CPB Netherlands Bureau for Economic Policy Analysis

CPB Discussion Paper | 335

The Impact of Uncertainty Shocks:

Continental Europe versus the Anglo-Saxon World

Joris de Wind
Katarzyna Grabska

The Impact of Uncertainty Shocks: Continental Europe versus the Anglo-Saxon World*

Joris de Wind[†] Katarzyna Grabska[‡]

September, 2016

Abstract
We show empirically that uncertainty shocks are followed by deeper recessions in Continental Europe than in the Anglo-Saxon World. Moreover, our variance decomposition indicates that the conditional variance of economic activity related to only uncertainty shocks is much larger in Continental Europe. We associate these findings with country heterogeneity in labor and capital market flexibility, since firms are less capable to deal with uncertain situations when investment and hiring decisions are less easy to reverse, as suggested by Bloom (2009) and consistent with the findings of Bartelsman, Gautier, and de Wind (2016). Our empirical results are based on a structural Vector Autoregression with a similar specification as Bloom (2009).

*We would like to thank our discussant Massimo Giuliodori for his comments and suggestions. Views expressed in this paper are our own and do not necessarily reflect those of the CPB.

[†]CPB Netherlands Bureau for Economic Policy Analysis, email: j.de.wind@cpb.nl.

[‡]CPB Netherlands Bureau for Economic Policy Analysis, email: k.grabska@cpb.nl.

Nederlandse samenvatting (summary in Dutch)

Voor het voorspellen van de economie is het belangrijk om te weten hoe de economie reageert op grote gebeurtenissen zoals het uiteenvallen van de Sovjet-Unie, de Griekse schuldencrisis, de recente terroristische aanslagen in Europa en de Brexit. Naast mogelijke directe effecten neemt de onzekerheid over de economie toe met als gevolg dat bedrijven hun investeringen uitstellen. Dit paper kwantificeert de economische gevolgen van de toegenomen onzekerheid.

Voor verschillende landen, waaronder de Verenigde Staten, Canada, het Verenigd Koninkrijk, Duitsland, Frankrijk, Nederland en een aantal kleine open economieën, hebben wij de effecten van een toename in onzekerheid in kaart gebracht. Onze empirische resultaten laten zien dat de negatieve gevolgen van een toename in onzekerheid groter zijn in continentaal Europa dan in de Angelsaksische wereld. Een mogelijke verklaring voor deze tweedeling is dat arbeids- en kapitaalmarkten over het algemeen flexibeler zijn in de Angelsaksische wereld dan in continentaal Europa, waardoor bedrijven in continentaal Europa meer moeite hebben om te opereren in tijden van onzekerheid. Deze verklaring is consistent met de theoretische analyse van Bloom (2009) en ook met de resultaten van Bartelsman, Gautier en de Wind (2016).

De empirische resultaten uit dit paper kunnen worden gebruikt voor het maken van scenario's, zie bijvoorbeeld Veenendaal, Grabska, Lanser, Ligthart en de Wind (2014) voor een CPB scenario over het conflict tussen Rusland en de Oekraïne. Bovendien zal de komende MEV-raming een scenarioanalyse bevatten over de korte termijn na de Brexit en die analyse is gedeeltelijk gebaseerd op de in dit paper gekwantificeerde onzekerheidseffecten.

1 Introduction

We want to know how the real economy responds to events like the 9/11 attacks, fall of the Soviet Union, Greek debt crisis, and the Brexit. An important part of the story is that uncertainty increases dramatically after such major events, which is an obvious concern to policymakers since an increase in uncertainty may harm the real economy. The seminal paper of Bloom (2009) shows that an uncertainty shock is followed by a short sharp recession and a quick recovery, based on a structural framework and a structural Vector Autoregression (VAR) estimated on post World War II data for the United States.

Policymakers in other countries need to know how well the estimates for the US generalize to their own economy. It appears that uncertainty shocks are highly correlated across countries but that does not necessarily mean that also the real economic effects are of the same order of magnitude. To analyze this matter, we extend the analysis of Bloom (2009) by looking at a panel of countries, including the United States, the United Kingdom, Canada, Germany, France, the Netherlands, Belgium, Ireland, Denmark, Sweden, and Finland.[1] For each country, we estimate a separate structural VAR based on data that are constructed in a consistent manner for the various countries. Using the same identification scheme as Bloom (2009), we can confirm on an international basis that an uncertainty shock causes a short sharp recession that is followed by a quick recovery. However, the size of the effects differs substantially between countries.

In particular, we have found that there is a dichotomy between Continental Europe and the Anglo-Saxon World, with deeper recessions in Continental Europe. This result makes sense intuitively as Continental European countries have in general less flexible labor and capital markets than countries in the Anglo-Saxon World, which makes it harder for firms to deal with uncertain situations. Moreover, the dichotomy is consistent with the comparative statics of the theoretical model of Bloom (2009), in which labor and capital adjustment costs play a key role.

Uncertainty reduces the willingness of firms to make investments and hire new employees. In fact, with greater uncertainty firms prefer to postpone investment and hiring decisions until the uncertainty materializes, i.e. wait-and-see behavior (Dixit and Pindyck, 1994). This is especially true when decisions are irreversible or when it is expensive to reverse decisions, which crucially depends on the level of labor and capital adjustment costs. Adjustment costs are in general higher in Continental Europe than in the Anglo-Saxon World, which is a possible explanation for the result that firms are less capable to deal with uncertainty in

[1] The literature is already crowded with estimates for several countries. However, this mainly concerns large economies such as Germany (van Roye, 2011; Popescu and Smets, 2010), the United Kingdom (Mumtaz and Theodoridis, 2012; Denis and Kannan, 2013), G7 countries (Gourio, Siemer, and Verdelhan, 2013), or country aggregates (Carrière-Swallow and Céspedes, 2013; Fornari and Stracca, 2012). As a matter of fact, there are only few available estimates for small open economies and there is not much literature about systematic differences between groups of countries.

Continental Europe than in the Anglo-Saxon World.[2]

In addition, greater uncertainty might also increase the amount of precautionary savings by households (Carroll, 1997), although this mechanism is found to be less important than the wait-and-see behavior of firms, see e.g. Beetsma and Giuliodori (2012) as well as Carrière-Swallow and Céspedes (2013). For a general overview of the uncertainty literature, including a detailed exposition of the theoretical mechanisms, see the excellent survey paper of Bloom (2014). The rest of this paper is concerned with the empirical question whether there is heterogeneity between countries in the response to uncertainty shocks.

Although we wanted to follow Bloom (2009) as close as possible for our country comparison exercise, our preferred VAR specification is somewhat different. In particular, our preferred identification scheme allows financial variables to respond contemporaneously on shocks originating from the real economy, whereas the identification scheme used by Bloom (2009) implies the exact opposite. Moreover, we are not satisfied with the dummy volatility indicator used by Bloom (2009), because the dummy approach artificially switches off the persistence in volatility, induces a shift in the timing of uncertainty shocks, and makes the identification more fragile. With the dummy approach it also becomes very difficult to warrant consistency between countries, since the dummy construction requires some fine-tuning by hand. We therefore prefer to use the full volatility time series instead, as suggested earlier by Beetsma and Giuliodori (2012) among others. Our key finding, the dichotomy between Continental Europe and the Anglo-Saxon World, is clearly robust to the various alternative specifications.

The organization of the rest of this paper is as follows. In section 2, we replicate the structural VAR results of Bloom (2009) and set out the setup for our country comparison exercise. In section 3, we describe the data and explain how we have extrapolated the volatility indices. In section 4, we present our country comparison exercise including the main results for the eleven countries in our data set and results for several alternative specifications. Finally, section 5 concludes.

2 Bloom (2009) replication

In subsection 2.1, we replicate the structural VAR results of Bloom (2009) and elaborate on the impulse response function of industrial production to an uncertainty shock. The VAR model of Bloom (2009) is quite large—eight variables and twelve lags—and is estimated

[2]The idea that firms are less capable to deal with uncertainty in Continental Europe than in the Anglo-Saxon World is consistent with the findings of Bartelsman, Gautier, and de Wind (2016), who show that it has been more advantageous for firms in the US to exploit new risky opportunities that arose from the arrival of new Information and Communication Technologies (ICT) since the mid-1990s than for firms in the EU who face more stringent employment protection legislation (EPL). In particular, they show that high-risk sectors, which contribute strongly to aggregate productivity growth, are relatively small and have relatively low productivity growth in countries with strict EPL.

on monthly data for the post World War II period in the United States. The goal of the current paper is to extend the results of Bloom (2009) by looking at a panel of countries, including the United States, the United Kingdom, Canada, Germany, France, the Netherlands, Belgium, Ireland, Denmark, Sweden, and Finland. Because of data availability and comparability we are constrained to use a somewhat more parsimonious VAR model with a smaller set of variables and fewer lags than Bloom (2009). Nevertheless, we show in subsection 2.2 that we can economize on the VAR model with hardly any influence on the impulse response function of industrial production to an uncertainty shock. Furthermore, we also show that we can (and prefer to) use the full volatility time series rather than the dummy indicator with similar results, so that we can use the more parsimonious VAR model with full volatility time series for our country comparison exercise in section 4.

2.1 Full VAR model specification of Bloom (2009) and data description

The reduced-form VAR model of Bloom (2009) contains monthly US data on the following variables: the S&P 500 stock market index, a stock market volatility indicator (described below), the Federal Funds Rate, average hourly earnings, the consumer price index, average hours, employment, and industrial production in the manufacturing sector. The sample period runs from July 1962 up to June 2008 inclusive. All variables are log-transformed except the Federal Funds Rate and average hours. To abstain from low-frequency properties of the data, all variables are pre-filtered using the Hodrick-Prescott filter with smoothing parameter $\lambda = 129,600$.[3]

The variables enter the vector of endogenous variables y_t in the order given above. The vector of endogenous variables is modeled by the VAR equation

$$y_t = c + \sum_{i=1}^{p} B_i y_{t-i} + u_t \tag{1}$$

where c is an $n \times 1$ vector of constants and $\{B_i\}_{i=1}^{p}$ are $n \times n$ matrices of autoregressive parameters in the model with $n = 8$ variables and $p = 12$ lags. The $n \times 1$ vector u_t with reduced-form error terms is assumed to be distributed according to the normal distribution $u_t \sim N(0, \Sigma)$.

The relationship between the reduced-form error terms u_t and the $n \times 1$ vector with structural shocks ε_t is given by

$$u_t = A\varepsilon_t \tag{2}$$

where A is assumed to be an $n \times n$ lower-triangular matrix such that $\Sigma = AA'$. This identi-

[3]Where applicable, the variables are seasonally adjusted at their source. For more details about the data, see Bloom (2009).

fication assumption, made by Bloom (2009), and the ordering of variables given above imply that real variables can respond contemporaneously on structural financial shocks, whereas financial variables cannot respond contemporaneously on structural shocks originating from the real economy. In our replication exercise, we adhere to the ordering of Bloom (2009), although our preferred ordering is the opposite.[4]

The stock market volatility indicator is constructed from the VXO index of percentage implied volatility from the Chicago Board of Options Exchange. Because the VXO index is only available from 1986 onwards, Bloom (2009) has extrapolated the VXO index backwards to 1962 based on the percentage actual volatility. For details about the extrapolation, see Bloom (2009) or section 3. The stock market volatility indicator is then constructed from the extrapolated VXO index using an indicator function that takes on the value 1 when the HP-filtered VXO index is more than 1.65 standard deviations above its mean and 0 otherwise. Moreover, from each episode of increased uncertainty only the month with the largest peak is taken to be 1.[5,6] This yields a dummy variable that takes on the value 1 in only 17 months, which is about 3% of the sample period of 46 years, so that the identification of the uncertainty shock is based on only the most extreme events. The extrapolated VXO index is plotted in figure 1 together with a label for each episode of increased uncertainty. The figure clearly illustrates that there is a strong relationship between the VXO index and major shocks such as, for example, the assassination of John F. Kennedy, Black Monday, the 9/11 attacks, and the 2007–2008 Credit Crunch.

Figure 2 plots the impulse response function of industrial production to an uncertainty shock. The impulse is scaled such that the volatility indicator, which is a dummy variable, increases by 1 on impact, so that the shock is representative for the events labeled in figure 1.[7] The response of industrial production displays a rapid fall of almost 1% within a few months, followed by a quick recovery and a strong rebound.[8]

[4]For an extensive discussion about the ordering of variables, see subsection 4.1. In that subsection, we explain our preferred identification scheme and show that our main results are robust along this dimension.

[5]Bloom (2009) also presents several alternative volatility indicators (taking the first month of each episode of increased uncertainty, scaling the dummy by the HP-filtered VXO index, using the full time series with the HP-filtered VXO index) and shows that his main results are robust along this dimension.

[6]It is not always clear what defines an episode of increased uncertainty, i.e. it requires judgement to classify cases such as e.g. high volatility in January, just below the volatility cutoff in February, and again high volatility in March as a single period of increased uncertainty or as two separate ones.

[7]In terms of p-value, the impulse is comparable in magnitude to an impulse of about two standard deviations when the volatility indicator would have been normally distributed (instead of binomially).

[8]It is remarkable that the positive contribution of the overshoot appears to be larger than the negative contribution of the preceding recession. This would suggest that either (i) uncertainty shocks are good for the real economy in the medium run or (ii) the impulse response function is not precisely estimated. In our country comparison exercise in section 4 we find not much evidence for the former, which is consistent with the literature that has emerged following Bloom (2009).

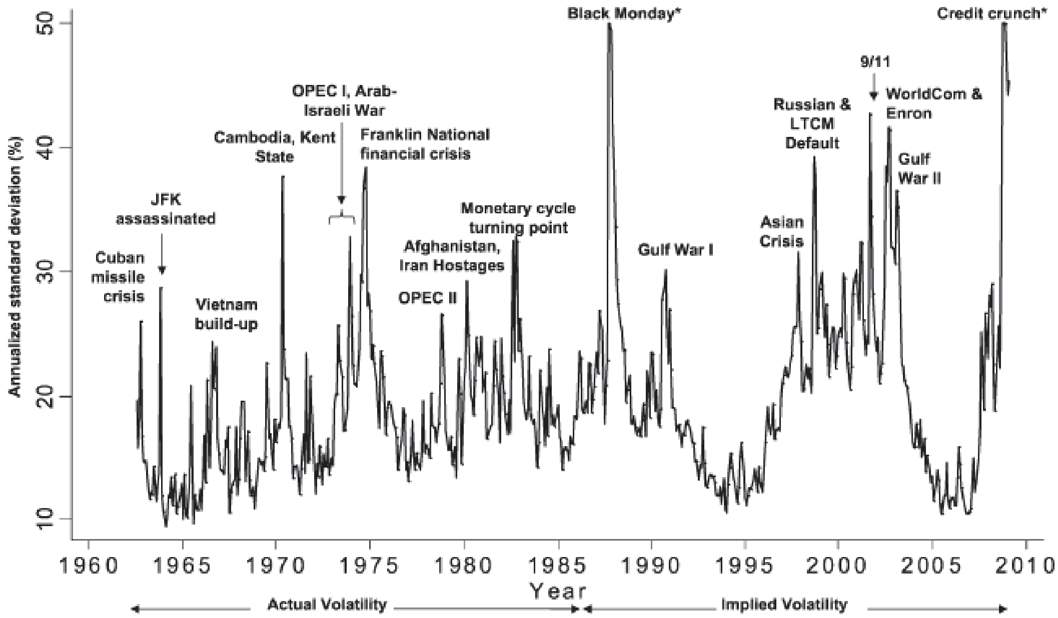

Figure 1: Extrapolated VXO index and episodes of increased uncertainty; source: Bloom (2009)

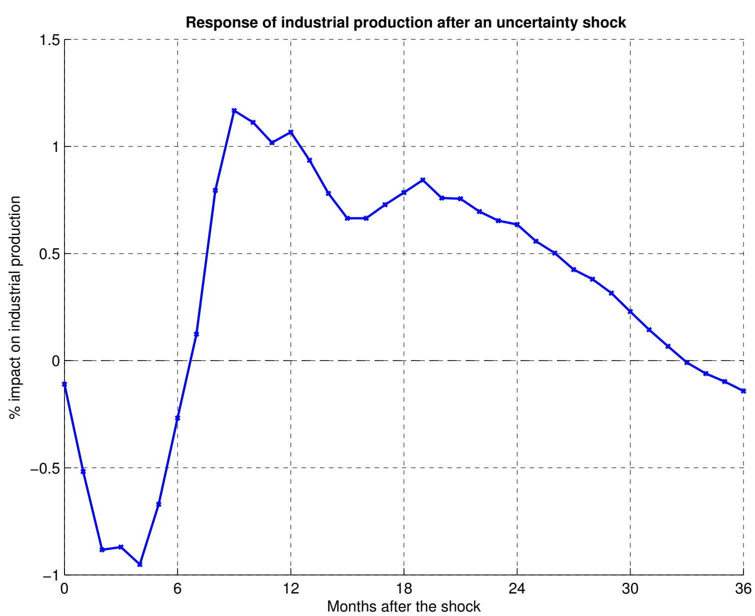

Figure 2: Impulse response function of industrial production to an uncertainty shock, impulse scaled such that the volatility indicator increases by 1 on impact; Bloom (2009) replication, full VAR model specification

2.2 More parsimonious VAR model specification

In this subsection, we show that we can economize on the full VAR model specification of Bloom (2009) with hardly any influence on the impulse response function of industrial production to an uncertainty shock. This is good news for our country comparison exercise in section 4, since we are constrained to use a somewhat more parsimonious VAR model because of data availability and comparability.

The more parsimonious VAR model consists of the following variables (in this order): the stock market index, the volatility indicator, the Federal Funds Rate, the consumer price index, and industrial production in the manufacturing sector.[9] Although we could have economized even further, we wanted to start off from a relatively standard macro VAR and then add the stock market index and the volatility indicator, which are, of course, crucial for the current application.[10] Furthermore, the more parsimonious VAR model includes four lags, which is a compromise between Bloom (2009) and several information criteria.

Figure 3 plots the impulse response function of industrial production to an uncertainty shock. The impulse response function is very similar to the one presented in figure 2, both in terms of timing as well as magnitude. In fact, like with the full VAR model specification of Bloom (2009), the response of industrial production displays a rapid fall of about 1% within a few months, followed by a quick recovery and a strong rebound.

When we replace the dummy volatility indicator by the full volatility time series we obtain the impulse response function as plotted in figure 4. The impulse is scaled as a two standard deviation uncertainty shock, which in terms of p-value is comparable to the dummy shock, yet the magnitude of the response is nevertheless smaller. The initial recession is less strong but lasts longer, presumably because in this specification the persistence of volatility is taken into account. Therefore, in our opinion and following Beetsma and Giuliodori (2012) among others, we prefer to use the full volatility time series. As a matter of fact, in view of our country comparison exercise it would be problematic anyway to use the dummy volatility indicator, since this makes the identification more fragile for the countries with relatively short sample periods and with the dummy approach it also becomes very difficult to warrant consistency between countries since the dummy construction requires some fine-tuning by hand, as explained in footnote 6.

Altogether, since the main results are robust to the various specifications, we can employ the more parsimonious VAR model with full volatility time series in our country comparison exercise in section 4.

[9] So essentially we have excluded all labor market variables. Even though Bloom (2009) has found a significant labor market response to uncertainty shocks, we had to exclude the labor market variables because of data availability and comparability as well as shorter sample periods.

[10] Our baseline VAR is comparable to Beetsma and Giuliodori (2012), whose quarterly VAR contains real per capita GDP growth, CPI inflation, the Federal Funds Rate, the volatility of the Dow Jones index, and the return on the Dow Jones index.

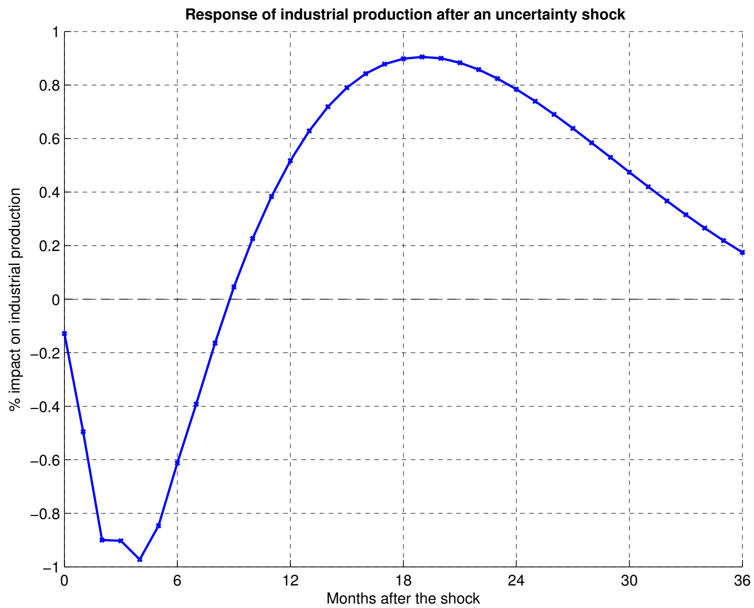

Figure 3: Impulse response function of industrial production to an uncertainty shock, impulse scaled such that the volatility indicator increases by 1 on impact; Bloom (2009) replication, more parsimonious VAR model specification

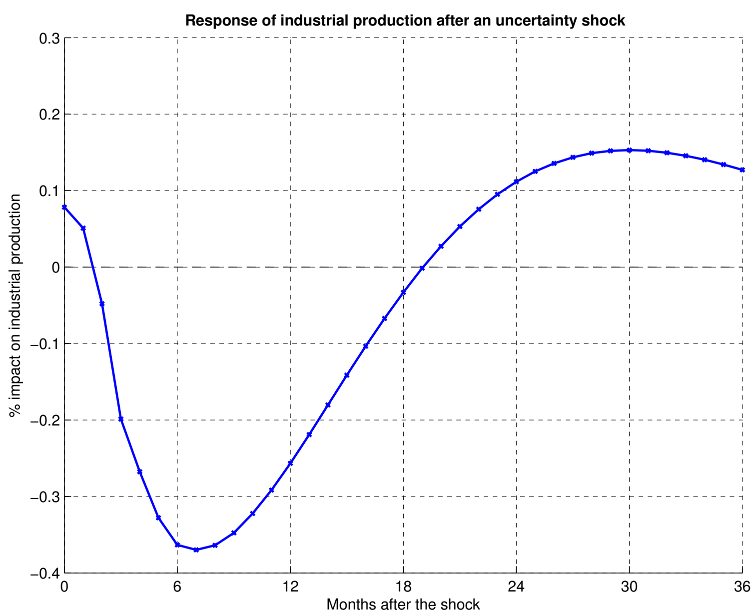

Figure 4: Impulse response function of industrial production to an uncertainty shock, two standard deviation impulse based on full volatility time series; Bloom (2009) replication, more parsimonious VAR model specification

3 Data description

The country comparison exercise in section 4 is conducted on monthly data that are comparable for the various countries. For each country, we obtain industrial production of total industry in constant prices, the short-term interest rate (three month interbank rate), and the consumer price index from Datastream.[11,12] We also take the national stock market indices (end-of-month value) and the associated volatility indices from Datastream.[13] All variables are log-transformed (except the short-term interest rate) and pre-filtered using the Hodrick-Prescott filter with smoothing parameter $\lambda = 129,600$. In subsection 3.1, we extensively discuss how we extrapolate the volatility indices.

The sample period is unbalanced. For the United States, we have data from the mid-1960s; for the Netherlands, Ireland, and the United Kingdom, we have data from the mid-1980s; for Canada, Germany, France, and Finland, we have data from the beginning of the 1990s; and for Belgium, Sweden, and Denmark, we have data from the beginning of the 2000s.

3.1 Volatility indicator

Although the concept of uncertainty seems to be illusive, there are several established measures of uncertainty. The most popular measure is a volatility index computed on the basis of options on the stock market index, based on the VIX methodology of the Chicago Board of Options Exchange. As an alternative measure, Baker, Bloom, and Davis (2016) have developed an economic policy uncertainty index for several major countries, which is based on the prevalence of specific keywords about uncertainty.[14] Other measures such as the degree of disagreement in the predictions of professional forecasters are subject to very limited country availability.

We follow Bloom (2009) and use the volatility index that is computed on the basis of options on the stock market index. Bloom (2009) uses the VXO index of the Chicago Board of Options Exchange. However, the VXO index is not available on an international basis and is, as a matter of fact, replaced by the VIX index as the preferred measure of volatility. So we take the VIX index of the Chicago Board of Options Exchange for the United States

[11] We use the data that are reported by the national statistical offices but access those through Datastream. An exception is industrial production for Canada which is provided by the IMF. Industrial production and the consumer price index are seasonally adjusted at their source.

[12] Note that the data for the United States are different from Bloom (2009). In particular, Bloom (2009) uses industrial production in the manufacturing sector, whereas we use industrial production in the total industry (because of data availability for the other countries).

[13] For the United States we take the S&P 500, for Germany the DAX, for France the CAC, for the United Kingdom the FTSE 100, for Canada the S&P/TSX 60, for the Netherlands the AEX, for Belgium the BEL 20, for Ireland the ISEQ 20, for Denmark the OMX 20, for Sweden the OMX, and for Finland the OMX 25.

[14] The economic policy uncertainty index is updated regularly and is available from www.policyuncertainty.com.

and equivalent indices based on the VIX methodology for the other countries. The various stock market volatility indices are available from Datastream.[15]

The national volatility indices are available from the beginning of the 1990s for the United States and Germany and from the beginning of the 2000s for the other countries. As in Bloom (2009), we have extrapolated the volatility indices backwards based on the actual volatility of monthly returns on the stock market indices. For each country, the actual volatility of monthly returns is first normalized to have the same mean and standard deviation as the volatility index (i.e. the series from Datastream) over the sample period for which the volatility index is available and then combined with the volatility index to obtain a much longer volatility series since the introduction of the national stock market index. The extrapolation is quite accurate, since for each country the actual volatility of monthly returns has a strong correlation with the volatility index over their common sample period (i.e. > 0.9).

4 The impact of uncertainty shocks in various countries

In this section, we investigate how different countries respond differently to uncertainty shocks. For this purpose, we estimate a separate structural VAR for the various countries in our data set. We start with the baseline VAR model specification as given in subsection 2.2 and use data that are constructed in a consistent manner for the various countries as described in section 3. After presenting the baseline results, in subsection 4.1 we present results for the VAR with our preferred ordering of variables and in subsection 4.2 we present several robustness checks. We would like to emphasize that the baseline results are based on the VAR specification that is the most comparable to Bloom (2009) but nevertheless eligible for our country comparison exercise (hence more parsimonious and no dummy approach), while our preferred VAR specification employs a different identification scheme than Bloom (2009).

Starting with the baseline VAR, figure 6 plots the impulse response function of industrial production to an uncertainty shock for the largest countries in our data set, including the United States, United Kingdom, Canada, Germany, France, and the Netherlands. The estimates clearly show that uncertainty shocks cause deeper recessions in Continental Europe (right column) than in the Anglo-Saxon World (left column). This finding is also confirmed by figure 7 which plots the impulse response functions for the smaller countries in our data set, including the Netherlands (again), Ireland, Sweden, Belgium, Finland, and Denmark.[16]

[15] The abbreviations in Datastream are the following: CBOEVIX for the United States, VDAXNEW for Germany, CACVOLI for France, VFTSEIX for the United Kingdom, GSPVIXC for Canada, AEXVOLI for the Netherlands, BELVOLI for Belgium, and SIXVXVL for Sweden. The volatility indices for Ireland, Denmark, and Finland are entirely reconstructed.

[16] For comparison purposes, the Netherlands is included in both figures.

Apart from Ireland and Denmark, for which we have problems with the identification of the uncertainty shock, the estimates clearly suggest that uncertainty shocks are followed by deeper recessions in Continental Europe than in the Anglo-Saxon World.[17]

This finding is not surprising since Continental European countries have in general less flexible labor and capital markets than countries in the Anglo-Saxon World, which makes it harder for firms to deal with uncertain situations. See, for example, figure 5 for a world map with the OECD indicator on employment protection legislation (EPL) showing that EPL is much stronger (implying less flexible labor markets) in Continental Europe than in the Anglo-Saxon World.[18] As a matter of fact, labor and capital adjustment costs play a key role in the theoretical model of Bloom (2009) and the comparative statics are consistent with the dichotomy between Continental Europe and the Anglo-Saxon World. Intuitively, the idea is that uncertainty reduces the willingness of firms to make investments and hire new employees and with greater uncertainty firms prefer to postpone investment and hiring decisions until the uncertainty materializes, i.e. wait-and-see behavior. This is especially true when decisions are irreversible or when it is expensive to reverse decisions, which crucially depends on the level of labor and capital adjustment costs. Adjustment costs are in general higher in Continental Europe than in the Anglo-Saxon World, which is a likely explanation for the result that firms are less capable to deal with uncertainty in Continental Europe than in the Anglo-Saxon World.

This idea is consistent with Bartelsman, Gautier, and de Wind (2016), who are concerned with the interaction between technology choice and labor market flexibility. They show that firms opt for relatively safe (but less rewarding) technologies in countries with strict EPL, so that innovative high-risk sectors are smaller in the EU than in the US.

Although the dichotomy between Continental Europe and the Anglo-Saxon World clearly emerges from the estimates, we do not have a clear idea about the ranking within the two groups of countries. For example, it is not clear why France displays a relatively small response to uncertainty shocks in comparison with the other Continental European countries. Other characteristics such as country size and openness might be important too, and besides, the estimation uncertainty is probably also too large to tell a meaningful story at the country level.

[17] The figures in this section contain 90% bootstrapped confidence intervals, which indicate the statistical significance of the impulse response functions. Nevertheless, the confidence intervals are not appropriate for testing the statistical significance of the observed heterogeneity between countries, for which we would need a coherent econometric framework, such as a panel structural VAR, with which we can jointly estimate the impulse response functions for the various countries. Exploiting the cross-country dimension of the data would also yield more efficient estimates making the confidence intervals more narrow.

[18] Similar statistics are available for capital markets and are supplied by, for example, the World Bank.

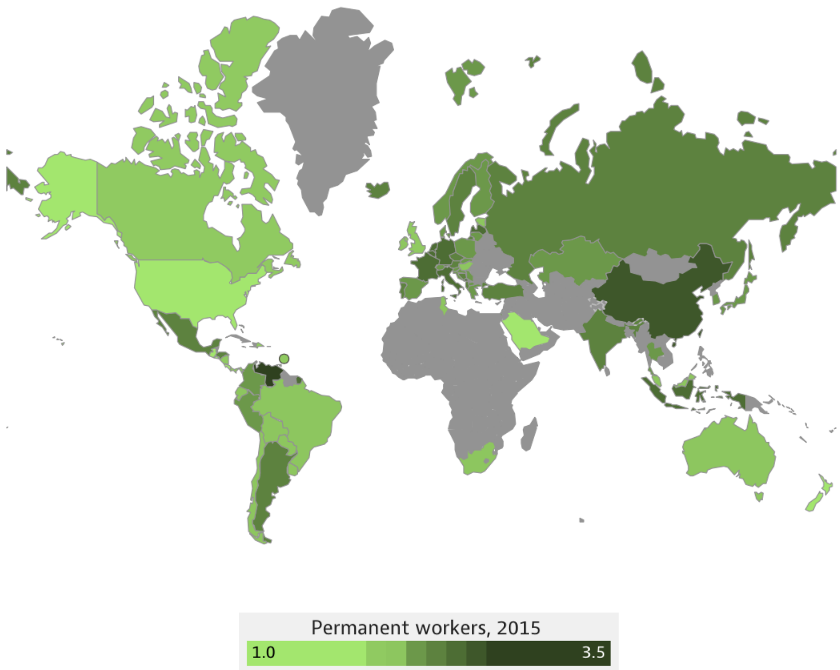

Figure 5: World map with OECD indicator on EPL, a darker color means stronger EPL; source: OECD

4.1 Our preferred identification scheme

Recall that Bloom (2009) first puts the financial variables in the vector of endogenous variables and after that the real variables. This ordering of variables, together with the Cholesky decomposition, implies that real variables can respond contemporaneously on structural financial shocks, whereas financial variables cannot respond contemporaneously on structural shocks originating from the real economy. We think this is implausible and believe, in contrast, that financial variables are much *quicker* than real variables.

Furthermore, within the financial block, Bloom (2009) first puts the stock market index and after that the volatility indicator. However, this ordering of variables is inconsistent with the finance literature, which suggests that stock market volatility has a contemporaneous effect on stock market returns, but not the other way around. See, for example, Guo (2002) for empirical evidence. Therefore, ordering the stock market index after the volatility indicator seems to be more plausible, an argument put forward by Beetsma and Giuliodori (2012).

Our preferred ordering of variables is therefore the exact opposite from the ordering used before in the baseline VAR, that is to say industrial production, the consumer price index, the short-term interest rate, the volatility indicator, and the stock market index. The rest of the model specification is unchanged relative to the baseline VAR.

Figures 8 and 9 plot the impulse response function of industrial production to an uncertainty shock for the largest and smallest countries in our data set, respectively. Our key finding that uncertainty shocks are followed by deeper recessions in Continental Europe than in the Anglo-Saxon World is somewhat stronger in our preferred VAR specification than in the baseline VAR.

An important exception, however, is the United States, but it turns out that this can be explained by the unbalancedness in the sample periods. In particular, we have re-estimated the VARs with our preferred identification scheme over a common sample period from January 1992 up to June 2014 inclusive. This yields a more consistent country comparison at the cost of throwing away a substantial amount of data.[19] Figures 10 and 11 plot the impulse response functions for the various countries. The estimates for the United States are now consistent with the dichotomy between Continental Europe and the Anglo-Saxon World. As a matter of fact, we can observe a noticeably smaller recession for the United States when the impulse response function is estimated over the shorter common sample period as opposed to the longer sample period that starts in the mid-1960s. This finding is consistent with the sample-split results of Beetsma and Giuliodori (2012), who have found that the negative response of US GDP growth to uncertainty shocks has become smaller over time.

As a further check on sample instability, we have again estimated the VARs with our preferred identification scheme over the common sample period but now excluding the financial crisis, i.e. from January 1992 up to December 2007 inclusive. Figures 12 and 13 plot the impulse response functions for the various countries. Although the impulse response functions are less significant than before, the estimates are highly suggestive that our key finding, the dichotomy between Continental Europe and the Anglo-Saxon World, is not driven (or at least not driven solely) by heterogeneous responses to the financial crisis. Clearly, uncertainty shocks have played an important role in the financial crisis and, not-surprisingly, the impulse response functions are estimated to be less strong in the sample without the financial crisis. This is, however, the case both for Continental European countries as well as for countries in the Anglo-Saxon World, and consequently our main result survives the exclusion of the financial crisis.

Finally, we have computed the variance decomposition of industrial production based on our preferred VAR specification in the common sample period and the results are given in table 1.[20] The first column provides the percentage of the total variance of industrial pro-

[19] In particular, for some countries like the United States the sample period becomes much shorter and for some other countries our data series only start at the beginning of the 2000s so that we had to exclude them from the common sample analysis. This is the case for Belgium, Sweden, and Denmark.

[20] The variance decomposition is related to the forecast error variance decomposition, i.e. they coincide

duction that is explained by uncertainty shocks, while the second column gives the absolute variance of industrial production conditional on only uncertainty shocks. It follows from the first column that uncertainty shocks have been relatively unimportant in the United States in comparison to Germany and the Netherlands, but the ranking for the other countries in our data set is not entirely consistent with the dichotomy between Continental Europe and the Anglo-Saxon World. Nevertheless, it becomes very clear from the second column that the variance of industrial production conditional on only uncertainty shocks is much larger in Continental Europe than in the Anglo-Saxon World. The reason that the first column does not seem to fully support our main result is simply that in some countries (like France and even more so in Finland) other shocks have been relatively more important too (which suppresses the percentage in the first column).

Altogether, the estimated impulse response functions and conditional variances are clearly supporting our finding that uncertainty shocks have larger effects and have played a more important role in Continental Europe than in the Anglo-Saxon World.

	% of variance of industrial production explained by uncertainty shocks	Variance of industrial production conditional on uncertainty shocks
US	5.76	0.0047
UK	11.34	0.0037
Canada	10.68	0.0020
Germany	14.47	0.0189
France	9.72	0.0079
Netherlands	17.36	0.0130
Ireland	1.96	0.0050
Finland	7.35	0.0121

Table 1: Variance of industrial production that is explained by uncertainty shocks; percentage of total variance of industrial production in the first column, absolute variance conditional on only uncertainty shocks in the second column; based on preferred VAR specification and common sample

A remaining question is whether the dichotomy between Continental Europe and the Anglo-Saxon World is driven by larger uncertainty shocks or simply by stronger responses to uncertainty shocks of similar magnitude. In the impulse response functions displayed so far the uncertainty shocks have been scaled as two standard deviation shocks, which may lead to larger increases in the volatility indicator in some countries than in other countries. Figures 14 and 15 plot the impulse response functions for the various countries for the case where for each country the uncertainty shock is scaled such that the volatility indicator increases by 10 points on impact. These normalized impulse response functions still display

when the forecast horizon goes to infinity. The variance decomposition and the forecast error variance decomposition give a very similar picture in the current example, and therefore we only present results based on the variance decomposition.

a dichotomy between Continental Europe and the Anglo-Saxon World, but the gap between the two country groups has become somewhat smaller. This suggests that the dichotomy is driven both by larger uncertainty shocks as well as by stronger responses to equally-sized uncertainty shocks in Continental European countries in comparison with countries in the Anglo-Saxon World.[21]

4.2 Robustness checks

In addition to the various alternative specifications presented above, we have checked for robustness along two additional dimensions and the results are found to be robust. In particular, we have checked what is the influence of HP-filtering the data and we have also checked whether changing the number of lags is of influence. The robustness checks presented in this subsection are relative to the preferred VAR specification and are estimated over the common sample period. For space considerations, we only present results for the large countries in our data set (but other results are available on request).

Figure 16 presents the impulse response functions when we estimate the VAR in levels rather than using HP-filtered data. Like before, the estimates suggest that uncertainty shocks are followed by deeper recessions in Continental Europe (right column) than in the Anglo-Saxon World (left column). Figure 17 presents the impulse response functions when we only include two lags in the VAR (as suggested by several information criteria) and again the estimates show a clear dichotomy between Continental Europe and the Anglo-Saxon World.

5 Concluding remarks

The key finding of our cross-country comparison exercise is that uncertainty shocks are followed by deeper recessions in Continental Europe than in the Anglo-Saxon World. Consistent with this dichotomy, we have also found that the conditional variance of economic activity (measured by industrial production) related to only uncertainty shocks is much larger in Continental Europe. Although our paper is mainly empirical, we associate the dichotomy with country heterogeneity in labor and capital market flexibility since firms are less capable to deal with uncertain situations when investment and hiring decisions are less easy to reverse, as suggested by Bloom (2009) and consistent with the findings of Bartels-

[21] We would like to note that at many occasions the same event causes a period of increased uncertainty in many countries, e.g. the events labeled in figure 1. At such occasions, the underlying common shock is the same but it nevertheless induces larger volatility increases in some countries than in other countries. Therefore, our finding that the dichotomy between Continental Europe and the Anglo-Saxon World is partly explained by larger uncertainty shocks needs some caution. As a matter of fact, the larger uncertainty shocks in Continental Europe are already a response to the same common events that also induce volatility increases in the Anglo-Saxon World. The key point here is that the common events cause larger volatility increases in Continental Europe than in the Anglo-Saxon World.

man, Gautier, and de Wind (2016). This suggests that countries with more flexible labor and capital markets are better capable to deal with uncertainty shocks.

At the CPB, we have used the results of the current paper at various occasions including a scenario analysis what we expected that would happen when the conflict between Russia and the Ukraine would have further escalated in the summer of 2014, see Veenendaal, Grabska, Lanser, Ligthart, and de Wind (2014). Moreover, shortly after the publication of this working paper, CPB will publish a new macroeconomic outlook (MEV) including various scenarios on the short run after the Brexit.

The estimated effects of increased uncertainty on the real economy have been important ingredients for the aforementioned scenarios. The idea is simply to take the estimated impulse response function for the Netherlands (based on our preferred VAR specification) and calibrate the size of the impulse based on a comparison with previous episodes of increased uncertainty, such as the ones labeled in figure 1. This yields a time path with dynamic effects for industrial production, which can then be mapped into effects for GDP and further decomposed in consumption and investment effects, based on rule-of-thumbs.[22] Of course, it would be more ideal to re-estimate the VARs with consumption and investment, but unfortunately those variables are not available (for a sufficiently long time period) on a monthly frequency for all the countries in our data set, as is the case for the Netherlands.

Finally, although it was very convenient to estimate separate structural VARs for the various countries, it is left for future research to more efficiently combine the various countries in a single framework such as a panel structural VAR. That would also yield a coherent econometric framework that is suitable to assess whether the estimated impulse response functions and conditional variances are significantly different from each other.

[22]In the Ukraine scenario, we have assumed that the GDP effect is about 70% of the effect on industrial production and following Beetsma and Giuliodori (2012) and Carrière-Swallow and Céspedes (2013) we have assumed that the investment effect is an order of magnitude more important than the consumption effect.

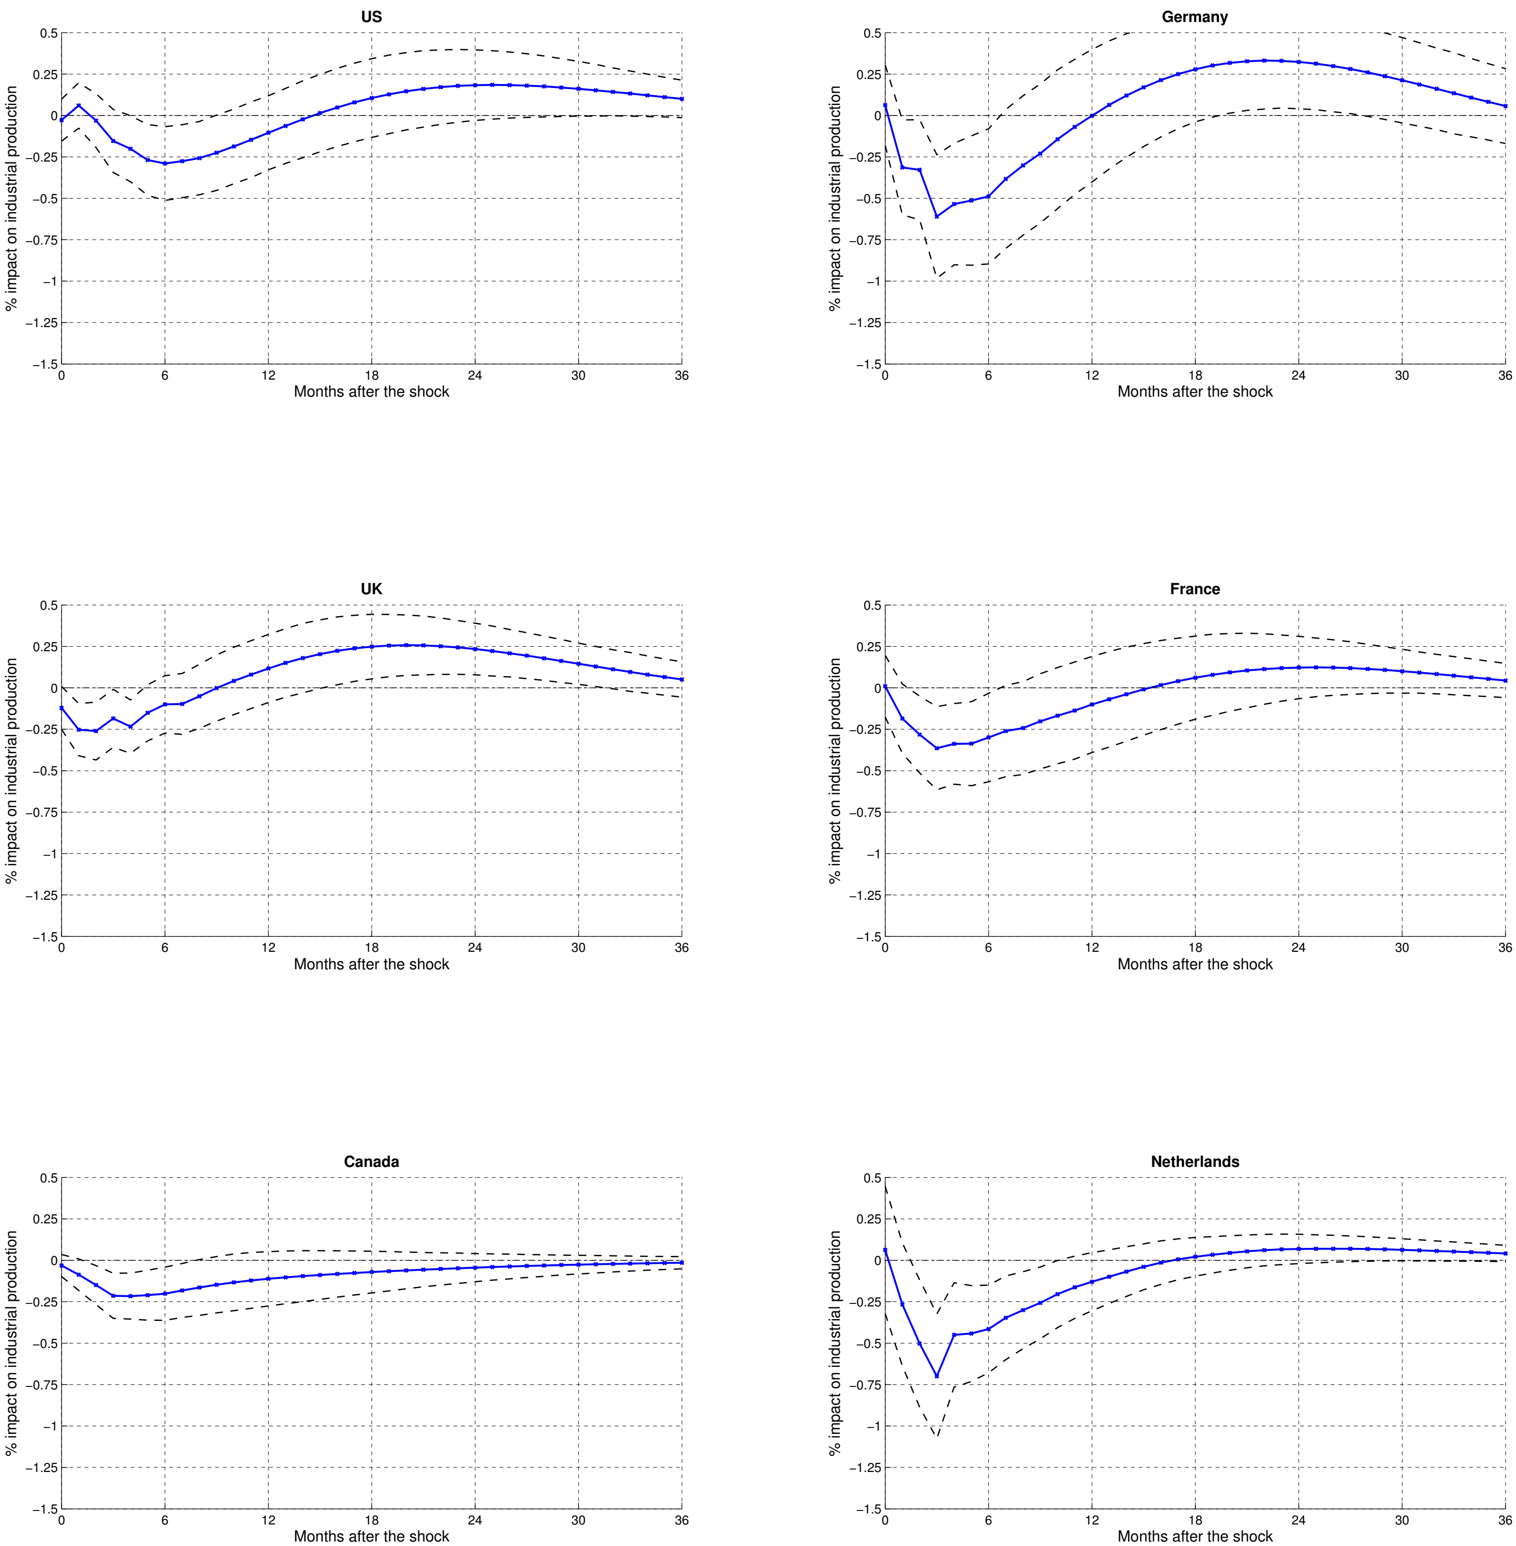

Figure 6: Impulse response function of industrial production to a two standard deviation uncertainty shock, baseline specification with full volatility time series, large countries; 90% bootstrapped confidence interval

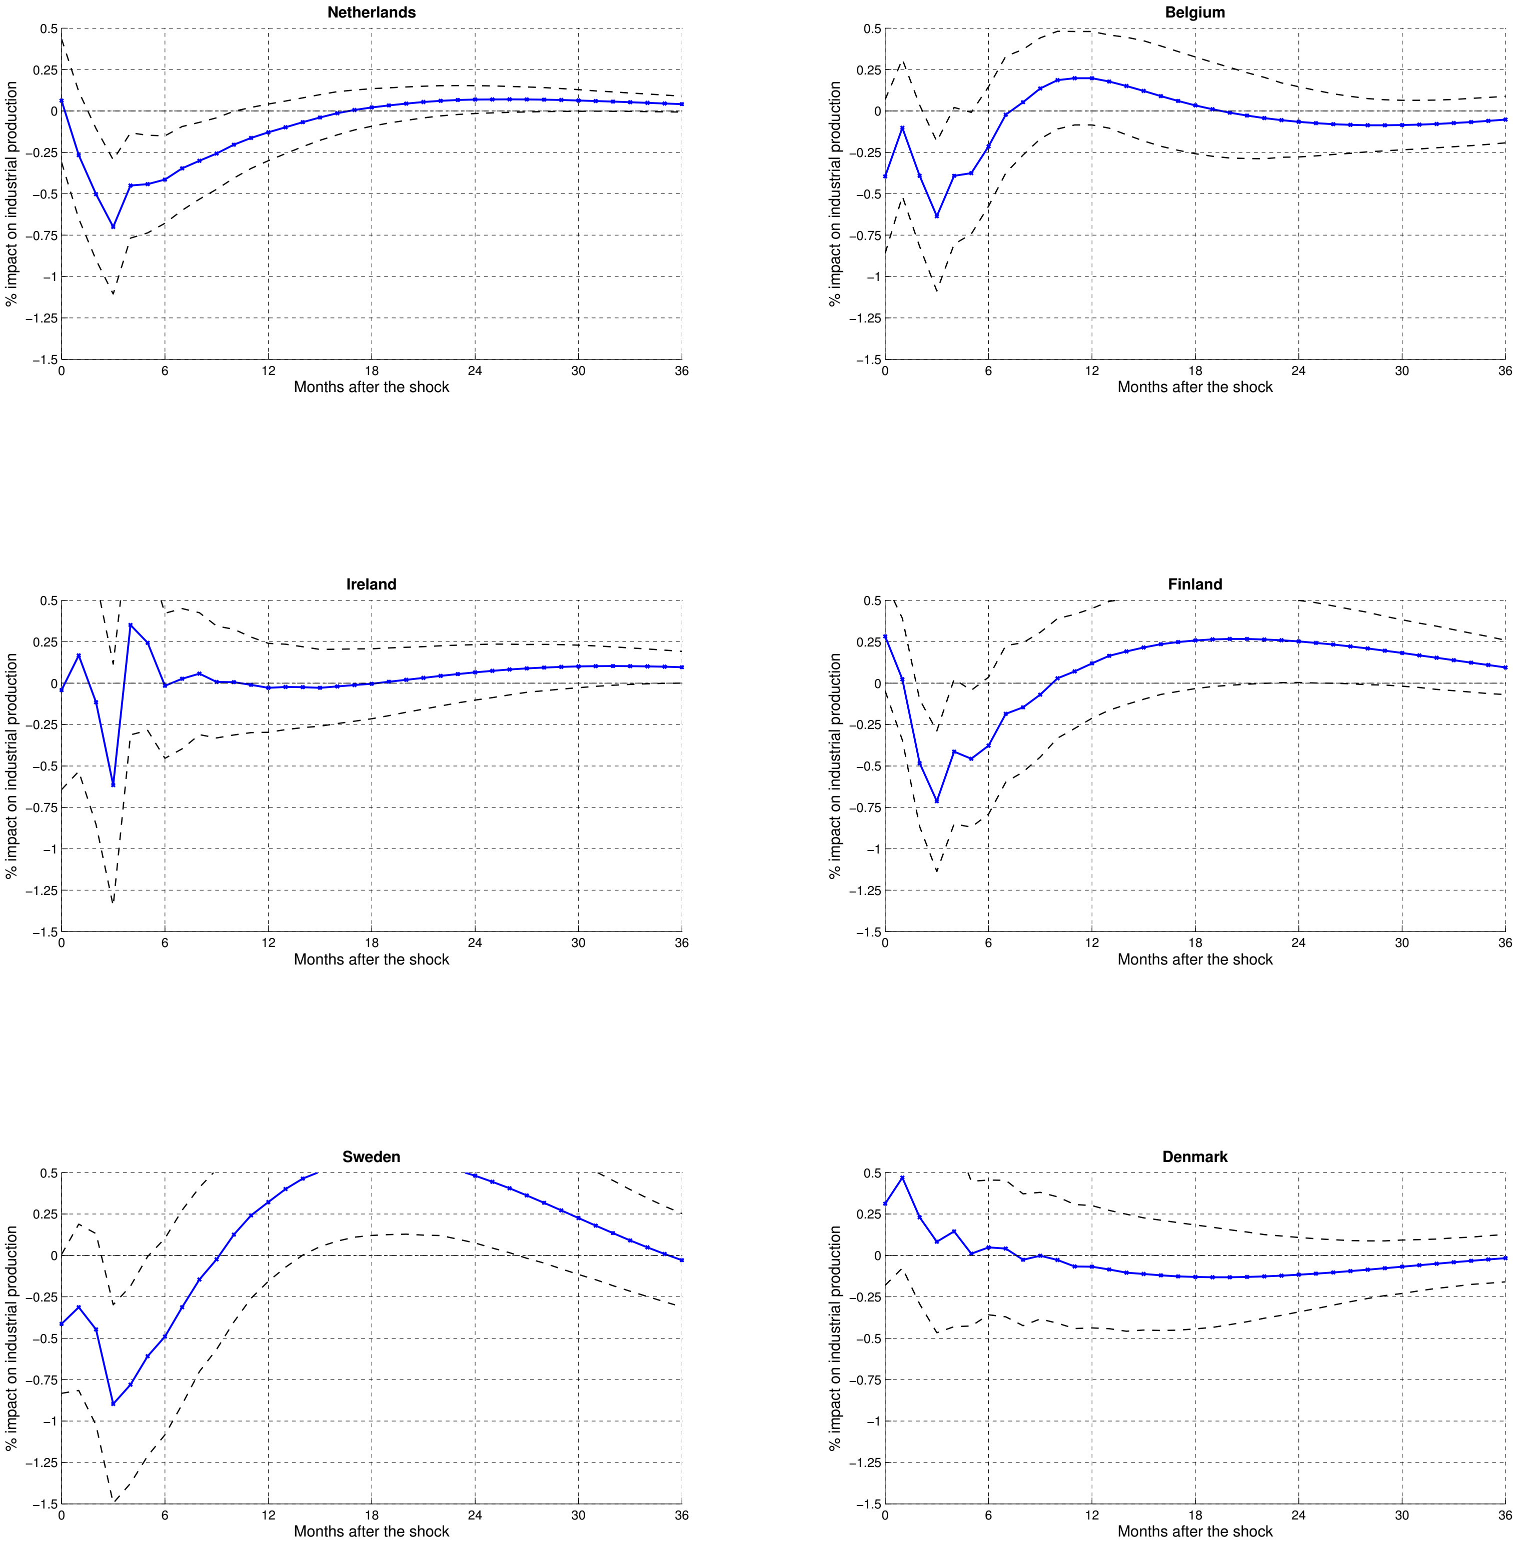

Figure 7: Impulse response function of industrial production to a two standard deviation uncertainty shock, baseline specification with full volatility time series, small countries; 90% bootstrapped confidence interval

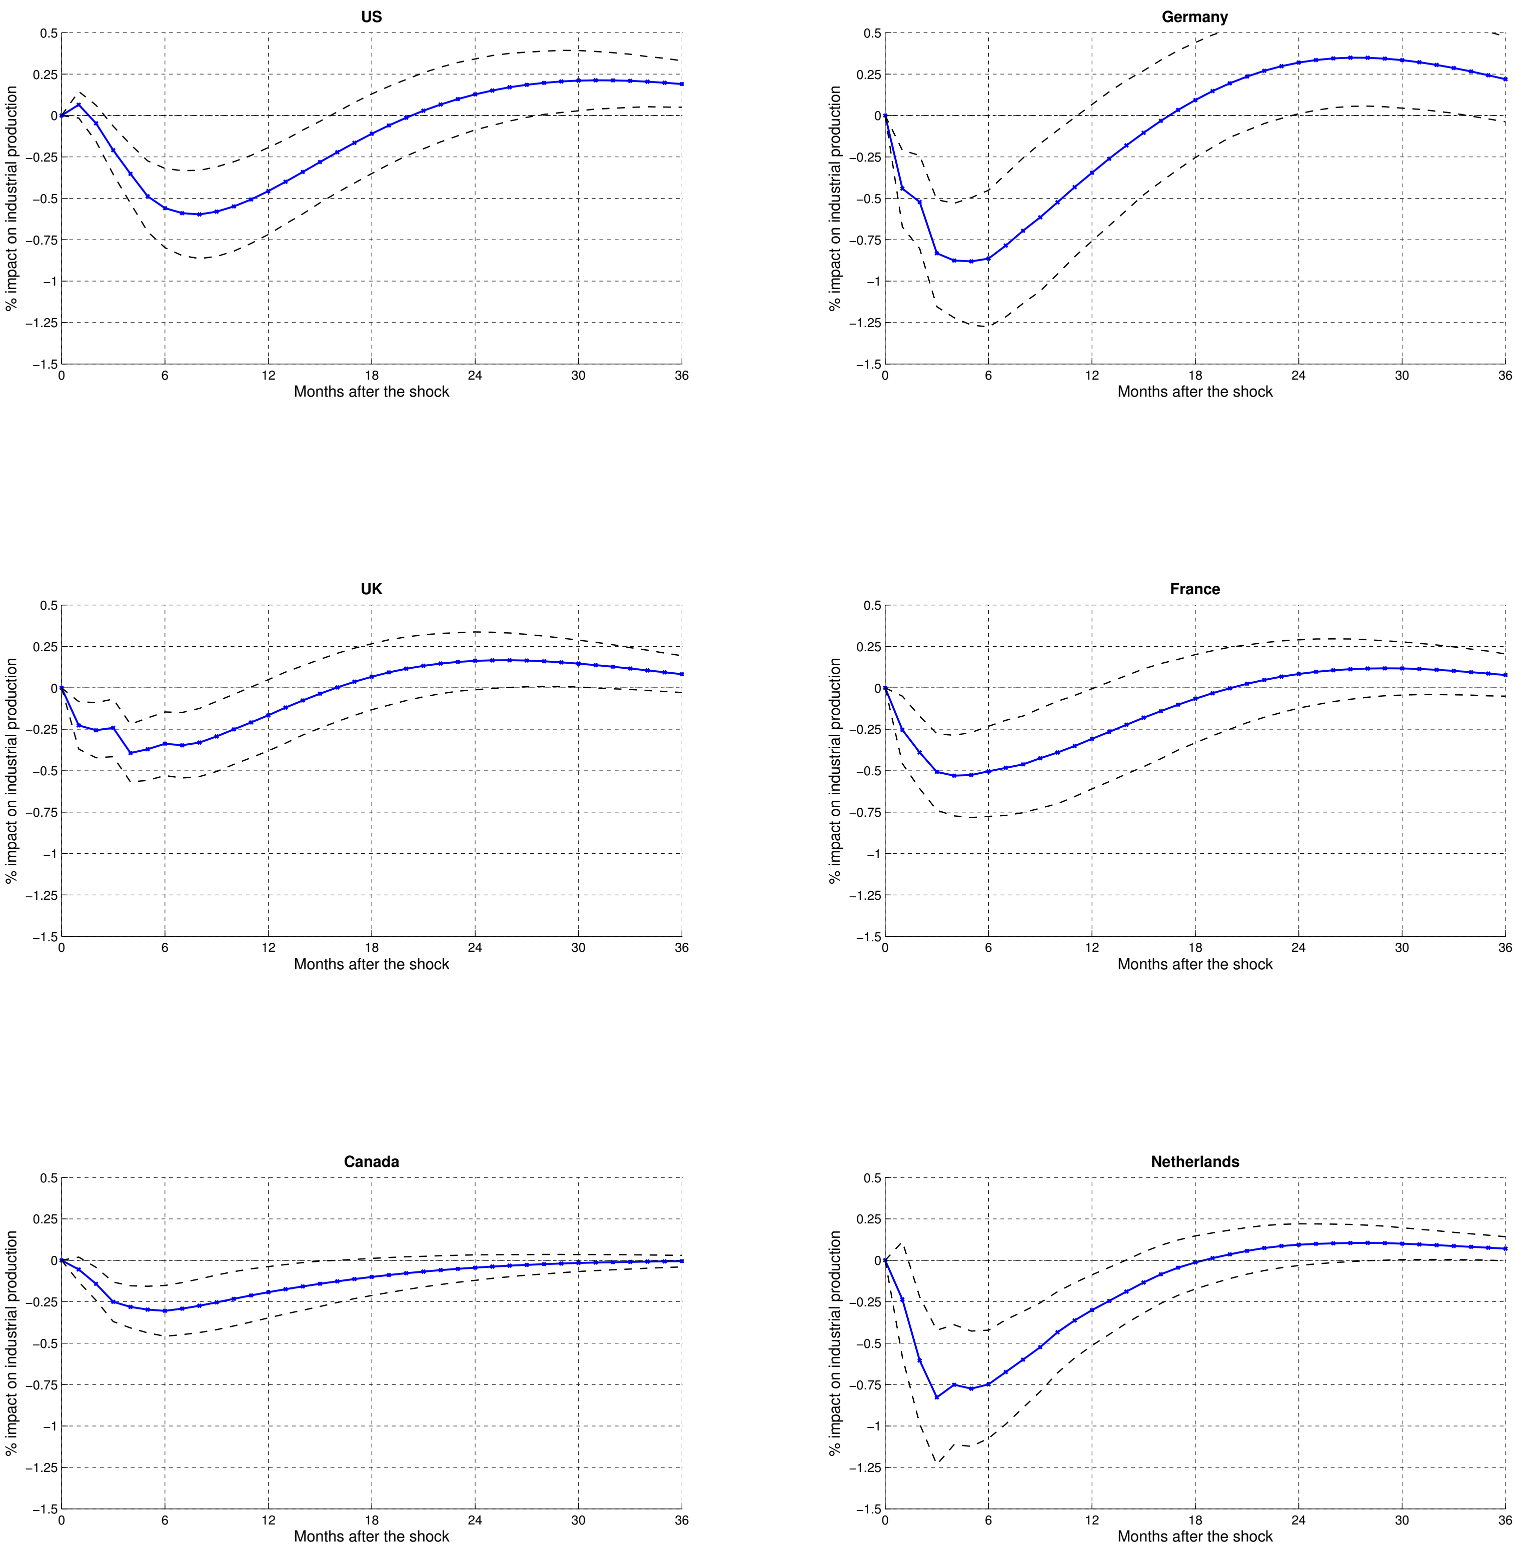

Figure 8: Impulse response function of industrial production to a two standard deviation uncertainty shock, preferred identification scheme with full volatility time series, large countries; 90% bootstrapped confidence interval

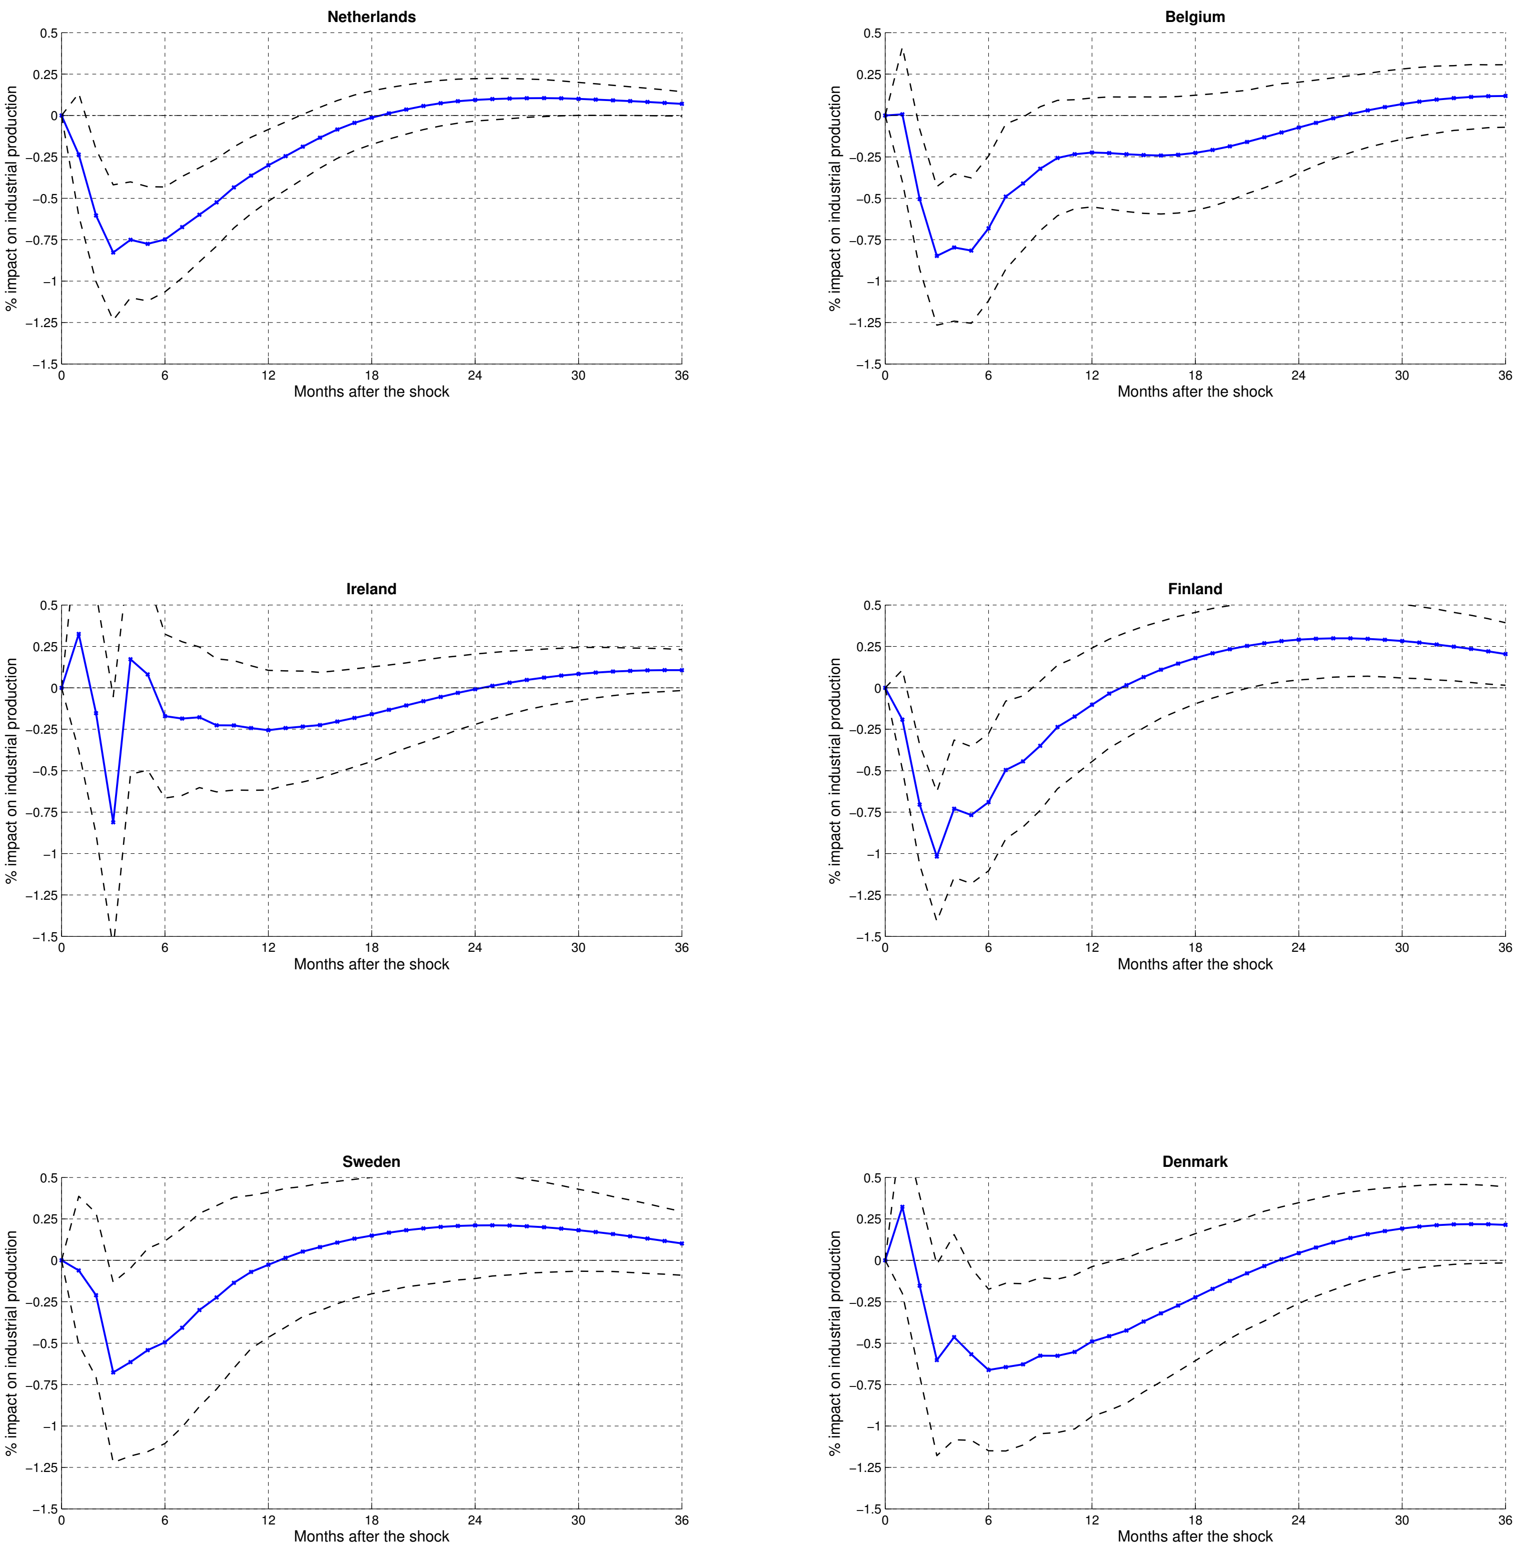

Figure 9: Impulse response function of industrial production to a two standard deviation uncertainty shock, preferred identification scheme with full volatility time series, small countries; 90% bootstrapped confidence interval

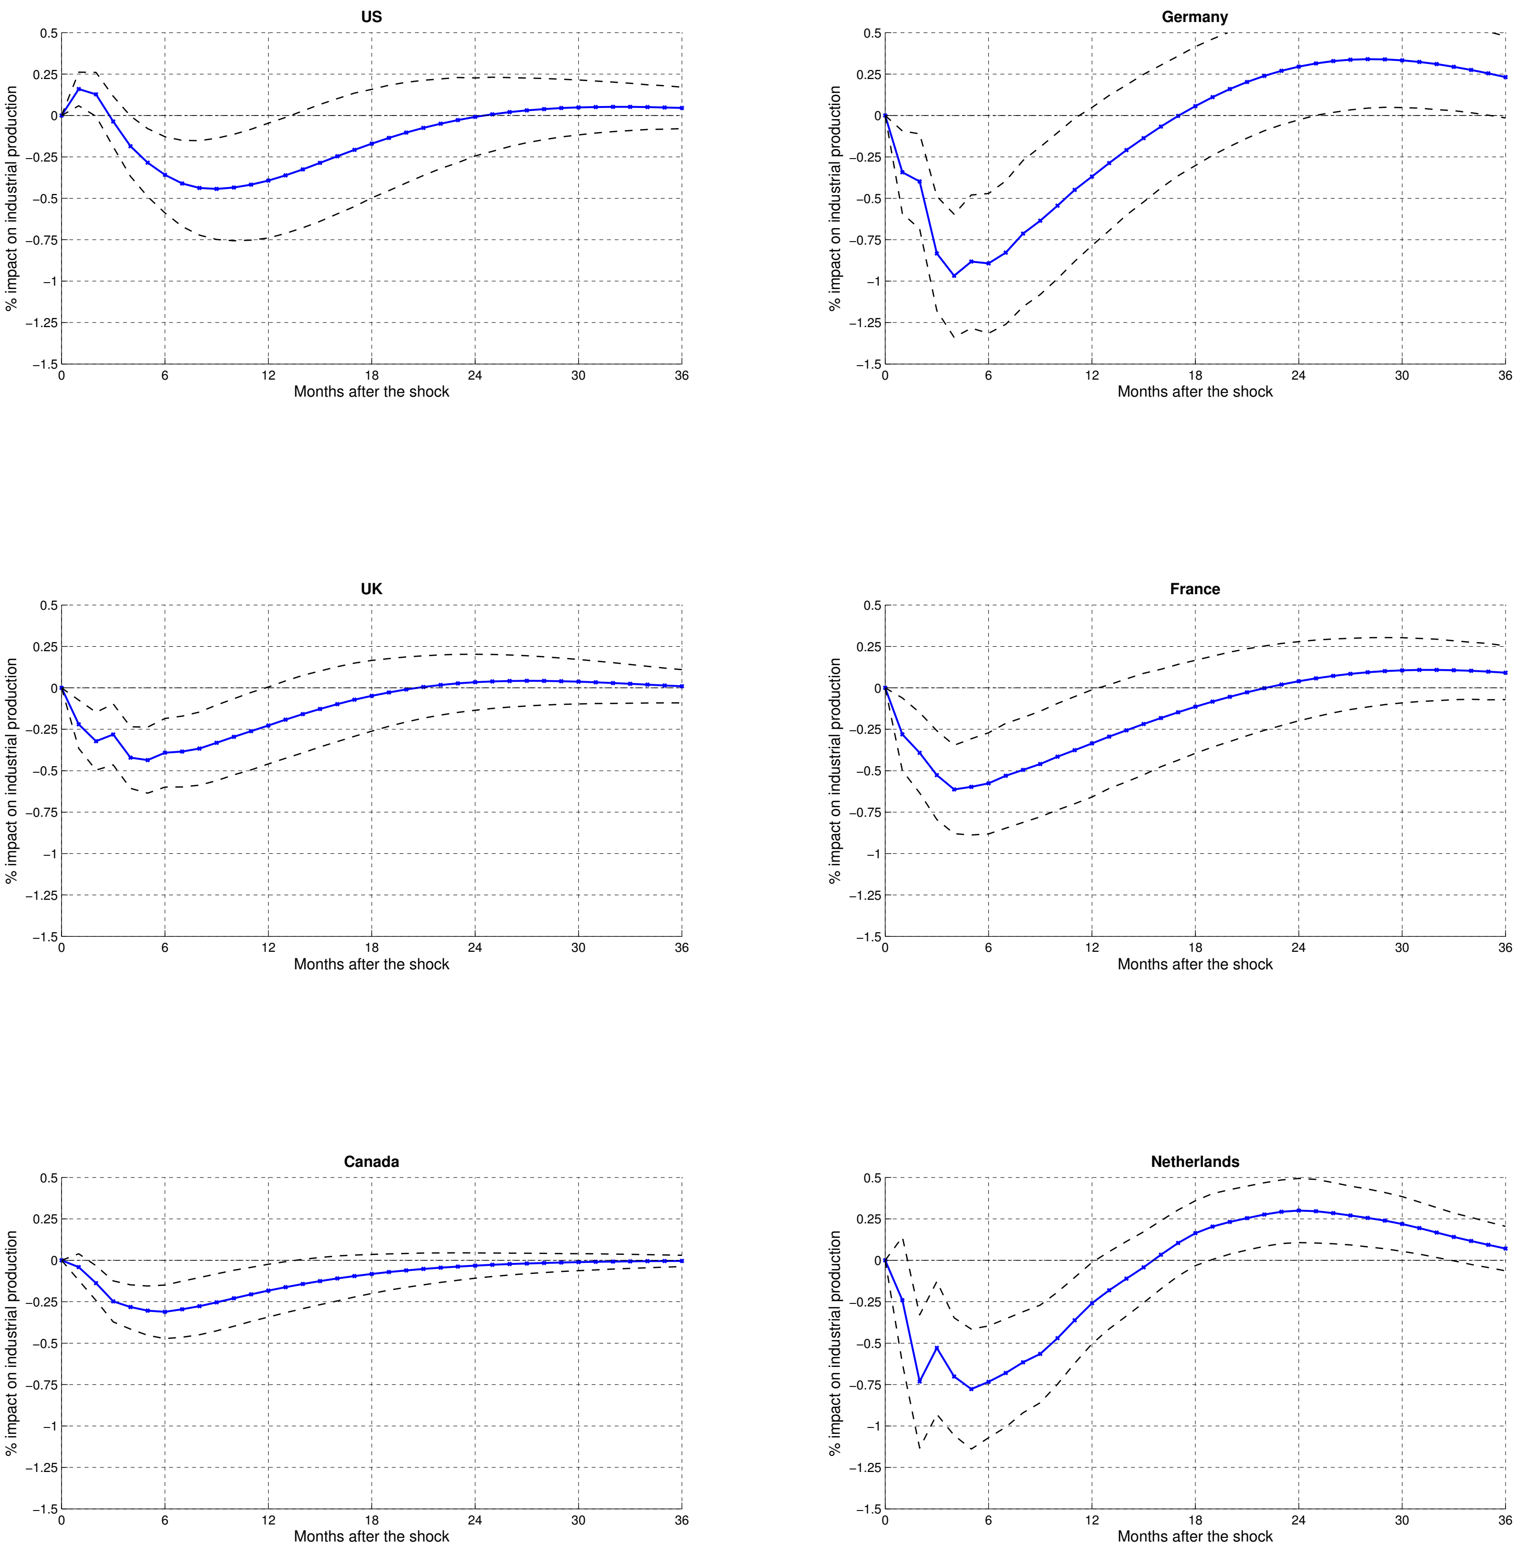

Figure 10: Impulse response function of industrial production to a two standard deviation uncertainty shock, preferred identification scheme with full volatility time series, large countries, common sample; 90% bootstrapped confidence interval

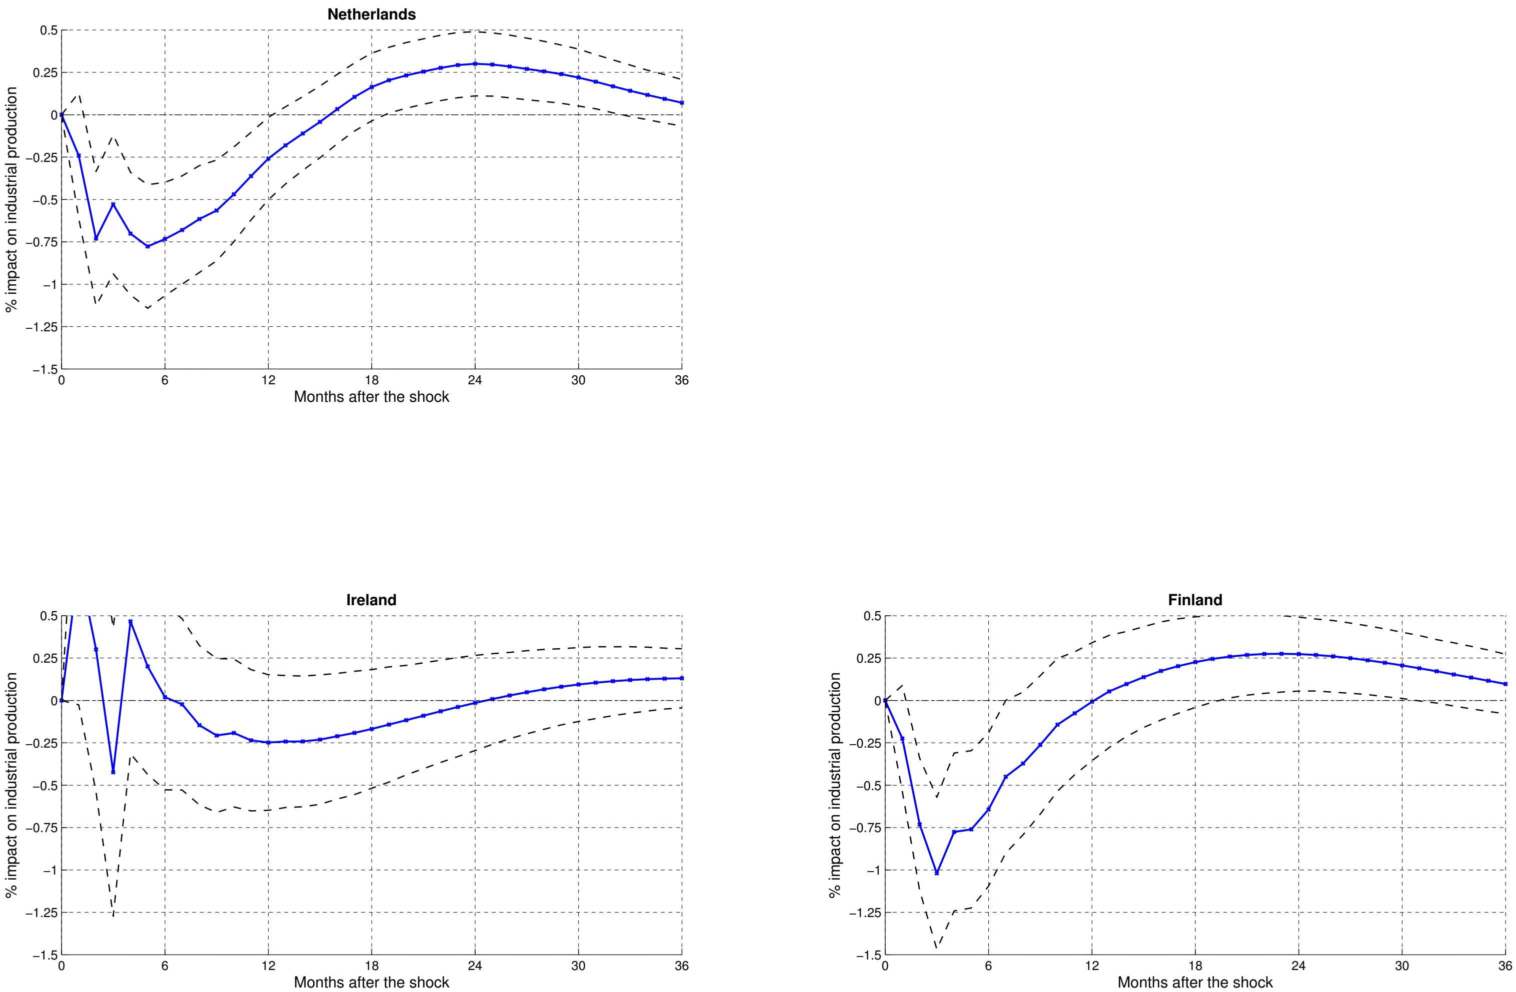

Figure 11: Impulse response function of industrial production to a two standard deviation uncertainty shock, preferred identification scheme with full volatility time series, small countries, common sample; 90% bootstrapped confidence interval

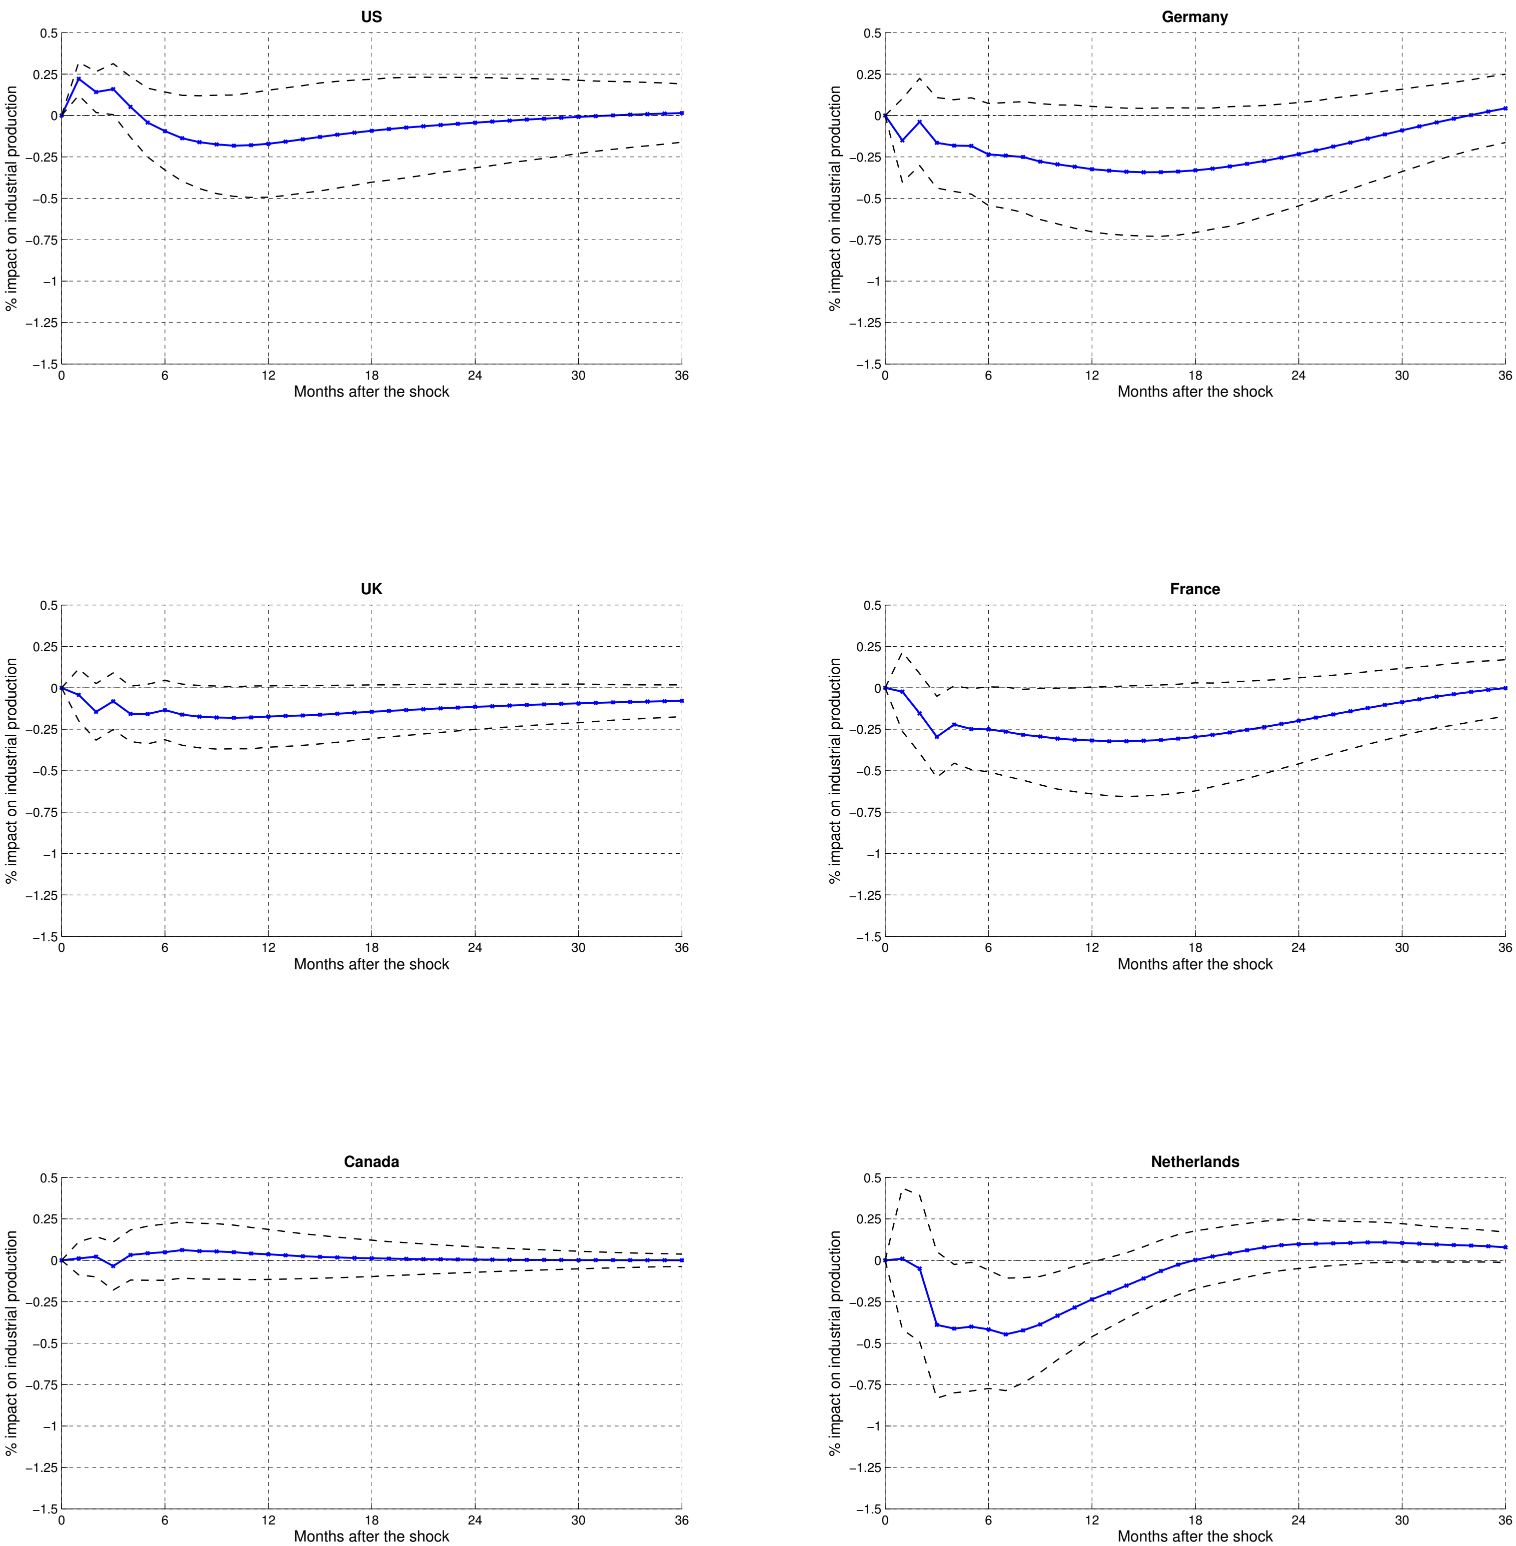

Figure 12: Impulse response function of industrial production to a two standard deviation uncertainty shock, preferred identification scheme with full volatility time series, large countries, common sample excluding financial crisis; 90% bootstrapped confidence interval

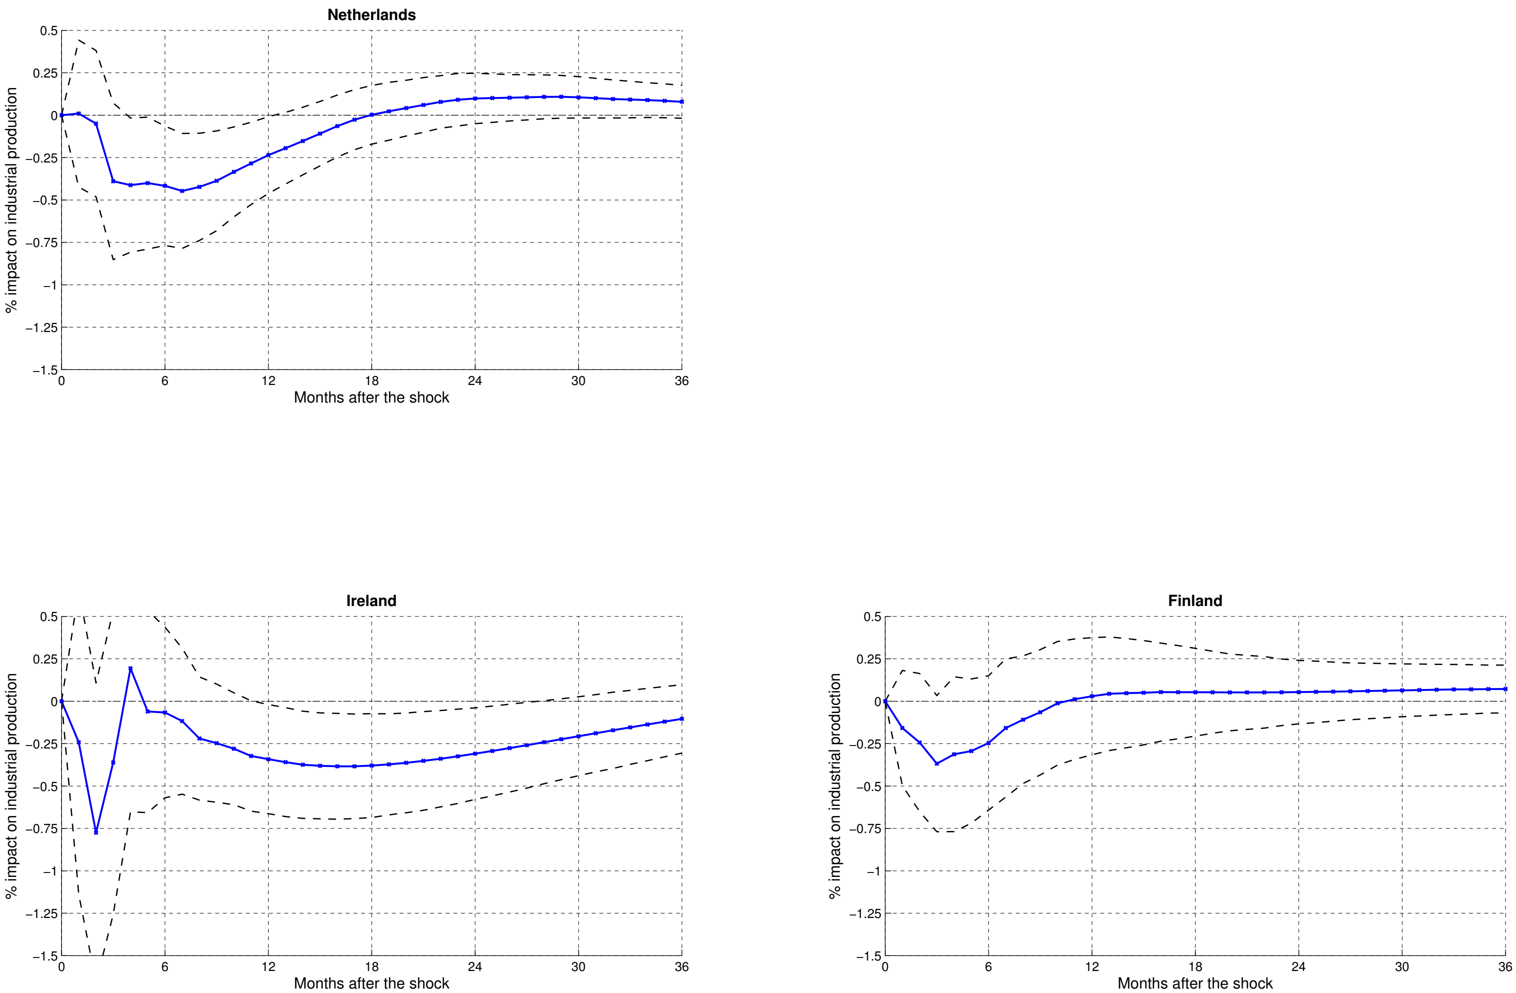

Figure 13: Impulse response function of industrial production to a two standard deviation uncertainty shock, preferred identification scheme with full volatility time series, small countries, common sample excluding financial crisis; 90% bootstrapped confidence interval

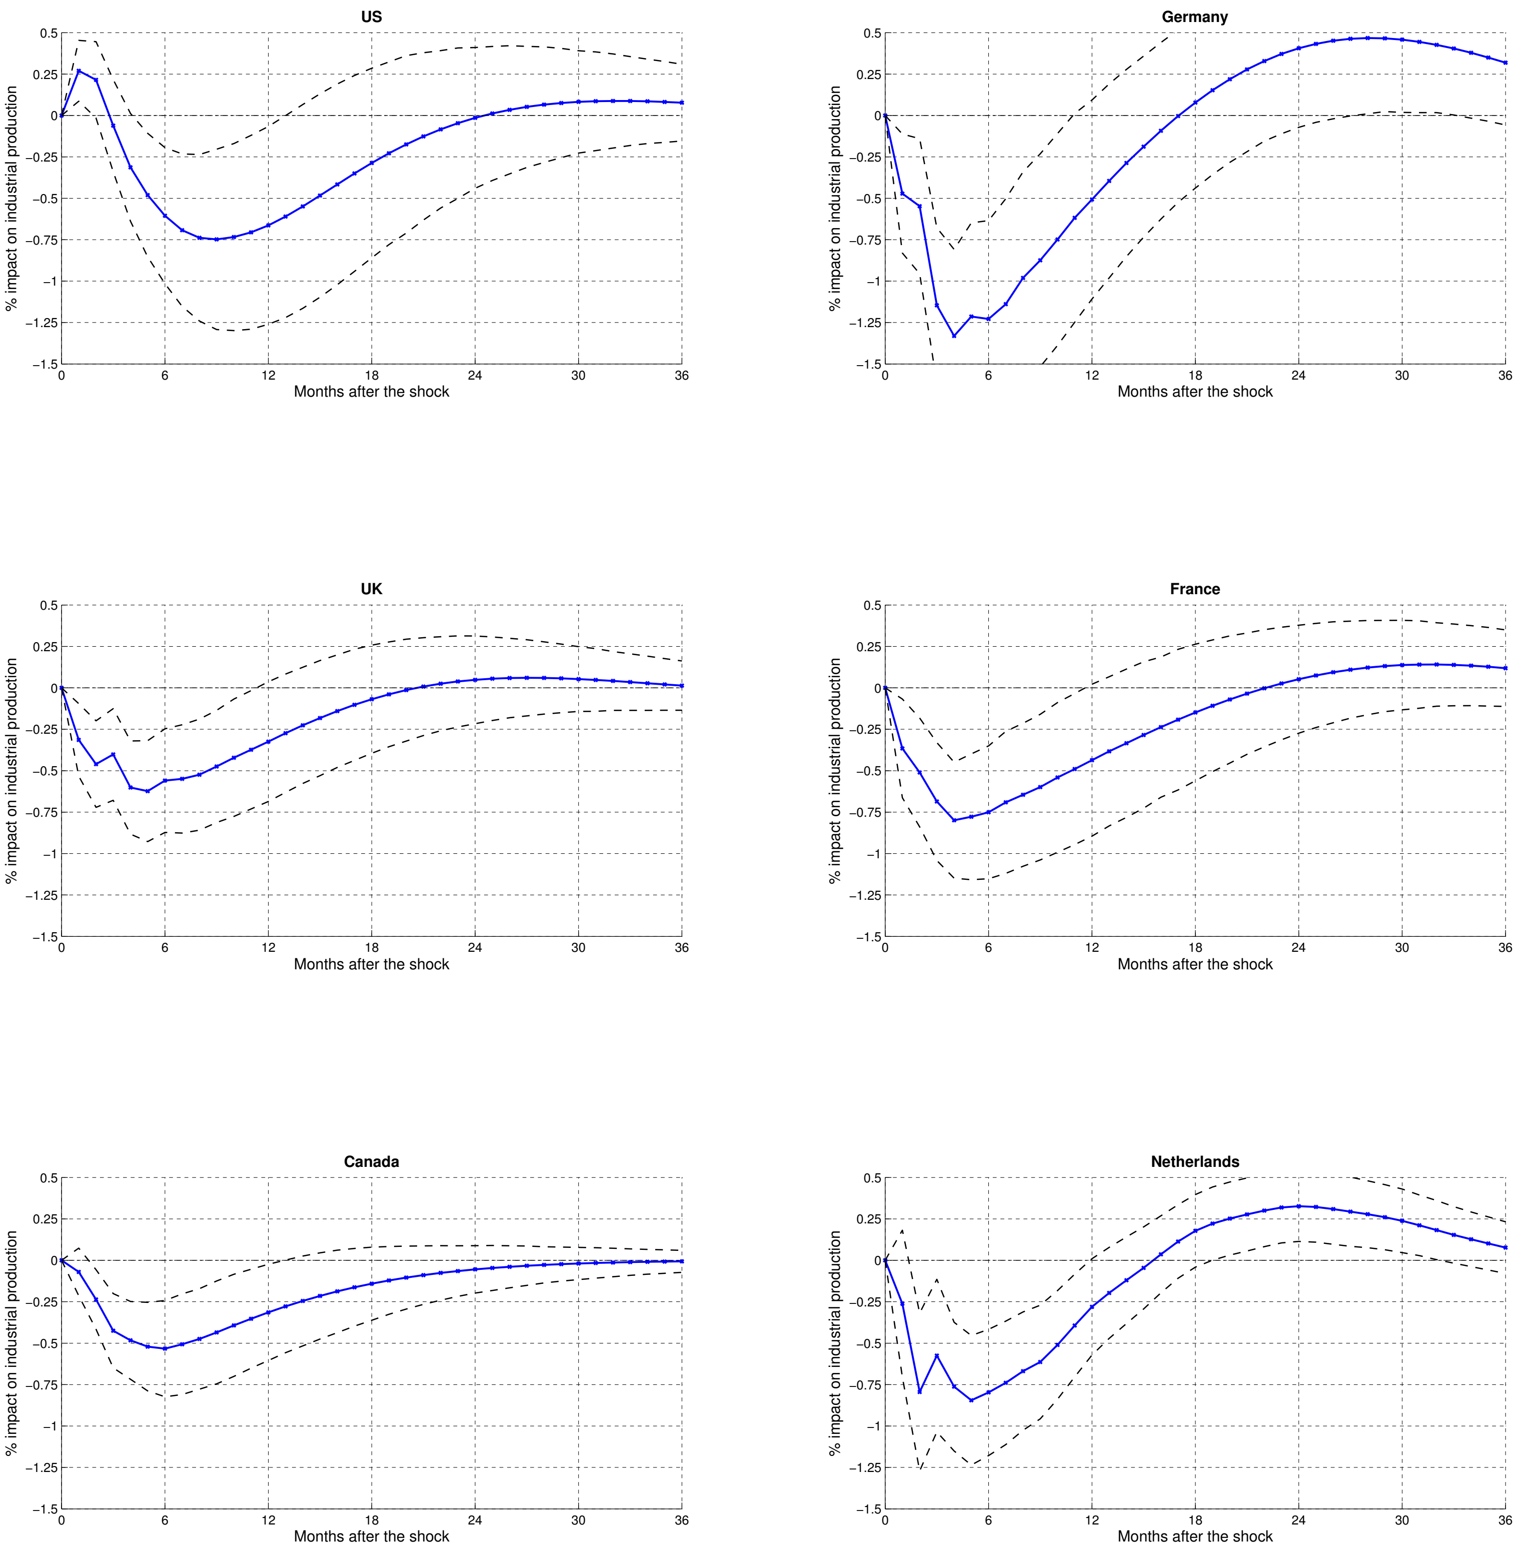

Figure 14: Impulse response function of industrial production to an uncertainty shock that is scaled such that volatility increases by 10 points on impact, preferred identification scheme with full volatility time series, large countries, common sample; 90% bootstrapped confidence interval

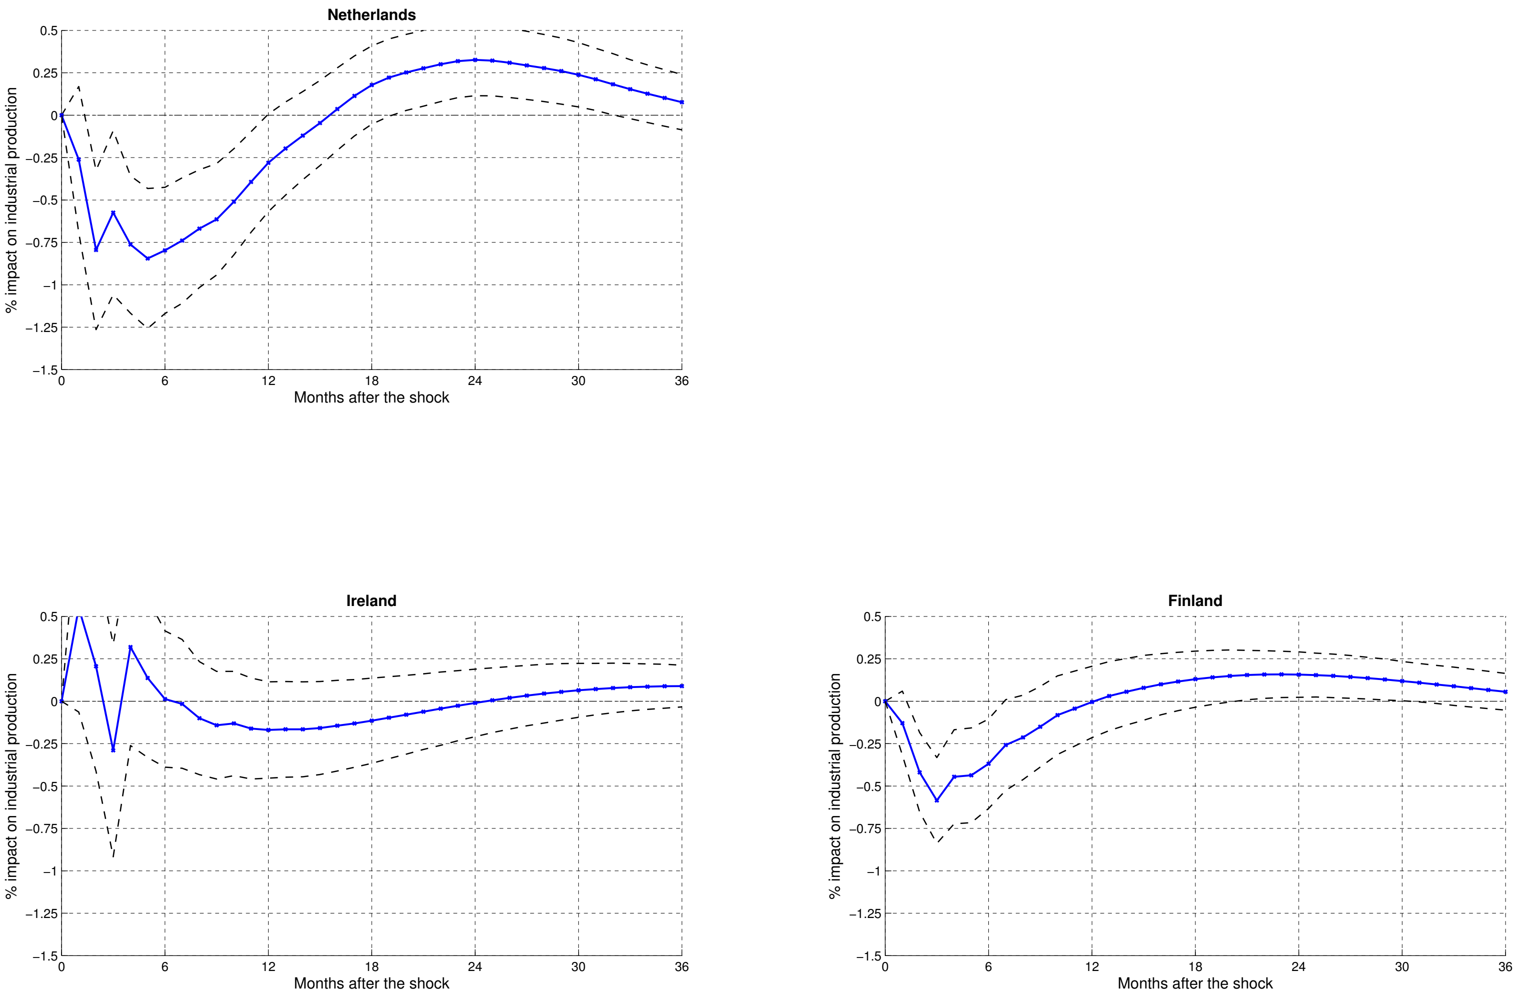

Figure 15: Impulse response function of industrial production to an uncertainty shock that is scaled such that volatility increases by 10 points on impact, preferred identification scheme with full volatility time series, small countries, common sample; 90% bootstrapped confidence interval

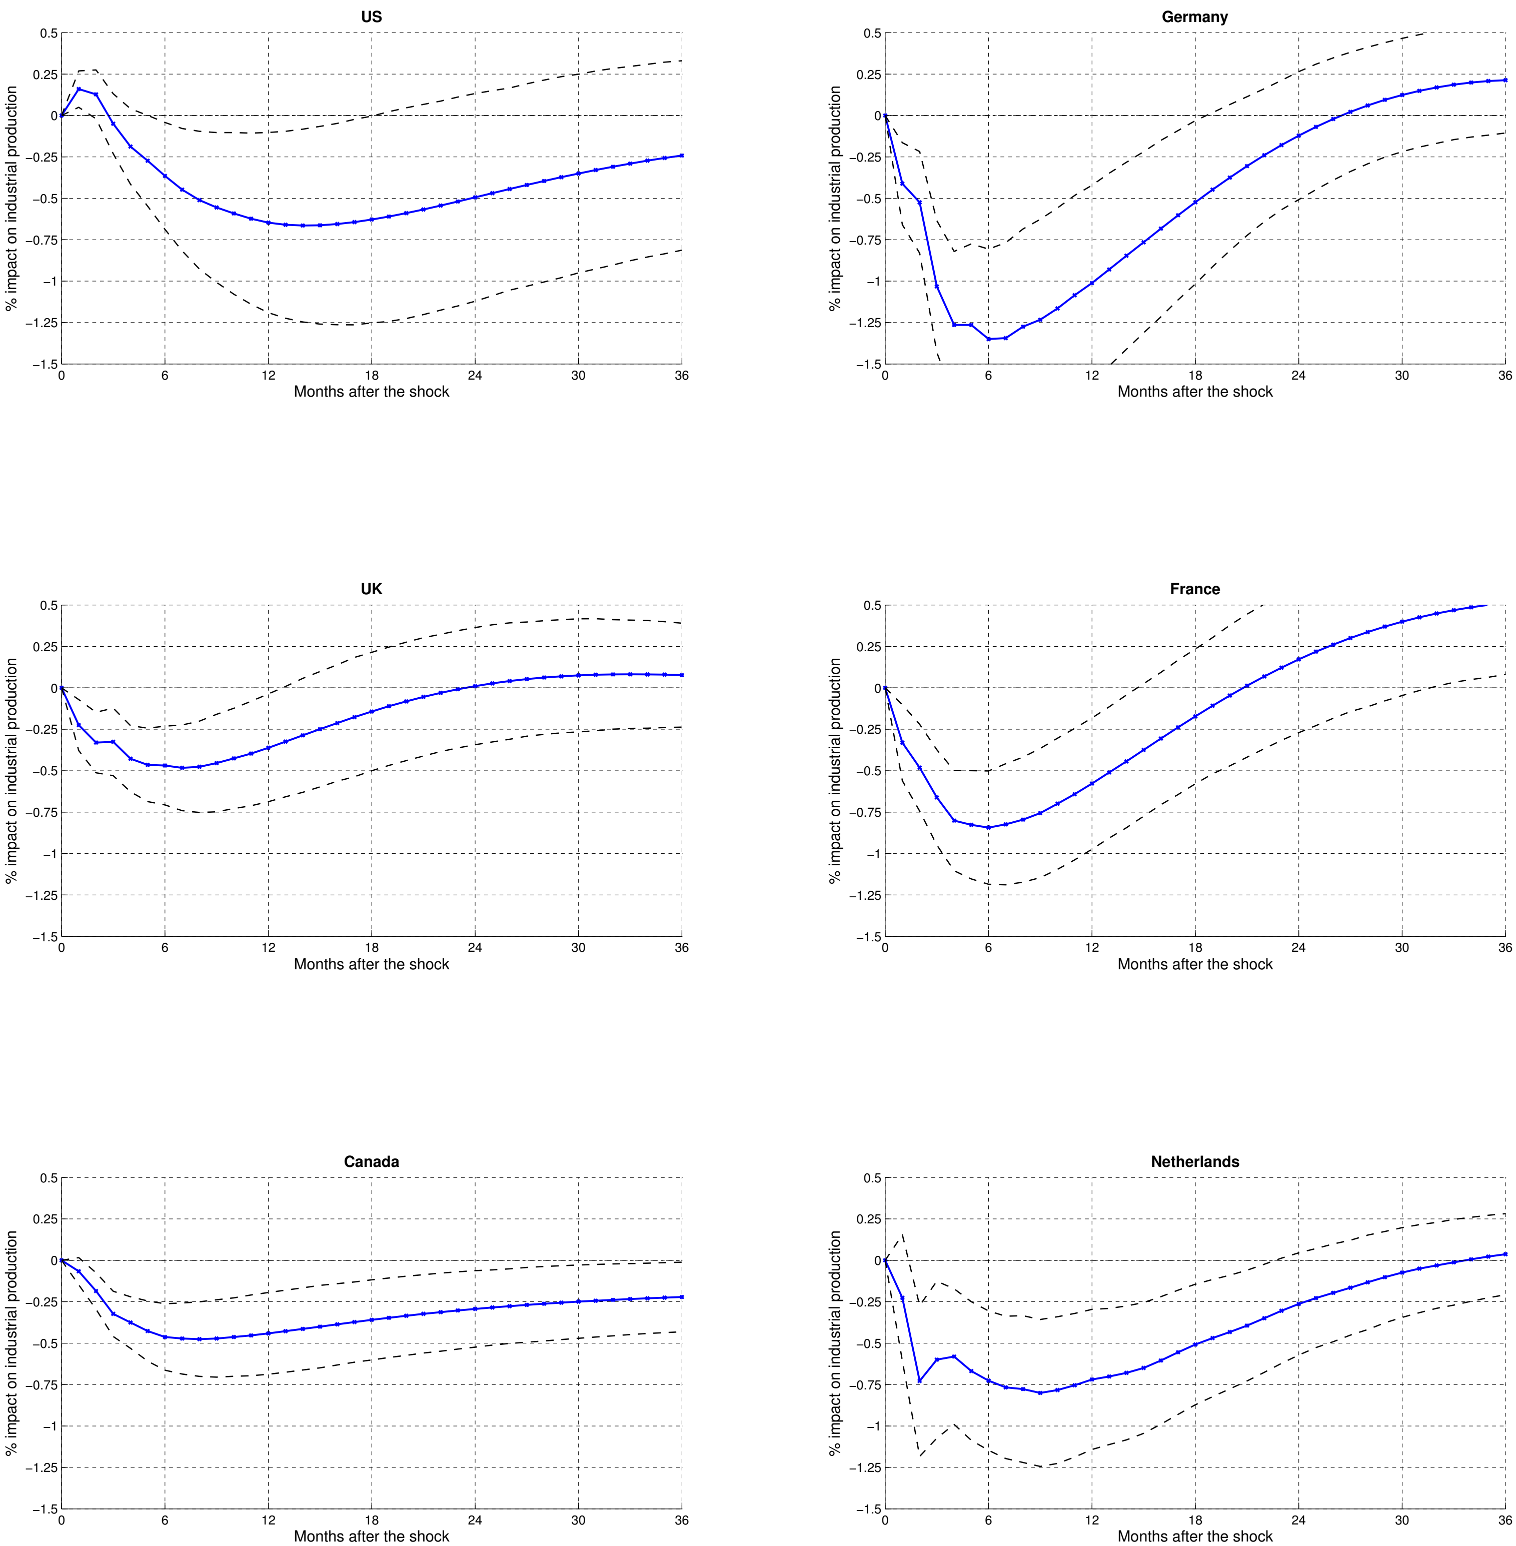

Figure 16: Impulse response function of industrial production to a two standard deviation uncertainty shock, preferred identification scheme with full volatility time series, large countries, common sample; robustness check without HP filter; 90% bootstrapped confidence interval

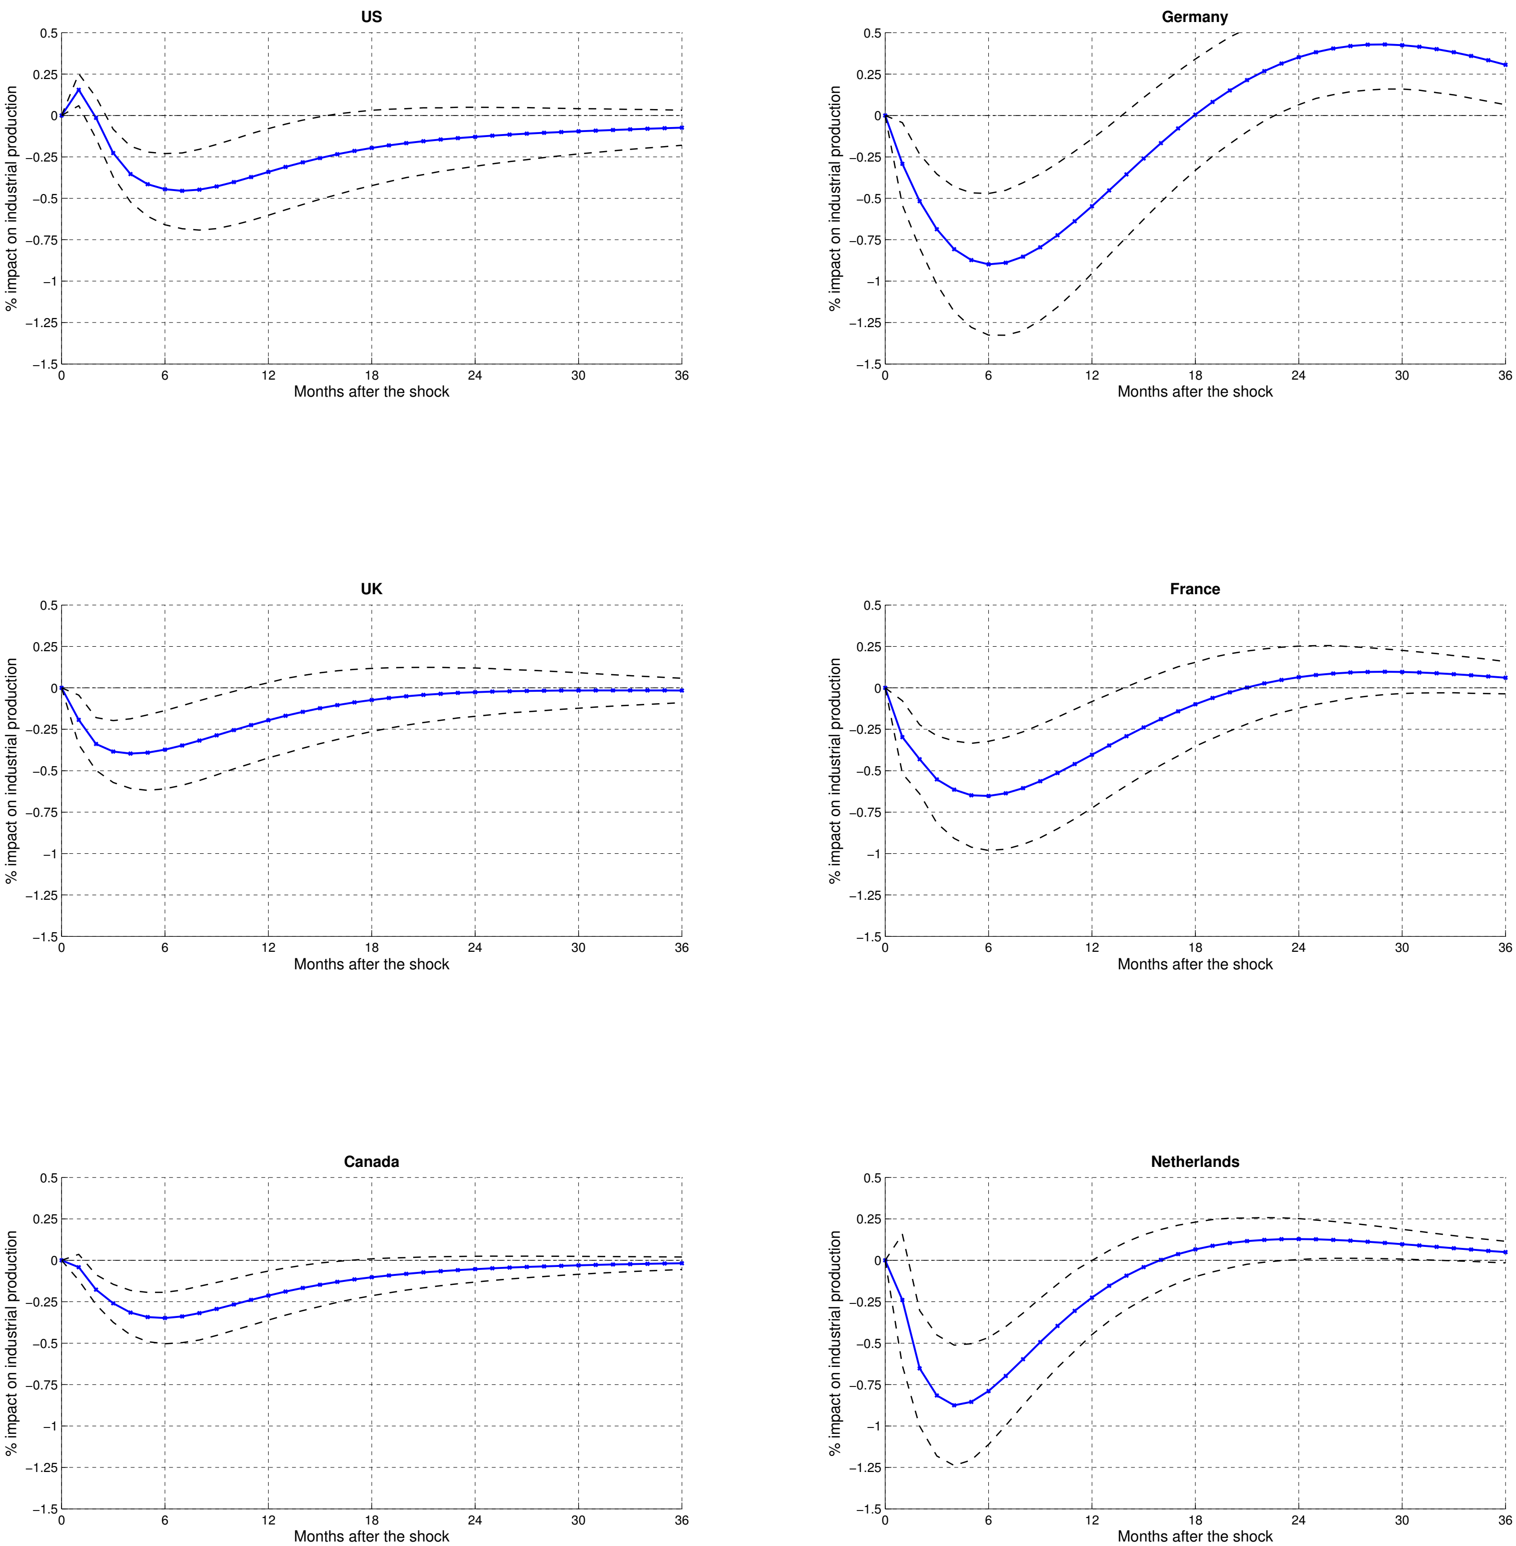

Figure 17: Impulse response function of industrial production to a two standard deviation uncertainty shock, preferred identification scheme with full volatility time series, large countries, common sample; robustness check with two lags; 90% bootstrapped confidence interval

References

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring Economic Policy Uncertainty," unpublished manuscript.

Bartelsman, E. J., P. A. Gautier, and J. de Wind (2016): "Employment Protection, Technology Choice, and Worker Allocation," *International Economic Review*, 57(3), pp. 787–826.

Beetsma, R., and M. Giuliodori (2012): "The changing macroeconomic response to stock market volatility shocks," *Journal of Macroeconomics*, 34(2), pp. 281–293.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), pp. 623–685.

——— (2014): "Fluctuations in Uncertainty," *Journal of Economic Perspectives*, 28(2), pp. 153–76.

Carrière-Swallow, Y., and L. F. Céspedes (2013): "The impact of uncertainty shocks in emerging economies," *Journal of International Economics*, 90(2), pp. 316–325.

Carroll, C. D. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, 112(1), pp. 1–55.

Denis, S., and P. Kannan (2013): "The Impact of Uncertainty Shocks on the UK Economy," IMF Working Papers No. 13/66, International Monetary Fund.

Dixit, A. K., and R. S. Pindyck (1994): *Investment under Uncertainty*. Princeton University Press, Princeton.

Fornari, F., and L. Stracca (2012): "What does a financial shock do? First international evidence," *Economic Policy*, 27(71), pp. 407–445.

Gourio, F., M. Siemer, and A. Verdelhan (2013): "International risk cycles," *Journal of International Economics*, 89(2), pp. 471–484.

Guo, H. (2002): "Stock Market Returns, Volatility, and Future Output," *Federal Reserve Bank of St. Louis Review*, 84, pp. 75–86.

Mumtaz, H., and K. Theodoridis (2012): "The impact of US real activity uncertainty on the UK economy," unpublished manuscript.

Popescu, A., and F. R. Smets (2010): "Uncertainty, Risk-taking, and the Business Cycle in Germany," *CESifo Economic Studies*, 56(4), pp. 596–626.

van Roye, B. (2011): "Financial stress and economic activity in Germany and the Euro Area," Kiel Working Papers No. 1743, Kiel Institute for the World Economy.

Veenendaal, P., K. Grabska, D. Lanser, M. Ligthart, and J. de Wind (2014): "Onzekerheidsvarianten bij verdere escalatie Oekraïne-conflict," CPB Background Document, CPB Netherlands Bureau for Economic Policy Analysis.

Publisher:

CPB Netherlands Bureau for Economic Policy Analysis
P.O. Box 80510 | 2508 GM The Hague
T (070) 3383 380

September 2016 | ISBN 978-90-5833-742-9