we want to replicate as closely as possible the first differenced data, hence the matching on ΔX_0. The drawback of this methodology is similar to the one discussed in the previous subsection

for the matching on levels or growth rate of consumption. By replicating the first differenced data we may find some countries with similar year-to-year changes, but very distant fundamentals from the actual series to be an excellent match. In order to shield against this possibility we keep some predictors in levels and match with a relatively homogeneous non-euro area group, namely OECD countries.

The results of this estimation are displayed in Tables 9 and 10. As can be observed, all our main results are confirmed, though their magnitude is partially reduced. This leads to a reduction in the ability of private savings to absorb income variation by around 9% and results in an overall increase in the unsmoothed component of 8%.

4.4 Placebo test

A standard check to evaluate the robustness of an estimated treatment effect are placebo tests. In our case this involves a match on the pool of non euro area countries that have never adopted the euro, as if they had actually adopted it. Hence, in this section we try to find the best match for a country like the US, which has never adopted the euro, as a linear combination of other countries that have never done so. The idea behind this methodology is that if we were to find any effect of the adoption of the euro on countries that have never been adopted it, then it is possible that our euro effect is picking up some spurious correlation.

Figure 7 displays the pre euro and post euro trends of $\Delta\ln(C+G)$. The series behave almost identically in both periods, confirming the robustness of our estimated effect of the adoption of the euro. After building a synthetic dataset for all our OECD non euro area countries, we run the same risk sharing decomposition we used for the euro area countries. The results of this estimation are displayed in Table 11. All our difference in difference estimators are extremely close to zero and never significant, meaning that we find no effect of the adoption of the euro on our non euro area group.

4.5 Year of adoption of the euro

One of the identifying assumptions of SCM is that the covariates on which the matching is carried out are not affected by the adoption of the euro. If this assumption is violated the matrix of weights may be biased by the matching on series which already incorporate the effect of the adoption of the euro. It is not unlikely that some effect of the introduction of the euro took place between the announcement and the actual introduction of the physical currency. In this sense our approach is already conservative as it uses 1999 as year of the adoption of the euro. This year corresponds to the introduction of the euro as an accounting currency, while physical euro coins and banknotes entered into circulation only in 2002. Evidence of anticipation effects has however been already found – see Frankel (2010) for an application to trade. For this reason, we run our analysis again using 1997 as the year of adoption. The results of this estimation are displayed in Table 12. Our estimates are in line

with the ones presented earlier as main result. In this case, we estimate a significant increase of 15% in the unsmoothed component of an income shock.

4.6   EU Member exclusion

An additional robustness check is the exclusion of EU countries outside of the euro area from our non euro area group. The rationale for this is that countries geographically in Europe may have endogenously decided not to join the common currency, as UK, or simply be indirectly affected by the existence of the euro. For this reason we exclude these countries from our non euro area group and run the decomposition. The results are displayed in Table 13. Our estimates are very close to our main results but have now turned out not significantly different from zero. A possible explanation is that our non euro area group is now very limited since it only includes 7 OECD countries.

4.7   Exclusion of the financial crisis period

One could argue that, despite the time fixed effects, our result of a decrease in the unsmoothed component of the shock might be driven by the turbulent period of the financial crisis starting in 2008. Therefore, in our last robustness check, we exclude the financial crisis period in running the regressions. Table 14 shows that even when we restrict our sample period to 1990-2007, thus excluding the financial crisis, our baseline results are robust. In fact, the coefficient β^u for the unsmoothed component grows even further. In particular, we find that 27% of an idiosyncratic variation to GDP would remain unsmoothed against the 20% found for the full sample. This would seem to suggest that the crisis period has actually increased consumption smoothing. Table 15 and Table 16 show the regressions with the exclusion of the financial crisis for the core and the periphery countries respectively. While the coefficient for the unsmoothed component is not significant for the core countries, it is positive and very significant for the periphery countries (0.42). This means that the increase in the unsmoothed component for the full sample is driven by the periphery countries. In particular, this decrease in consumption smoothing is mainly due to a reduction in the shock absorption capacity through the private savings channel (-35%).

5   Conclusion

This paper assesses the effect of the adoption of the common currency on the ability of euro area member states to smooth consumption and share risk. We do so by building a dataset of counterfactual macroeconomic variables for the euro area countries without the adoption of the euro via the Synthetic Control Method. We run a number of econometric procedures, namely the evaluation of correlations between consumption and GDP within countries, bilateral correlations of consumption and BCS indices across pairs of countries, and the GDP variance decomposition introduced by Asdrubali et al. (1996), to evaluate the existence of this effect and the channels through which it may have occurred.

Our main result, which is robust to our different risk sharing measures and specifications, is that *international* risk sharing across euro area member states through both private and public channels has not changed after the adoption of the common currency. At the same time, we show evidence of a decrease in consumption smoothing across euro area countries for the period after the introduction of the euro. In particular, pairwise consumption correlations and BCS indices have decreased, while correlation between consumption and GDP within each country has increased. Both indicate that with the introduction of the euro, consumption smoothing has diminished. We report that the adoption of the common currency has had a positive effect on GDP growth, which has been accompanied by an increase in output volatility. We interpret our result on the lower level of consumption smoothing due to the adoption of the common currency as driven by larger shocks to GDP, which agents are not able to insure against. We provide evidence of heterogeneous effects for member states by splitting the sample into core and periphery countries. We show that the euro has not affected the consumption smoothing of core countries significantly, whereas the aggregate negative effect that we find is due to the decrease in consumption smoothing that has happened in periphery member states. Finally, we show that our results are robust to changes in the matching strategy, exclusion of potentially affected units from the group of non euro area countries, and changes of the year of the euro adoption.

A Synthetic Control Method

A.1 Synthetic matching

Figure 1 – Actual and synthetic series
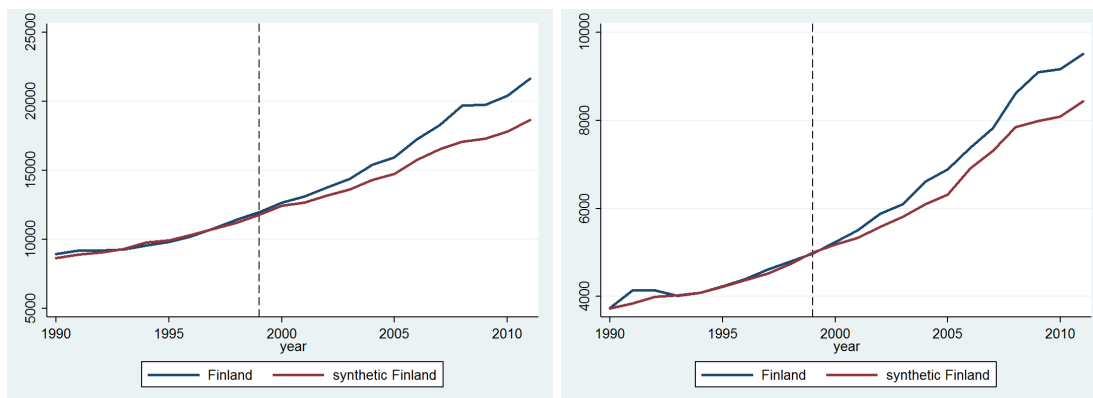
(a) Finnish household consumption   (b) Finnish government consumption

Note: The matching window is 1990-1998. The figure shows the actual series (blue lines) to be matched and the synthetic series (red lines), which maximise the matching with the blue series over the period 1990-1998.

A.2 Matrices of weights

Table 1 – Matrix of weights using OECD data (percentage points)

Non euro area	Austria	Belgium	Finland	France	Germany	Greece	Ireland	Italy	Netherlands	Portugal	Spain
Australia	33			25.70				26.70			
Canada		33									
Japan	2	39.40			44.10			2.500			
Korea							50.40			12	
Mexico	2.400	2.900	11.30	12.20	5.700	44.50		7.100	0.500	40.30	27.70
New Zealand		1.800			8.300						
Sweden	32.90	14	40.20	33.20	6.700	10.40		46.40	16.20	47.70	38.80
Switzerland	24	9		28.90	35.10			17.40			
UK			48.50				12.90		56		33.40
US	5.600					45.10	36.70		27.30		

Table 2 – Matrix of weights using World Bank data (percentage points)

Non euro area	Austria	Finland	France	Germany	Italy	Netherlands	Portugal	Spain
Brazil		14.30						
Cameroon					0.100			
Central African Republic								3
Chile	3.100					1.600	13.30	
Comoros	1.200			4.900				
Costa Rica					1			
Denmark	59		50.70	55.50		65.50	42.60	31.30
Japan	14.50		2.700	26.40		5.900		
Jordan							19	3.300
Lebanon						3		
Madagascar	1.600							
Mexico			3.800		12.80			2.200
Rwanda					1.700		2.400	
Senegal			6.800					9.900
Sweden		70.70	3.700		85.50		18.80	48.20
Switzerland	18.30	15	32.40	13.20		17		0.500
Turkey	2.200					6.100	4	1.700

Note: Table 1 and Table 2 show the matrix of weights used to generate the best linear combination of GDP of non euro area countries to reproduce the GDP of euro area countries over the matching window 1990-1998. For example, the Finnish GDP using OECD data is best reproduced by a vector of Mexican, Swedish, and British GDP in the percentages of 11.3, 40.2, and 48.5.

B Preliminary measures of risk sharing

Table 3 – Difference in difference estimates

	All countries 1990-2012	All countries 1990-2007	Core countries 1990-2012	Periphery countries 1990-2012
$\text{Corr}(C_i, Y_i)$	0.19***	0.15***	0.16***	0.23***
	(6.03)	(4.38)	(14.78)	(13.49)
$\text{Corr}(C_i, C_j)$	-0.02	-0.02	-0.02	-0.10
	(-0.29)	(-0.28)	(-0.13)	(-0.72)
$\text{BCS}_{i,j}$	-0.07	-0.18***	-0.02	-0.08**
	(-1.47)	(-2.62)	(-0.27)	(-2.05)

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis and are obtained through bootstrap by resampling the donor pool of countries to generate the synthetic variables. The table reports the average across panel of the difference in difference estimates of the measures listed in the first column – $\text{Corr}(C_i, C_j)$, $\text{Corr}(C_i, Y_i)$, and bilateral BCS Index. In order to get them, we proceed in two steps. First, we take the difference between the measure obtained using the actual and the synthetic data both pre and post euro. Then, we take the difference between the post euro and the pre euro differences. All estimates point towards a reduction in risk sharing after the adoption of the euro.

C Risk sharing channels

C.1 Estimates of risk sharing channels

Figure 2 − $\Delta\log(C+G)$ for actual and synthetic series

Note: The matching window is 1990-1998. The figure shows the actual (blue line) and synthetic (red line) series of $\Delta\log(C+G)$. The straight lines are the fitted trends to both the actual and the synthetic series before and after the adoption of the euro.

Table 4 – OLS estimated risk sharing channels - sample period 1990-2011

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre euro	Synthetic	-0.02	0.04***	0.14***	0.35***	0.50***
		(-0.20)	(3.50)	(4.43)	(3.10)	(6.57)
	Actual	-0.05	-0.01	0.00	0.06	-0.00
		(-0.37)	(-0.21)	(0.03)	(0.38)	(-0.00)
Post euro	Synthetic	0.12	-0.01	-0.11***	-0.03	0.03
		(0.97)	(-0.52)	(-3.01)	(-0.28)	(0.31)
	Actual	-0.00	-0.02	0.01	-0.17	0.18**
		(-0.02)	(-0.57)	(0.17)	(-1.23)	(2.18)
N		462	462	462	462	462
R^2		0.20	0.11	0.67	0.54	0.95

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis.

Table 5 – PCSE (het) estimated risk sharing channels - sample period 1990-2011

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre euro	Synthetic	-0.02	0.04*	0.14***	0.35***	0.50***
		(-0.30)	(1.87)	(5.28)	(5.21)	(7.18)
	Actual	-0.05	-0.01	0.00	0.06	-0.00
		(-0.59)	(-0.20)	(0.05)	(0.63)	(-0.00)
Post euro	Synthetic	0.12	-0.01	-0.11***	-0.03	0.03
		(1.46)	(-0.49)	(-3.49)	(-0.39)	(0.39)
	Actual	-0.00	-0.02	0.01	-0.17*	0.18*
		(-0.03)	(-0.53)	(0.24)	(-1.70)	(1.73)
N		462	462	462	462	462
R^2		0.20	0.11	0.67	0.54	0.95

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis.

Note: Table 4 and 5 display the results of our estimations over the window 1990-2011 for the *Actual* and the *Synthetic* series in the period before (*Pre euro*) and after (*Post euro*) the introduction of the euro. Table 4 shows the OLS estimates with clustered standard errors, while Table 5 displays the OLS estimates with panel correlated standard errors. The row *Post euro Actual* displays the effect of the adoption of the euro. With clustered standard errors, none of the coefficients for the smoothing channels is significant, while the coefficient for the unsmoothed component is significant. We find that the introduction of the euro increases the unsmoothed component of the shock by 18%. With panel correlated standard errors also the private saving channel is significant. As a result of the reduction of smoothed variations in the private savings channel the overall ability to absorb income variations is reduced by about 18%.

C.2 GDP Growth and Volatility

Figure 3 – GDP growth in euro area countries
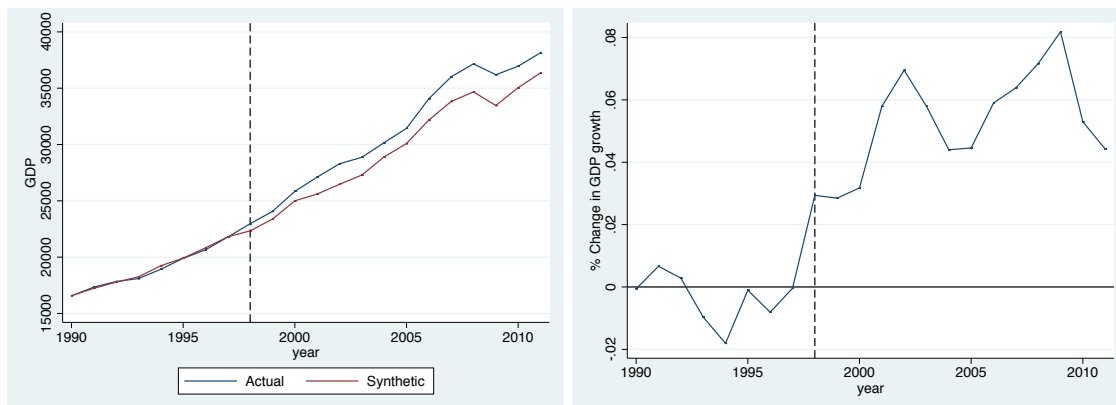
(a) Actual and Synthetic   (b) Percentage difference

Note: The left-hand side chart shows the actual and the synthetic series of GDP. The right-hand side chart exhibits the percentage difference between the actual and the synthetic series of GDP averaged across countries over the period 1990-2012. Since the adoption of the euro the percentage difference between the two has risen sharply.

Figure 4 – Variance of GDP
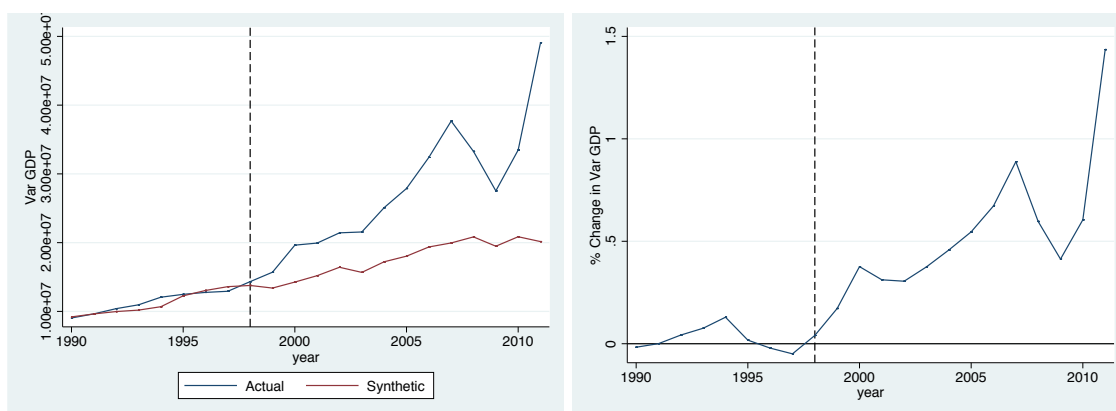
(a) Actual and Synthetic   (b) Percentage difference

Note: The left-hand side chart exhibits the actual and synthetic data variances, while the right-hand panel shows the percentage difference in volatility between the actual and the synthetic series of GDP detrended with a linear quadratic trend.

Figure 5 – Coefficient of Variation of GDP
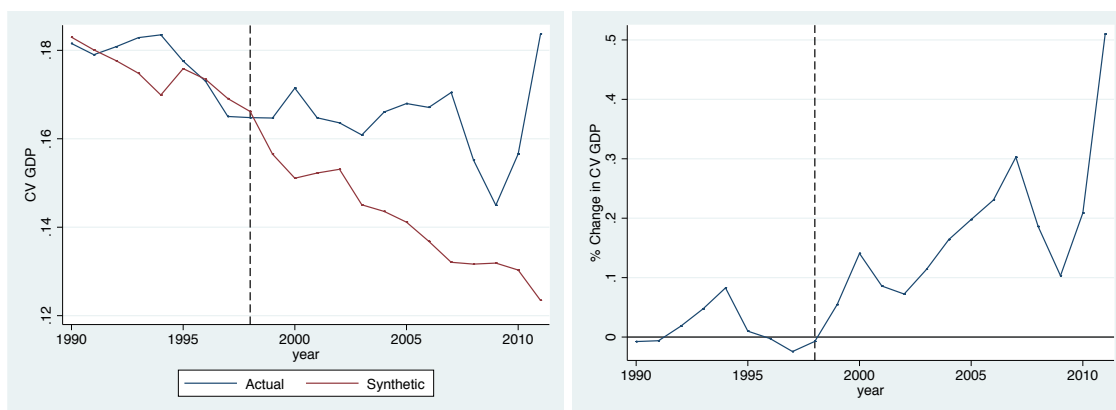
(a) Actual and Synthetic

(b) Percentage difference

Note: The left-hand side chart shows actual and synthetic coefficient of variation of detrended GDP. The right-hand panel portrays the percentage difference of the coefficient of variation of detrended GDP computed from the actual and the synthetic series. The coefficient is obtained as the volatility of detrended GDP scaled by each subsample average GDP.

Table 6 – GDP Growth and Volatility - sample period 1990-2011

		GDP Growth	GDP Variance	GDP Coeff Var
Pre euro	Synthetic	31023.08***	$3.45e + 07$***	0.15***
		(33.90)	(10.26)	(20.13)
	Actual	-1.88	320298.27***	0.00***
		(-0.00)	(6.50)	(13.46)
Post euro	Synthetic	5365.39***	$-5.03e + 06$***	-0.01**
		(5.63)	(-3.32)	(-2.35)
	Actual	7092.90***	$5.35e + 06$***	0.01***
		(4.01)	(3.53)	(2.84)
N		462	484	484
R^2		0.51	0.91	0.89

Note: In columns (2) and (3) GDP is detrended. *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis. The table displays regressions of respectively GDP growth, variance of GDP and coefficient of variation of GDP on a set of dummies spanning the possible combinations of pre euro/post euro and actual/synthetic. The results show that the adoption of the euro had a positive and significant effect on GDP growth, but also on measures of volatility.

C.3 Sample split in core and periphery countries

Core countries

Table 7 – OLS estimated risk sharing channels - sample period 1990-2011

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre euro	Synthetic	-0.17	0.04**	0.24***	0.63***	0.27
		(-1.76)	(2.52)	(8.54)	(4.25)	(1.60)
	Actual	-0.07	0.01	0.00	-0.08	0.14
		(-0.53)	(0.29)	(0.03)	(-0.37)	(0.77)
Post euro	Synthetic	0.27	-0.01	-0.17***	-0.45***	0.37**
		(1.76)	(-0.57)	(-3.82)	(-3.87)	(2.35)
	Actual	-0.02	-0.04	0.02	0.18	-0.15
		(-0.11)	(-0.75)	(0.38)	(0.92)	(-0.79)
N		252	252	252	252	252
R^2		0.32	0.16	0.74	0.62	0.96

Note: The countries included as core are: Austria, Belgium, Finland, France, Germany, Netherlands. *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis.

Periphery countries

Table 8 – OLS estimated risk sharing channels - sample period 1990-2011

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre euro	Synthetic	0.06	0.05**	0.10***	0.24**	0.56***
		(0.54)	(2.86)	(3.56)	(2.32)	(8.67)
	Actual	-0.00	-0.02	-0.03	0.06	-0.01
		(-0.02)	(-0.90)	(-0.80)	(0.55)	(-0.11)
Post euro	Synthetic	-0.02	-0.02	-0.07*	0.29***	-0.17
		(-0.13)	(-0.73)	(-2.14)	(3.27)	(-1.41)
	Actual	0.04	-0.00	0.04	-0.41***	0.34***
		(0.25)	(-0.12)	(0.88)	(-4.84)	(4.26)
N		210	210	210	210	210
R^2		0.34	0.23	0.67	0.63	0.96

Note: The countries included as periphery are: Greece, Ireland, Italy, Portugal, Spain.
*, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis.

Note: Tables 7 and 8 show the following facts. First, before the adoption of the euro core countries were able to smooth a larger share of output variations than the periphery counterpart. Second, the adoption of the euro has not affected consumption smoothing in the core countries, while it has decreased consumption smoothing in the periphery countries by 34%. This is mainly due to a reduction in the smoothing through private savings.

D Robustness Checks

D.1 Match on consumption growth

Figure 6 – Actual and synthetic series of Greek consumption growth
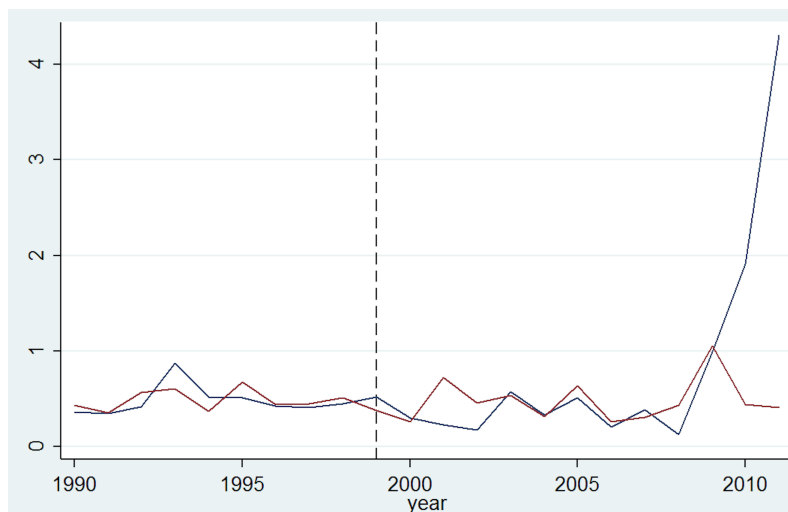
Note: The matching window is 1990-1998. The blue line is the actual series of Greek consumption growth to be matched, while the red line is the synthetic series of Greek consumption growth, which maximises the matching with the blue series over the period 1990-1998.

D.2 Match on first differences

Table 9 – OLS estimated risk sharing channels - sample period 1990-2011

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre euro	Synthetic	-0.10*	0.05**	0.12***	0.41***	0.52***
		(-1.89)	(2.18)	(4.28)	(7.06)	(9.87)
	Actual	0.01	-0.00	-0.01	0.01	-0.01
		(0.39)	(-0.29)	(-0.66)	(0.30)	(-0.21)
Post euro	Synthetic	0.12**	-0.04	-0.08**	-0.11*	0.11*
		(1.97)	(-1.39)	(-2.30)	(-1.66)	(1.78)
	Actual	0.00	-0.00	0.01	-0.09*	0.08*
		(0.10)	(-0.25)	(0.42)	(-1.90)	(1.87)
N		484	484	484	484	484
R^2		0.20	0.06	0.47	0.53	0.95

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis.

**Table 10 – PCSE (het) estimated risk sharing channels - sample period 1990-2011**

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre-tr	Synthetic	-0.10*	0.05**	0.12***	0.41***	0.52***
		(-1.71)	(2.34)	(3.65)	(6.51)	(8.57)
	Actual	0.01	-0.00	-0.01	0.01	-0.01
		(0.39)	(-0.29)	(-0.64)	(0.30)	(-0.21)
Post-tr	Synthetic	0.12*	-0.04	-0.08*	-0.11	0.11
		(1.85)	(-1.55)	(-1.89)	(-1.55)	(1.57)
	Actual	0.00	-0.00	0.01	-0.09*	0.08*
		(0.10)	(-0.25)	(0.41)	(-1.93)	(1.89)
N		484	484	484	484	484
R^2		0.20	0.06	0.47	0.53	0.95

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis.

Note: Tables 9 and 10 show that all our main results from Tables 7 and 8 are confirmed, though their magnitude is partially reduced. In particular, there is a reduction in the ability of private savings to absorb income variation by around 9%, which results in an overall increase in the unsmoothed component of 8%.

D.3   Placebo Studies

Figure 7 – Placebo Studies Δlog(C+G) for actual and synthetic series
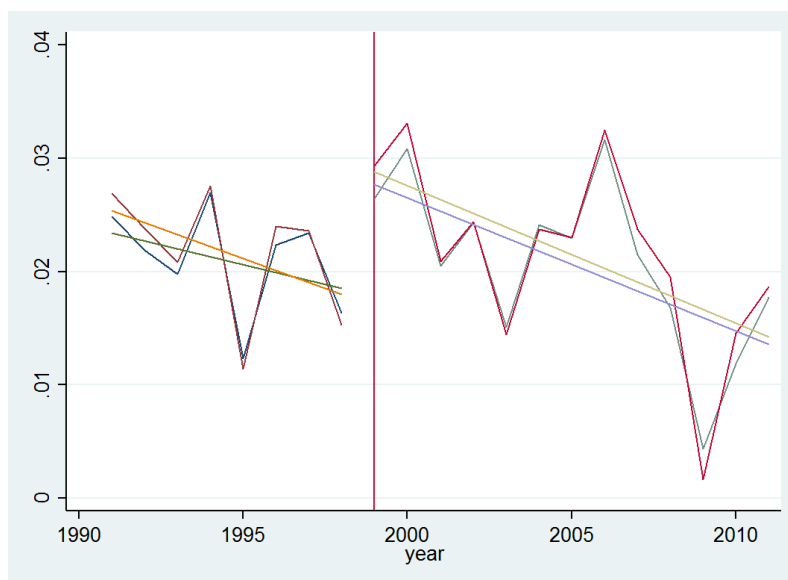
Note: The matching window is 1990-1998.

Note: The figure displays the pre euro and post euro trends of Δlog(C+G). The series behave almost identically in both periods, confirming the robustness of our estimated effect of the adoption of the euro.

Table 11 – Placebo OLS estimated risk sharing channels - sample period 1990-2011

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre-tr	Synthetic	-0.12*	0.07***	0.06***	-0.08	1.06***
		(-1.94)	(4.15)	(2.74)	(-1.04)	(14.91)
	Actual	-0.07	-0.03	0.02	0.10	-0.02
		(-0.99)	(-1.39)	(0.66)	(1.08)	(-0.17)
Post-tr	Synthetic	0.12	0.01	0.01	0.11	-0.25**
		(1.45)	(0.41)	(0.29)	(1.05)	(-2.53)
	Actual	0.03	-0.01	-0.02	-0.02	0.02
		(0.33)	(-0.53)	(-0.46)	(-0.14)	(0.14)
N		420	420	420	420	420
R^2		0.17	0.15	0.57	0.31	0.93

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis. The table shows that all our difference in difference estimators are extremely close to zero and never significant, meaning that we find no effect of the adoption of the euro on our non euro area group.

D.4   Matching over the period 1990-1997

Table 12 – OLS estimated risk sharing channels - sample period 1990-2011

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre-tr	Synthetic	-0.10	0.04**	0.16***	0.39***	0.51***
		(-1.03)	(2.20)	(6.47)	(4.77)	(9.11)
	Actual	0.01	0.00	-0.02	0.01	-0.01
		(0.09)	(0.13)	(-0.32)	(0.10)	(-0.20)
Post-tr	Synthetic	0.09	-0.01	-0.11***	-0.05	0.08
		(0.80)	(-0.60)	(-3.34)	(-0.45)	(0.79)
	Actual	0.01	-0.02	0.01	-0.15	0.15**
		(0.09)	(-0.85)	(0.14)	(-1.11)	(2.34)
N		462	462	462	462	462
R^2		0.26	0.13	0.69	0.57	0.96

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis. The table shows that our estimates when matching over the period 1990-1997 are in line with the ones presented in Table 4 as main result. We find an increase in the unsmoothed component, which is close to our baseline result and still significant.

D.5 Exclusion of EU members from non euro area group of countries

Table 13 – OLS estimated risk sharing channels - sample period 1990-2011

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre-tr	Synthetic	-0.10	0.09***	0.18***	0.41***	0.42***
		(-1.36)	(3.64)	(6.04)	(5.21)	(5.96)
	Actual	0.01	-0.05*	-0.02	-0.00	0.07
		(0.06)	(-1.82)	(-0.68)	(-0.01)	(0.84)
Post-tr	Synthetic	0.15*	-0.05*	-0.13***	-0.04	0.07
		(1.71)	(-1.78)	(-3.63)	(-0.40)	(0.84)
	Actual	-0.03	0.03	0.02	-0.17	0.16
		(-0.29)	(0.80)	(0.38)	(-1.58)	(1.61)
N		462	462	462	462	462
R^2		0.20	0.16	0.69	0.55	0.95

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis. The table shows that, when excluding countries that are part of the EU but not of the euro area, our estimates are very close to our main results in Table 4, but are now not significantly different from zero. A possible explanation for this is that our non euro area group is now very limited since it only includes seven OECD countries.

D.6 Exclusion of the financial crisis period

Full sample

Table 14 – OLS estimated risk sharing channels - sample period 1990-2007

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre-tr	Synthetic	-0.02	0.04***	0.14***	0.34***	0.50***
		(-0.19)	(3.56)	(4.44)	(3.08)	(6.56)
	Actual	-0.06	-0.01	0.00	0.06	-0.00
		(-0.37)	(-0.22)	(0.04)	(0.38)	(-0.00)
Post-tr	Synthetic	0.32**	0.03	-0.14***	-0.19	-0.03
		(2.18)	(1.11)	(-3.14)	(-1.54)	(-0.28)
	Actual	-0.16	-0.05	0.02	-0.08	0.27***
		(-0.96)	(-1.55)	(0.33)	(-0.60)	(3.09)
N		374	374	374	374	374
R^2		0.15	0.13	0.40	0.37	0.96

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis. The table shows that, when excluding the financial crisis period, the unsmoothed component of the shock increases even further as compared to the full sample period case. This confirms that the result of a decrease in risk sharing is not due to the turbulences of the financial crisis.

Core countries

Table 15 – OLS estimated risk sharing channels - sample period 1990-2007

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre-tr	Synthetic	-0.17	0.04**	0.24***	0.63***	0.27
		(-1.76)	(2.52)	(8.52)	(4.24)	(1.59)
	Actual	-0.07	0.01	0.00	-0.08	0.14
		(-0.53)	(0.29)	(0.03)	(-0.37)	(0.77)
Post-tr	Synthetic	0.41**	0.04	-0.21***	-0.60***	0.35**
		(2.22)	(1.25)	(-3.43)	(-6.61)	(2.48)
	Actual	-0.25	-0.10	0.05	0.34*	-0.05
		(-1.42)	(-1.56)	(0.71)	(1.93)	(-0.35)
N		204	204	204	204	204
R^2		0.32	0.19	0.58	0.49	0.96

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis.

Periphery countries

Table 16 – OLS estimated risk sharing channels - sample period 1990-2007

		Capital Markets	International Transfers	Public Savings	Private Savings	Unsmoothed
Pre-tr	Synthetic	0.06	0.05**	0.10***	0.24**	0.56***
		(0.56)	(2.90)	(3.58)	(2.29)	(8.79)
	Actual	-0.00	-0.02	-0.03	0.06	-0.01
		(-0.03)	(-0.91)	(-0.80)	(0.56)	(-0.10)
Post-tr	Synthetic	0.16	0.01	-0.08	0.13	-0.22**
		(0.99)	(0.25)	(-1.78)	(1.28)	(-2.35)
	Actual	-0.06	-0.03	0.02	-0.35***	0.42***
		(-0.34)	(-0.93)	(0.45)	(-3.54)	(3.82)
N		170	170	170	170	170
R^2		0.32	0.26	0.37	0.49	0.97

Note: *, **, and *** denote significance at 10%, 5%, 1% respectively. t-statistics are in parenthesis.

Note: Tables 15 and 16 show that the increase in the unsmoothed component of the shock in the full sample – see Table 14 – is driven entirely by the periphery countries, whose unsmoothed component increases by 42%. The decrease in the consumption smoothing is mainly due to the private savings channel which decreases by 35%.

References

Abadie, A., Diamond, A., & Hainmueller, J. (2010, June). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, *105*(490), 493–505.

Abadie, A., & Gardeazabal, J. (2003). The Economic Costs of Conflict: A Case Study of the Basque Country. *The American Economic Review*, *93*(1), 113–132.

Abraham, A., Carceles-Poveda, E., Liu, Y., & Marimon, R. (2016). On the Optimal Design of a Financial Stability Fund. *Working Paper*.

Asdrubali, P., Sorensen, B., & Yosha, O. (1996, November). Channels of Interstate Risk Sharing: United States 1963-1990. *The Quarterly Journal of Economics*, *111*(4), 1081–1110.

Baxter, M., & Crucini, M. (1995). Business Cycles and the Asset Structure of Foreign Trade. *International Economic Review*, *36*(4), 821–854.

Billmeier, A., & Nannicini, T. (2012, October). Assessing Economic Liberalization Episodes: A Synthetic Control Approach. *Review of Economics and Statistics*, *95*(3), 983–1001.

Brandt, M. W., Cochrane, J. H., & Santa-Clara, P. (2006, May). International risk sharing is better than you think, or exchange rates are too smooth. *Journal of Monetary Economics*, *53*(4), 671–698.

Campos, N. F., Coricelli, F., & Moretti, L. (2014, May). Economic Growth and Political Integration: Estimating the Benefits from Membership in the European Union Using the Synthetic Counterfactuals Method. *IZA Discussion Paper No. 8162*.

Carnot, N., Evans, P., Fatica, S., & Mourre, G. (2015, March). Income Insurance: a Theoretical Exercise with Empirical Application for the Euro Area. *Economic Papers*, *546*.

Cochrane, J. H. (2001). *Asset Pricing*. Princeton University Press.

Frankel, J. (2010, May). The Estimated Trade Effects of the Euro: Why Are They Below Those from Historical Monetary Unions among Smaller Countries? In *Europe and the Euro* (p. 169-212). National Bureau of Economic Research, Inc.

Furceri, D., & Zdzienicka, A. (2015, May). The Euro Area Crisis: Need for a Supranational Fiscal Risk Sharing Mechanism? *Open Economies Review*, *26*(4), 683–710.

Kalemli-Ozcan, S., Luttini, E., & Sørensen, B. (2014, January). Debt Crises and Risk-Sharing: The Role of Markets versus Sovereigns. *The Scandinavian Journal of Economics*, *116*(1), 253–276.

Krueger, D., & Perri, F. (2006, January). Does Income Inequality Lead to Consumption Inequality? Evidence and Theory. *The Review of Economic Studies*, *73*(1), 163–193.

Poghosyan, T., Senhadji, A., & Cottarelli, C. (2016). The Role of Fiscal Transfers in Smoothing Regional Shocks: Evidence From Existing Federations and Implications for the Euro Area. *Working Paper*.

Roodman, D. (2009). How to do xtabond2: An introduction to difference and system GMM in Stata. *Stata Journal*, *9*(1), 86–136(51).

Rungcharoenkitkul, P. (2011, October). Risk Sharing and Financial Contagion in Asia: An Asset Price Perspective. (ID 1956389).

Saia, A. (2016). Choosing the open sea: The cost to the uk of staying out of the euro. *Working Paper*.

Svensson, L. (1988). Trade in Risky Assets. *American Economic Review*, *78*(3), 375–94.

van Beers, N., Bijlsma, M., & Zwart, G. (2014). Cross-country Insurance Mechanisms in Currency Unions: an Empirical Assessment. *Bruegel Working Paper*(04).