

The Costs of Economic Nationalism: Evidence from the Brexit Experiment*

Benjamin Born Gernot J. Müller Moritz Schularick
Petr Sedláček

May 4, 2019

Abstract

Economic nationalism is on the rise, but at what cost? We study this question using the unexpected outcome of the Brexit vote as a natural macroeconomic experiment. Employing synthetic control methods, we first show that the Brexit vote has caused a UK output loss of 1.7-2.5 percent by year-end 2018. An expectations-augmented VAR suggests that these costs are to a large extent driven by a downward revision of growth expectations in response to the vote. Linking quasi-experimental identification to structural time-series estimation allows us to not only quantify the aggregate costs but also to understand the channels through which expected economic disintegration impacts the macroeconomy.

Keywords: Economic nationalism, Brexit, Natural experiment, Synthetic control method, Anticipation effects, Economic policy uncertainty, Expectations-augmented VAR

JEL Codes: E65, F13, F42

*Born: Frankfurt School of Finance & Management and CEPR, b.born@fs.de, Müller: University of Tübingen and CEPR, gernot.mueller@uni-tuebingen.de, Schularick: University of Bonn and CEPR, schularick@uni-bonn.de, Sedláček: University of Oxford and CEPR, petr.sedlacek@economics.ox.ac.uk. We thank three anonymous referees and Morten Ravn for very useful comments. We also thank seminar audiences at the Bundesbank, the Dallas Fed, ifo Munich, Queen Mary University London, the University of Hamburg, and the LSE CEP Conference on the “The Economic Consequences of Brexit” for helpful questions and discussions. Special thanks go to Chris Redl for sharing data for the UK macroeconomic uncertainty index, and Rubén Domínguez Díaz and Emily Swift for excellent research and editorial assistance, respectively. The usual disclaimer applies.

It is the maxim of every prudent master of a family, never to make at home what will cost him more to make than to buy (...) What is prudence in the conduct of every private family, can scarce be folly in that of a great kingdom.

(Adam Smith, 1776)

The specter of economic nationalism is haunting the global economy. Supporters of the rule-bound liberal world economic order that was constructed after World War II are on the defensive. For economists, the recent rise of protectionism represents a particular challenge. From its beginnings, the benefits of an international division of labor have been a central tenet of the discipline. Foreshadowing a large literature, Adam Smith diagnosed, in disarmingly simple words, that foregoing the gains from trade would harm the wealth of nations.

It seems therefore plausible that the recent rise of economic nationalism could take a toll on future economic prosperity. And to the extent that market participants act in a forward-looking manner, expectations of economic disintegration and de-globalization could already affect investment and consumption today. In addition, as trade agreements are torn apart, old alliances nullified, and protectionist measures contemplated, policy uncertainty has increased substantially. Increased uncertainty, too, may impact the global economy adversely.

Can we measure the costs of economic nationalism? In this paper we make an attempt to do so as we exploit a unique natural experiment: the decision of the UK to leave the European Union. Two aspects are key for interpreting the vote for Brexit as a natural experiment. First, the outcome of the referendum on June 23, 2016 came as a major surprise. “Remain” was ahead in the voter polls for most of the time and betting markets indicated that it would win by a considerable margin. Second, the voting behavior was largely unrelated to UK’s recent macroeconomic performance. Rather, according to many observers, the case for Brexit was predominantly based on the political imperative to “take back control.”

The Brexit experiment allows us to measure the costs of economic nationalism because the (eventual) UK departure from the Single European Market would entail significant economic disintegration. The disintegration shock would extend beyond trade in goods and services. The British labor market may become less open to foreign workers, and capital markets would likely be affected through disintegration from the common European market for financial services. However, while the direction of the change is clear, the exact extent of disintegration remains uncertain, not least because the details of Brexit are still negotiated. Hence, the Brexit experiment nests both an expected disintegration shock and a policy uncertainty shock: it is a showcase of economic nationalism.

In addition to measuring the output costs of the Brexit vote, our paper makes two methodological contributions. First, it breaks new ground by combining two different approaches

in empirical macroeconomics: the synthetic control method and an expectations augmented vector autoregression (EVAR). In particular, we use the synthetic control method that was recently added to the toolbox of empirical macroeconomics by Abadie and Gardeazabal (2003) and Abadie *et al.* (2010, 2015) to identify, under fairly mild assumptions, the causal effect of the Brexit vote on UK’s macroeconomic performance since the referendum. But while the synthetic control approach exposes causal effects at the aggregate level, the underlying channels operate in the dark. We therefore map the results of the synthetic control method into a structural EVAR framework. This allows us to quantify the contribution of different channels to the overall impact of the Brexit vote estimated on the basis of the synthetic control method. It is the combination of both approaches that allows us to both identify the overall costs of the Brexit vote to the British economy, and to understand the channels through which these come about. Our second methodological contribution is to apply the “end-of-sample” test proposed by Andrews (2003) and discussed with respect to the synthetic control framework in Hahn and Shi (2017) and Ferman and Pinto (forthcoming) in order to establish the significance of the estimated causal effect of the Brexit vote. This is an important step forward in the synthetic control literature which has up until now relied almost exclusively on placebo tests to evaluate the credibility of the results.

More specifically, the synthetic control method makes it possible to measure the causal impact of the Brexit vote on the UK economy by estimating its synthetic “doppelgänger”. It does so by letting an algorithm determine which combination of “donor” economies matches the growth trend of the UK economy before the Brexit vote with the highest possible accuracy. The set of weights assigned to the donor economies is entirely data-driven. The better the algorithm constructs a doppelgänger for the UK economy as a weighted combination of other economies before the referendum, the more precise our results will be. In order to ensure that countries are sufficiently homogenous to begin with, we limit our analysis to OECD countries. We then rely on all available data to obtain the best match possible.

Comparing the evolution of this synthetic doppelgänger to actual data for the UK economy directly quantifies the aggregate costs of the Brexit referendum. Identification hinges on the very notion that the Brexit vote is a natural experiment: because the vote was unanticipated and unrelated to macroeconomic performance, the doppelgänger continues to evolve in the way the UK economy would have in absence of the referendum. The difference in output between the UK economy and its doppelgänger after the referendum is the causal effect of the experiment. Importantly, our approach does not depend on having the right economic model for the British, the European, or the global economy, nor do we need to assume a particular Brexit deal emerging from future negotiations.

We find that the economic costs of the Brexit vote are already visible and quite large:

there is a sizable output gap between the doppelganger and actual output in the UK. By the end of 2018, the “doppelganger gap” amounts to 2.4 percent in our baseline, and the cumulative loss of GDP is 55 billion pounds. Following Abadie *et al.* (2015), we also conduct a number of time- and country-placebo tests, reassuring us of the causal effect of the Brexit vote. In addition, we run a battery of robustness tests and find that the costs of the Brexit vote may lie in a range between 1.7 and 2.5 percent of GDP.

However, while the synthetic control method points to large causal effects of the Brexit vote on the UK economy, the underlying channels remain a black box. In order to open this black box we turn to a structural time-series framework. The starting point is the fact that the estimated aggregate costs have materialized before Brexit itself has actually taken place. Therefore, the impact of the Brexit referendum on UK’s macroeconomy must necessarily be caused by changes of expectations in response to the Brexit vote.

Yet, expectations may have changed in two distinct ways. On the one hand, households and firms may have revised downwards their expectations of future prosperity, because they expect economic disintegration to take its toll on “the wealth of the nation”. Such a downward revision induces an immediate reduction of consumption and investment spending (e.g., Blanchard *et al.*, 2013). On the other hand, market participants may also have become more uncertain about future income, not least because the details of Brexit are still unclear. Such uncertainty effects can also be detrimental to economic activity (e.g. Bloom, 2009; Born and Pfeifer, 2014; Fernández-Villaverde *et al.*, 2015; Baker *et al.*, 2016).

To dissect the doppelganger gap and unscramble anticipation and uncertainty effects we estimate an EVAR. It features quarterly data on output, interest rates, inflation, the exchange rate, but also a measure of economic policy uncertainty, and importantly, forecast revisions (“news”) regarding future output growth for various forecasting horizons. This approach, pioneered in the context of fiscal policy by Ramey (2011), Mertens and Ravn (2012) and others, allows us to directly capture the change in expectations due to the Brexit vote. Specifically, we use a unique data set which comprises output growth forecasts for the UK up until the year 2050. These forecasts have been substantially downgraded in response to the Brexit vote. In addition, we use the Economic Policy Uncertainty (EPU) index compiled by Baker *et al.* (2016). And again, this index reached an all-time high in the aftermath of the referendum.

This expectations-augmented VAR model serves two purposes. First, we use it to directly capture the effect of news on macroeconomic performance which a conventional VAR is ill-equipped to recover because of its backward-looking structure. Moreover, the EVAR allows us to purge the growth news of potential uncertainty effects: under our baseline identification scheme we permit uncertainty shocks to impact growth news contemporaneously, but not

vice versa, because forecasters are likely to downgrade their outlook if uncertainty is high and likely to hurt growth.

The second role of the EVAR is to identify uncertainty and growth-news shocks *caused* by the Brexit vote. We then use the estimated EVAR model to quantify the impact of these Brexit-related shocks on the time-path of real GDP. Specifically, we continue to rely on the Brexit vote being a natural experiment, which singles out structural shocks occurring in 2016Q3, the period right after the Brexit vote, as those caused by the referendum. We are then able to construct a counterfactual time-path for real GDP by “switching off” these Brexit-related uncertainty and growth-news shocks in the estimated EVAR.

It turns out that this EVAR-based counterfactual tracks the output path of the *doppelgänger* very closely. Because it is based on an altogether different approach and data set, the VAR analysis provides a valuable cross-check of the results obtained under the synthetic control technique. More importantly still, it allows us to separate anticipation and uncertainty effects. Overall, we find that the role of heightened uncertainty is fairly limited and downgrades of future output growth expectations account for the bulk of the estimated costs of the Brexit vote.

Our paper relates to work on the impact of (trade policy) uncertainty on international trade (see e.g. Novy and Taylor, 2014; Handley and Limão, 2017, 2015; Limão and Maggi, 2015). We also share a focus of analysis with studies of macroeconomic experiments at the aggregate level (Alesina and Fuchs-Schündeln, 2007; Fuchs-Schündeln and Hassan, 2016). Billmeier and Nannicini (2013), in particular, also use the synthetic control approach to study the impact of economic liberalizations. Finally, our paper complements a number of influential studies on the instantaneous macroeconomic impact of anticipated future (policy) changes or, more generally, “news” (see e.g. Barsky and Sims, 2011, 2012; Beaudry and Portier, 2006; Schmitt-Grohé and Uribe, 2012; Mertens and Ravn, 2011, 2012).

In a closely related—and yet quite distinct—study Campos *et al.* (forthcoming) also use the synthetic control method to estimate the growth effect of *joining* the EU. They find a positive and sizable effect of EU accession also for the UK, consistent with our results. Also, we stress that in this paper we focus on the consequences of the Brexit vote, rather than on actual Brexit. Saia (2017) uses the synthetic control approach to measure the costs of the UK of staying out of the euro. Had the UK joined the euro, trade flows would have been 16% higher, he finds.

A systematic analysis of the immediate implications of the Brexit vote has just begun.¹ An exception is Ramiah *et al.* (2016) who show that the response of cumulative abnormal returns

¹Instead, a number of authors have investigated actual Brexit scenarios on the basis of model simulations, see, for instance, Dhingra *et al.* (2017) and the studies surveyed by Sampson (2017).

in different sectors after the referendum is mostly negative. Davies and Studnicka (2018) also study the response of stock returns to the Brexit vote and find considerable heterogeneity. Breinlich *et al.* (2017) argue that the inflation increase following the post-referendum pound depreciation amounts to about a 400 pound consumption loss for the average British household. Finally, Berg *et al.* (2017) use a matching strategy to show that bank lending dropped by 20 percent in the syndicated loan market after the Brexit vote.

The remainder of this paper is organized as follows. In the following section, we provide more details to support the argument that the Brexit vote can be understood as a natural experiment. Section 2 then describes how we apply the synthetic control method to measure the output effect of the Brexit vote. Section 3 zooms in on the transmission mechanism and quantifies the roles of economic uncertainty and shifts in expectations. A final section concludes.

1 The Brexit vote as a natural experiment

The Brexit vote offers a rare opportunity to measure the costs of economic nationalism. As argued above, economic nationalism reduces international economic integration and raises policy uncertainty. In general—because of confounding factors—it is challenging to quantify the impact of these developments on economic activity. One strategy is to employ fully structural equilibrium models. For instance, following the seminal contributions of Eaton and Kortum (2002) and Melitz (2003), studies have attempted to measure how impediments to trade impact aggregate income. Similarly, Fernández-Villaverde *et al.* (2015) and Born and Pfeifer (2014) employ dynamic general equilibrium models in order to determine the extent to which policy uncertainty causes economic contractions. Overall, these studies have delivered valuable insights, but the results depend on restrictive assumptions and hence remain controversial.

A second strategy is to pursue a more data-driven approach. As far as economic integration is concerned, a long-standing literature has investigated the correlation between trade openness and growth. Here, the evidence often points towards a positive correlation between openness and growth, but while informative, identifying a causal effect remains a major challenge because trade policies are generally not determined randomly (Goldberg and Pavcnik, 2016).² Similarly, there is evidence that economic policy uncertainty causes output to contract in the short run, but identification remains challenging (Baker *et al.*, 2016).

²The historical record is mixed too as some of the greatest success stories in economic history, the rise of the US and German economies in the 19th, and Japan in the 20th century, partly occurred behind high tariff walls; see Schularick and Solomou (2011).




Figure 1: Left: odds for referendum outcome implied by online bets placed on Betfair exchange (source: BETdata). Right: Google search for “Brexit Leave” (source: Google trends).

Natural experiments, in contrast, “are situations in which we can argue that the change in policy is large relative to potential confounding factors that cannot be controlled for” (Nakamura and Steinsson, 2018). This holds true for the Brexit vote. However, also in this case the underlying identification assumptions have to be made explicit. Fuchs-Schündeln and Hassan (2016, p. 925) define “natural experiments as historical episodes that provide observable, quasi-random variation in treatment subject to a plausible identifying assumption.”³

That the UK has been subjected to the Brexit vote is indeed random as far as the macroeconomy is concerned, because macroeconomic developments were largely irrelevant for a) the decision to hold a referendum and b) its outcome. According to most observers political factors were the key in both instances. In 2013, then Prime Minister David Cameron promised to hold a referendum as a concession to the euro-sceptic wing of his party. This scepticism—which eventually prevailed in the referendum—is largely fueled by political considerations, rather than by concerns about economic growth or the business cycle. A key aspect was the idea to “take back control”, in turn due to concerns about political sovereignty, notably with regards to immigration and the rulings of the European Court of Justice (Sampson, 2017).

This is not to say that socio-economic characteristics are unrelated to individual voting behavior. For instance, voting behavior varied systematically in terms of educational attainment, demography, and regional industry structure (e.g., Alabrese *et al.*, 2019; Becker *et al.*, 2017). It is unlikely, however, that these factors impact economic performance systematically at the macroeconomic level. And what matters for our analysis is that the decision to hold

³“The “natural” in natural experiments indicates a researcher did not consciously design the episode to be analyzed, but can nevertheless use it to learn about causal relationships” (Fuchs-Schündeln and Hassan, 2016). In this regard natural experiments differ from controlled experiments, “the holy grail of empirical science” (Nakamura and Steinsson, 2018).

the referendum as well as its outcome are unrelated to macroeconomic performance.⁴

Moreover, we are able to date the “treatment” precisely, because the outcome of the referendum was largely unexpected. This is illustrated by Figure 1. The left panel shows the odds for the referendum outcome implied by bets offered on the Betfair exchange.⁵ Throughout our sample period odds were clearly stacked against “Leave”. Similarly, for the longest time prior to the referendum most polls suggested a victory for “Remain”.⁶ The right panel of Figure 1 shows the frequency of Google search incidents for “Brexit Leave”. Clearly, interest in the issue arose only after the referendum suggesting once more that the outcome of the Brexit vote took most people by surprise.

Finally, we note that the Brexit vote is a *unique* natural experiment because it involves changes at the *aggregate* level. Other experiments that are studied in macroeconomics do not directly allow to measure the macro impact of policies because treatment takes place at the household or individual level. For instance, an influential study of the US economic stimulus payments in 2008 by Parker *et al.* (2013), exploits the randomized timing of disbursements of payments to households. As a result it is possible to measure the effect of transfers on household consumption. This effect, however, is not directly informative about the macroeconomic effects of variations in aggregate transfers. Instead, an additional, model-based analysis is required (Fuchs-Schündeln and Hassan, 2016).⁷ The Brexit experiment, on the other hand, exposes an entire country to a “treatment” such that we are able to measure its macroeconomic effect directly.

2 The output effect of the Brexit vote

In order to evaluate the causal impact of the Brexit vote on the UK macroeconomy we need to define an appropriate comparison economy, a counterfactual benchmark. Since our focus is on the dynamic effects of the Brexit vote on UK output, we require the comparison economy to track the actual UK economy as closely as possible prior to the referendum. At the same time, it must be left unaffected by the Brexit vote.

We follow Abadie and Gardeazabal (2003); Abadie *et al.* (2010, 2015) and use synthetic

⁴Reassuringly, Becker *et al.* (2017) find that immigrant share at the local authority level does not predict vote shares for “Leave”. This suggests that the result of the vote is unrelated to the increasing foreign labour supply that is a macroeconomic trend that has been somewhat specific to the UK. Fetzer (2018), in turn, argues that the outcome of the referendum is closely associated with fiscal austerity. This appears plausible. However, many countries in our donor pool have also been subjected to austerity. Hence, with regard to austerity the UK has not been experiencing an idiosyncratic macroeconomic development.

⁵Clearly, these odds need not reflect actual public opinion at the time.

⁶An exception was a brief period in early June when “Leave” was ahead in the poll of polls, see <https://whatukthinks.org/eu/opinion-polls/poll-of-polls/>.

⁷See also the approach and the discussion in Nakamura and Steinsson (2014).

control methods to construct precisely such a doppelganger to the UK economy. Our identifying assumption is that the UK economy would have developed as the doppelganger, had it not been for the Brexit vote. This assumption is plausible to the extent that, given economic fundamentals, the UK economy and its doppelganger were equally likely to obtain the “treatment” of the Brexit vote.

We can then directly quantify the costs of the Brexit vote as the “doppelganger gap”: the difference between UK’s actual output performance and that of the doppelganger economy. Lastly, we run a number of tests showing that our estimated effects indeed reflect a causal impact of the referendum shock.

2.1 Constructing the doppelganger

We construct the doppelganger as a synthetic control unit from a “donor pool”. In order to specify the donor pool we proceed as follows. First, we focus on OECD countries in order to ensure that countries are sufficiently homogenous to begin with. Second, we keep all OECD countries in the donor pool for which data on all relevant variables are available. For the baseline we do not restrict the donor pool further. Given this unrestricted pool, the construction of the doppelganger follows a strictly data-driven approach. However, below we also conduct an extensive robustness analysis in order to explore to which extent our results depend on individual countries being included in the donor pool.

Our approach leaves us with 23 countries and quarterly observations for the period from 1995Q1–2016Q2.⁸ Our procedure thus assumes that a possible treatment effect materializes after 2016Q2. Moreover, we assume that the countries in the donor pool are not affected by the treatment. We relax both assumptions in our analysis below.

The doppelganger is a weighted average of the countries in the donor pool. The weights are determined by minimizing the distance between real GDP of the UK and of the doppelganger prior to the treatment.⁹ Following Abadie *et al.* (2010); Abadie and Gardeazabal (2003), we also match the pre-Brexit-vote averages of a number of country characteristics.¹⁰ In our application, they are the GDP shares of consumption, investment, exports, and imports, plus labor productivity growth and the employment share in the population. Formally, we let \mathbf{x}_1 denote the (92×1) vector of 86 observations for real GDP and 6 covariate averages in the UK and let \mathbf{X}_0 denote a (92×23) matrix with observations in the countries included in the

⁸The reduction in donor countries compared to earlier versions of this paper is due to the inclusion of a number of covariates (see below).

⁹Specifically, we normalize real GDP to unity in 1995 in each country. See the online appendix for further details on the dataset.

¹⁰Averages of covariates are taken over the entire sample period 1995Q1 to 2016Q2. The online appendix shows that the results are robust to averaging over a period just before the Brexit-vote.

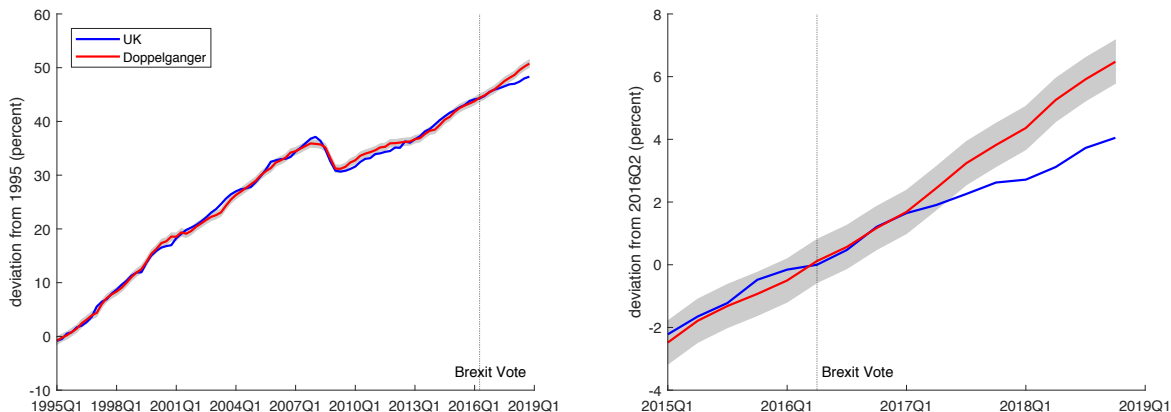


Figure 2: Real GDP of the UK. Actual data (blue line) vs doppelganger (red line). *Note:* shaded area is one standard deviation of difference prior to Brexit vote. Data source: OECD Economic Outlook.

donor pool. Finally, we let \mathbf{w} denote a (23×1) vector of weights w_j , $j = 2, \dots, 24$. Then, the doppelganger is defined by \mathbf{w}^* which minimizes the following mean squared error:

$$(\mathbf{x}_1 - \mathbf{X}_0 \mathbf{w})' \mathbf{V} (\mathbf{x}_1 - \mathbf{X}_0 \mathbf{w}), \quad (1)$$

subject to $w_j \geq 0$ for $j = 2, \dots, 24$ and $\sum_{j=2}^{24} w_j = 1$. In this expression, \mathbf{V} is a (23×23) symmetric and positive semidefinite matrix.¹¹

Turning to the results, the left panel of Figure 2 displays the time series for real GDP in the UK (blue line) and in the doppelganger economy (red line). The shaded area represents one standard deviation of the pre-treatment difference between the UK and its doppelganger. Note that the match is imperfect as our procedure determines 23 parameters (country weights) in order to match more than 90 observations. This being said, prior to the referendum both series display a very high degree of co-movement—both at low and high frequencies. Table 1 shows that the pre-Brexit vote averages of the additional covariates are also matched well. We are thus confident that the doppelganger provides a meaningful counterfactual which allows us to quantify the effect of the referendum shock on economic activity in the UK.¹²

Table 2 displays the country weights (rounded to the second digit) which define the

¹¹ \mathbf{V} is a weighting matrix assigning different relevance to the characteristics in \mathbf{x}_1 and \mathbf{X}_0 . Although the matching approach is valid for any choice of \mathbf{V} , it affects the weighted mean squared error of the estimator (see the discussion in Abadie *et al.* (2010), p. 496). Following Abadie and Gardeazabal (2003) and Abadie *et al.* (2010), we choose a diagonal \mathbf{V} matrix such that the mean squared prediction error of the outcome variable (and the covariates) is minimized for the pre-Brexit vote period. Including the covariates in the optimization differs from Kaul *et al.* (2018) who have raised concerns about including all pre-intervention outcomes together with covariates when using the SCM.

¹²In addition, Section 3.1 shows that the non-targeted time paths of other economic aggregates in our doppelganger economy display a similar behavior as their UK counterparts. This is reassuring as it suggests that the synthetic control economy indeed provides a good match to the UK.

Table 1: Matching of covariates

| | UK | Doppelganger |
|---------------------------|-------|--------------|
| Consumption / GDP | 65.53 | 62.10 |
| Investment / GDP | 16.79 | 20.73 |
| Exports / GDP | 25.44 | 24.43 |
| Imports / GDP | 25.63 | 25.61 |
| Labor productivity growth | 0.28 | 0.29 |
| Employment share | 63.42 | 60.18 |

Note: All numbers are in percent. Labor productivity growth is the log difference between quarterly real GDP and quarterly total employment; employment share is the ratio between total employment and the working age population.

Table 2: Composition of the doppelganger: country weights

| | | | | | | | |
|-----------|-------|-------------|-------|-----------------|-------|-------------|-------|
| Australia | <0.01 | Austria | <0.01 | Belgium | <0.01 | Canada | <0.01 |
| Finland | <0.01 | France | <0.01 | Germany | 0.05 | Hungary | 0.11 |
| Iceland | 0.01 | Ireland | 0.01 | Italy | 0.17 | Japan | <0.01 |
| Korea | <0.01 | Luxembourg | <0.01 | Netherlands | <0.01 | New Zealand | 0.14 |
| Norway | <0.01 | Portugal | <0.01 | Slovak Republic | <0.01 | Spain | <0.01 |
| Sweden | <0.01 | Switzerland | <0.01 | United States | 0.51 | | |

doppelganger economy. The United States and New Zealand, but also Italy and Hungary are assigned the largest weights. Together, Germany, New Zealand and the US account for 70 percent of the doppelganger dynamics. There are also smaller contributions from Iceland and Ireland. While these weights are plausible, given the position of the UK in the world economy and the fact that it operates within the EU, but outside the Euro area (like Hungary), in what follows we consider a battery of robustness checks related to the donor pool countries.

2.2 Measuring the immediate output effect of the Brexit vote

We are now in a position to quantify the output effect of the referendum shock. In order to do this we contrast the output performance in the UK and in the doppelganger economy in the quarters following the referendum shock. For this purpose the right panel of Figure 2 zooms in on the post-referendum period. As before, the shaded area corresponds to one standard deviation of the pre-treatment difference between output of the UK and the doppelganger. We loosely interpret a post-treatment path of GDP that leaves the shaded area as evidence of a significant output effect of the referendum shock and will conduct more sophisticated inference below. To facilitate the quantitative assessment we express output deviations vis-à-vis the UK level in 2016Q2.

A number of observations stand out. While throughout the second half of 2016 there

is hardly any effect of the referendum shock, a significant effect begins to materialize since 2017Q1. In fact, the UK seems to embark on a different growth trajectory relative to the doppelganger. By the end of 2018, output in the UK falls short of the doppelganger level by about 2.4 percent of GDP. The cumulative loss in terms of 2016 GDP equals approximately 55 billion pounds.

2.3 Inference

The shaded areas in Figure 2 quantify the standard deviation of the doppelganger gap prior to the Brexit vote. In other words, they are a measure of fit prior to the Brexit vote. The right panel of Figure 2 then highlights that the doppelganger quickly deviates from the realized path of UK GDP that far exceeds these bounds, indicating that such a deviation is non-standard compared to the pre-Brexit vote period.

While such bounds are indicative of strong effects, they are not a formal test of significance. Recently, Hahn and Shi (2017) have suggested that the Andrews (2003) end-of-sample instability test may be used to conduct inference in the context of the synthetic control method. On an intuitive basis, the instability test quantifies whether the post-referendum doppelganger gap and all the pre-referendum doppelganger gaps of the same length can be considered to come from the same distribution.¹³

We follow Andrews (2003) and apply the end-of-sample instability test to our baseline estimation. The results show that the output effects of the Brexit vote are statistically significant (p-value of 0.05). Therefore, we conclude that despite the relatively short post-Brexit vote period, our estimated output effects of the Brexit vote are not only large, but also statistically significant.

2.4 Causality

Are these effects causal? To back the notion that the doppelganger gap is indeed caused by the referendum shock, this subsection provides a number of placebo experiments (Abadie *et al.*, 2010, 2015). The basic idea of the placebos is very intuitive. We can be confident that the synthetic control estimator captures the causal effect of an intervention as long as similar magnitudes are not estimated in cases where the intervention did not take place. In addition, we corroborate the results of the placebo tests with data on GDP forecasts just before the Brexit referendum. If indeed the Brexit vote caused the divergence of the doppelganger from the realized path of UK GDP, and if the referendum outcome was unexpected, then this should have not been forecasted prior to June 2016.

¹³More details on the test can be found in the online appendix.




Figure 3: Placebo tests. *Note:* left panel shows real GDP of UK (blue line) and baseline doppelganger (red line), with grey lines representing time placebo doppelganger estimates with fictitious Brexit vote dates ranging from 2010Q1 to 2016Q1. Right panel shows the UK doppelganger *gap* (thick black line), with grey lines representing country placebo doppelganger gaps estimated by considering fictitious Brexit votes in donor pool economies. For comparability, all doppelganger gaps are normalized by their respective pre-Brexit standard deviations and centered around their 2015 means.

2.4.1 Placebo tests

First, we run twelve time-placebo tests for which we shift the treatment date artificially backward in time: we consider treatment dates in all quarters from 2010Q1 to 2016Q1. In each instance, we construct a new doppelganger using exactly the same approach as in the benchmark specification. These doppelgangers are bound to differ from the baseline doppelganger, because the pre-treatment sample is shorter. Yet if there is indeed a causal effect of the actual treatment, then we should not observe a decline of UK output relative to these doppelgangers prior to the Brexit vote, that is before the actual treatment took place.

The left panel of Figure 3 shows the results together with the series for actual GDP (blue line) and our benchmark doppelganger (red line). Each grey line represents the path of a doppelganger obtained for one placebo treatment. Reassuringly, despite the fact that the time-placebo studies work with earlier “fictitious” Brexit-vote dates, the resulting synthetic controls are essentially parallel to our baseline doppelganger series. They only exhibit a divergence from the actual UK data at the “true” Brexit vote date.

In a second set of tests, we estimate synthetic controls for the donor pool countries, while exposing each of them to a placebo treatment at the end of 2016Q2. Once again, if our benchmark estimate for the UK is picking up the causal effect of the referendum shock, its effect should dominate any possible impact of the fictitious Brexit votes in the donor pool countries.

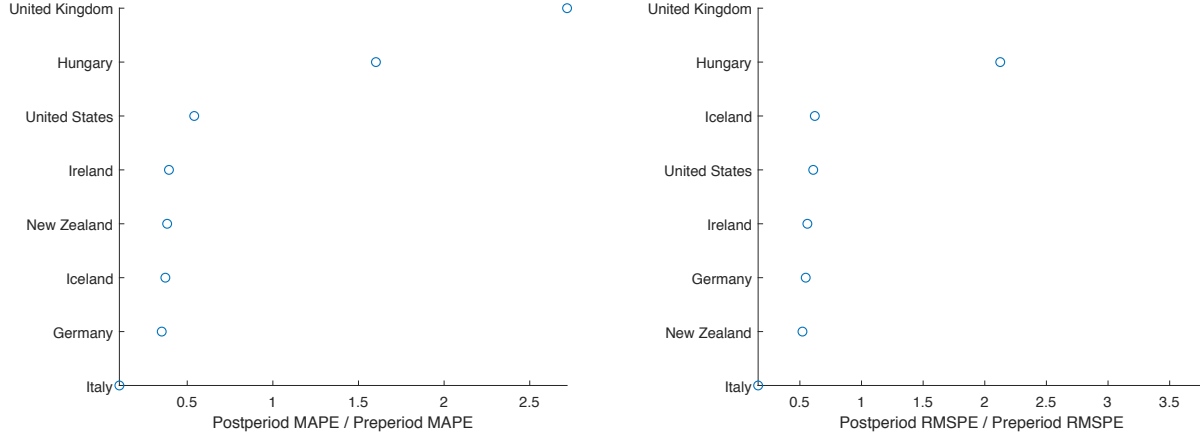


Figure 4: Relative measures of the pre- and post-treatment doppelganger gaps. *Note:* left panel shows the relative maximum absolute prediction error ρ_2 , the right panel shows the relative root mean squared prediction error ρ_1 .

The right panel of Figure 3 shows the UK doppelganger gap together with the doppelganger gaps of the seven countries which account for essentially all the weights in our baseline synthetic control estimate.¹⁴ For comparability, all doppelganger gaps are normalized by their respective pre-Brexit standard deviation and centered around their 2015 means. Relative to the country placebo estimates, the UK doppelganger gap stands out, both in terms of size and the systematic nature of the post-Brexit vote deviation.

An alternative way of quantifying the country placebo results is to compute statistics of relative pre- and post-treatment fit in the UK and the donor countries.¹⁵ Two such statistics are the relative root mean squared prediction error (RMSPE) and the maximum absolute prediction error (MAPE) defined as $\rho_1 = RMSPE_{post}/RMSPE_{pre}$ and $\rho_2 = MAPE_{post}/MAPE_{pre}$. Letting T denote the sample size and T_0 denote the period of treatment, i.e. the Brexit vote, the pre- and post-treatment measures of fit are defined as¹⁶

$$RMSPE_{pre} = \sqrt{\frac{1}{T_0 - 1} \sum_{t=1}^{T_0-1} (x_{1,t} - \mathbf{x}_{0,t}\mathbf{w})^2}, \quad (2)$$

$$MAPE_{pre} = \max |x_{1,t} - \mathbf{x}_{0,t}\mathbf{w}|, \quad t \in [1, T_0 - 1], \quad (3)$$

¹⁴The online appendix shows similar placebo results for all countries in the donor pool.

¹⁵Relative measures take into account heterogeneity in terms of pre-treatment fit of donor pool country synthetic controls.

¹⁶We normalize the post-treatment prediction error to zero at the treatment date to account for the possibility that the post-treatment time-path of the prediction error may be a continuation of previous trends rather than the result of the treatment.

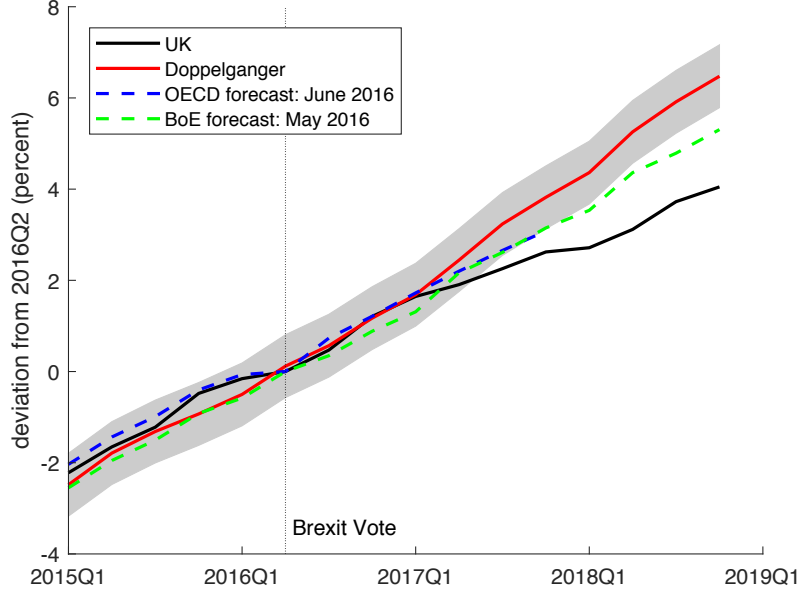


Figure 5: UK output: actual, doppelganger and two forecasts prior to the Brexit vote. *Note:* baseline doppelganger and actual real GDP together with real GDP predicted by the OECD (June 2016 Economic Outlook) and the Bank of England (May 2016 Inflation report). Both forecasts use the 2016 Economic Outlook data prior to 2016.

$$RMSP E_{post} = \sqrt{\frac{1}{T - T_0 - 1} \sum_{t=T_0}^T (x_{1,t} - \mathbf{x}_{0,t} \mathbf{w} - x_{1,T_0} + \mathbf{x}_{0,T_0} \mathbf{w})^2}, \quad (4)$$

$$MAPE_{post} = \max |x_{1,t} - \mathbf{x}_{0,t} \mathbf{w} - x_{1,T_0} + \mathbf{x}_{0,T_0} \mathbf{w}|, \quad t \in [T_0, T]. \quad (5)$$

Figure 4 depicts these two relative measures showing that the UK stands out with a particularly large post-treatment doppelganger gap.

2.4.2 The doppelganger and GDP forecasts prior to the Brexit vote

To corroborate the causal effect of the Brexit vote on the development of UK GDP, we can look at GDP forecasts just *before* the referendum. Given the unexpected nature of the Brexit vote outcome, and to the extent that this event had a causal effect on the subsequent evolution of GDP, one would expect that forecasts just prior to the referendum would not predict a slowdown in output growth but would rather be closer to our estimated doppelganger.

We verify this argument by using GDP forecasts from the June 2016 vintage of the OECD Economic Outlook and from the May 2016 Inflation Report of the Bank of England. Figure 5 then shows our baseline doppelganger, actual GDP and the GDP evolution based on the above two forecasts. Clearly, both forecasts are close to our estimated doppelganger providing further support of the causal nature of the Brexit vote on the development of UK GDP.

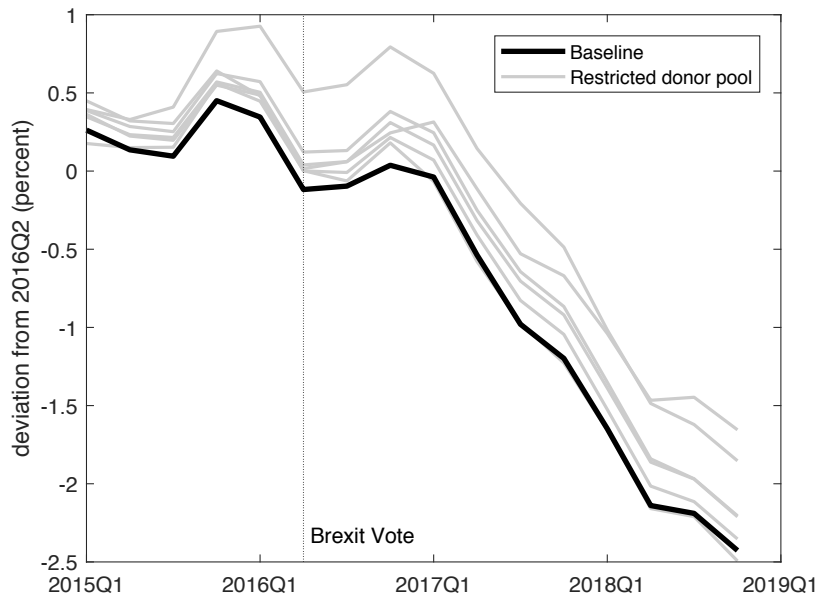


Figure 6: Baseline doppelganger gap and restricted donor pool doppelganger gaps. *Note:* baseline doppelganger gap with alternatives estimated by sequentially dropping each donor pool country which received a positive weight in the baseline estimates.

2.5 Effect of individual countries in the donor pool

Before moving on to understanding what drives the doppelganger gap, we assess the contribution of individual donor pool countries with non-zero weights to the doppelganger. Towards this end, we iteratively re-estimate our baseline model omitting in each iteration one of the countries that has a positive weight in the baseline estimation.

Figure 6 shows the baseline doppelganger gap together with the restricted donor pools, and Table 3 details the estimated weights in the restricted donor pool cases. While there is some variation, the overall conclusion remains unchanged: the Brexit vote caused a substantial output loss. Even in the estimation that shows the smallest effect, the output loss amounts to 1.7 percent of GDP at the end of 2018. This is the case when we omit Hungary.

However, it is important to recall that excluding countries from the donor pool also means that goodness of fit falls. For instance, when we exclude Hungary, the mean squared prediction error increases by 25 percent. This being said, the one-standard-deviation bands around the doppelganger gap in the baseline overlaps with the estimate that excludes Hungary. This implies that it is not possible to quantitatively distinguish the two estimates.

Table 3: Doppelganger weights: restricted donor pools

| | I | II | III | IV | V | VI | VII |
|-------------|--------|--------|--------|--------|--------|--------|--------|
| Australia | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | 0.12 |
| Austria | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 |
| Belgium | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 |
| Canada | < 0.01 | < 0.01 | < 0.01 | < 0.01 | 0.04 | < 0.01 | 0.09 |
| Finland | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 |
| France | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 |
| Germany | NA | < 0.01 | < 0.01 | 0.04 | 0.07 | 0.05 | < 0.01 |
| Hungary | 0.13 | NA | 0.11 | 0.13 | 0.13 | 0.15 | < 0.01 |
| Iceland | < 0.01 | 0.01 | NA | 0.01 | < 0.01 | 0.01 | 0.09 |
| Ireland | < 0.01 | 0.01 | 0.01 | NA | < 0.01 | < 0.01 | 0.04 |
| Italy | 0.17 | 0.25 | 0.19 | 0.16 | NA | 0.12 | 0.15 |
| Japan | 0.01 | < 0.01 | < 0.01 | < 0.01 | 0.09 | < 0.01 | 0.14 |
| Korea | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 |
| Luxembourg | < 0.01 | 0.04 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 |
| Netherlands | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 |
| New Zealand | 0.13 | 0.19 | 0.13 | 0.13 | 0.07 | NA | 0.23 |
| Norway | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 |
| Portugal | < 0.01 | < 0.01 | < 0.01 | < 0.01 | 0.05 | < 0.01 | 0.15 |
| Slovak Rep. | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 |
| Spain | < 0.01 | < 0.01 | < 0.01 | < 0.01 | 0.05 | < 0.01 | < 0.01 |
| Sweden | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 | < 0.01 |
| Switzerland | 0.02 | < 0.01 | 0.03 | < 0.01 | < 0.01 | 0.04 | < 0.01 |
| U.S. | 0.54 | 0.50 | 0.53 | 0.53 | 0.48 | 0.63 | NA |

Note: doppelganger weights in seven restricted donor pools. In each of the seven cases (I to VII) we omit one of the donor countries that received a positive weight in our baseline specification.

3 What drives the doppelganger gap?

By year-end 2018 the doppelganger gap amounts to 1.7-2.5 percent of GDP. This result emerges robustly from our synthetic control approach. We now seek to shed some light on the specific channels through which the Brexit vote has been impacting the UK economy. We proceed in two steps. First, we decompose the response of GDP into its components and contrast the evolution of these components in the UK to that of the doppelganger's GDP components. This simple accounting exercise shows that investment and, in particular, consumption have been particularly responsive to the Brexit vote.

Second, we note that the doppelganger gap emerged in response to the Brexit *vote*, before actual Brexit has taken place. Hence, the Brexit vote must have triggered a change in expectations which, in turn, had an effect on the economy prior to actual Brexit. However, expectations may change in two distinct ways. The Brexit vote may have changed the outlook for the UK economy (first moment) or may have simply increased policy uncertainty (second moment). We use an EVAR to explore this issue formally.

3.1 The components of GDP

We now decompose GDP into its components, both for the UK and the doppelganger. This exercise serves two purposes. First, it reassures us that the doppelganger mimics the behavior of the UK prior to the referendum, not only in terms of GDP, but also in terms of its components. This is important because the time path of GDP served as a target as we picked the weights that define the doppelganger. The time paths of the GDP components, however, have not been targeted. A good fit in this regard can therefore not be taken for granted. Second, the adjustment of the components of GDP in the UK relative to the doppelganger since the referendum provides some indication about the channels through which the Brexit vote has impacted the economy.

Specifically, we compute the components of GDP for the doppelganger for each of the real GDP components using our estimated baseline weights.¹⁷ Figure 7 shows these component for the UK and the doppelganger. Prior to the referendum, all components behave quite similarly in the UK and in the doppelganger economy, perhaps with the exception of real government consumption (and real exports after 2010). In addition, the bottom right panel shows the time path of the “components-based doppelganger” which is constructed by summing the individual component doppelgangers weighted by their respective average shares in GDP.¹⁸ This last panel shows that the discrepancies between the component doppelgangers do not cumulate to generate an unrealistic time path for real GDP.

Figure 7 shows that there is a widening gap between the UK and the doppelganger for all GDP components after the Brexit vote. This is particularly true for private consumption, investment and imports. While the contribution of consumption to the doppelganger gap starts almost immediately after the Brexit vote, the contribution of investment sets in more gradually. But the contributions of both variables gain on strength over time. Especially the slowdown in consumption throughout 2017 is an important driver of the doppelganger gap. This is in line with findings in Breinlich *et al.* (2017), who document that the large depreciation of the pound, in response to the referendum, induced consumer prices to rise mainly in 2017. On the other hand, the slowdown in imports relative to its doppelganger, reflecting the depreciation of the pound, contributed towards a reducing of the doppelganger gap.

Overall, our findings are consistent with the notion—central to modern macroeconomics—that economic agents respond in a forward-looking manner to an anticipated policy change.¹⁹

¹⁷In constructing the component doppelgangers, we rescale their levels such that their means prior to the Brexit vote match those of the data.

¹⁸Due to changing component shares over time, we adjust the level of the components-based doppelganger to match that of real GDP in the data, prior to the Brexit vote.

¹⁹For evidence on how the Brexit vote impacts firms’ financing decisions, see Berg *et al.* (2017).

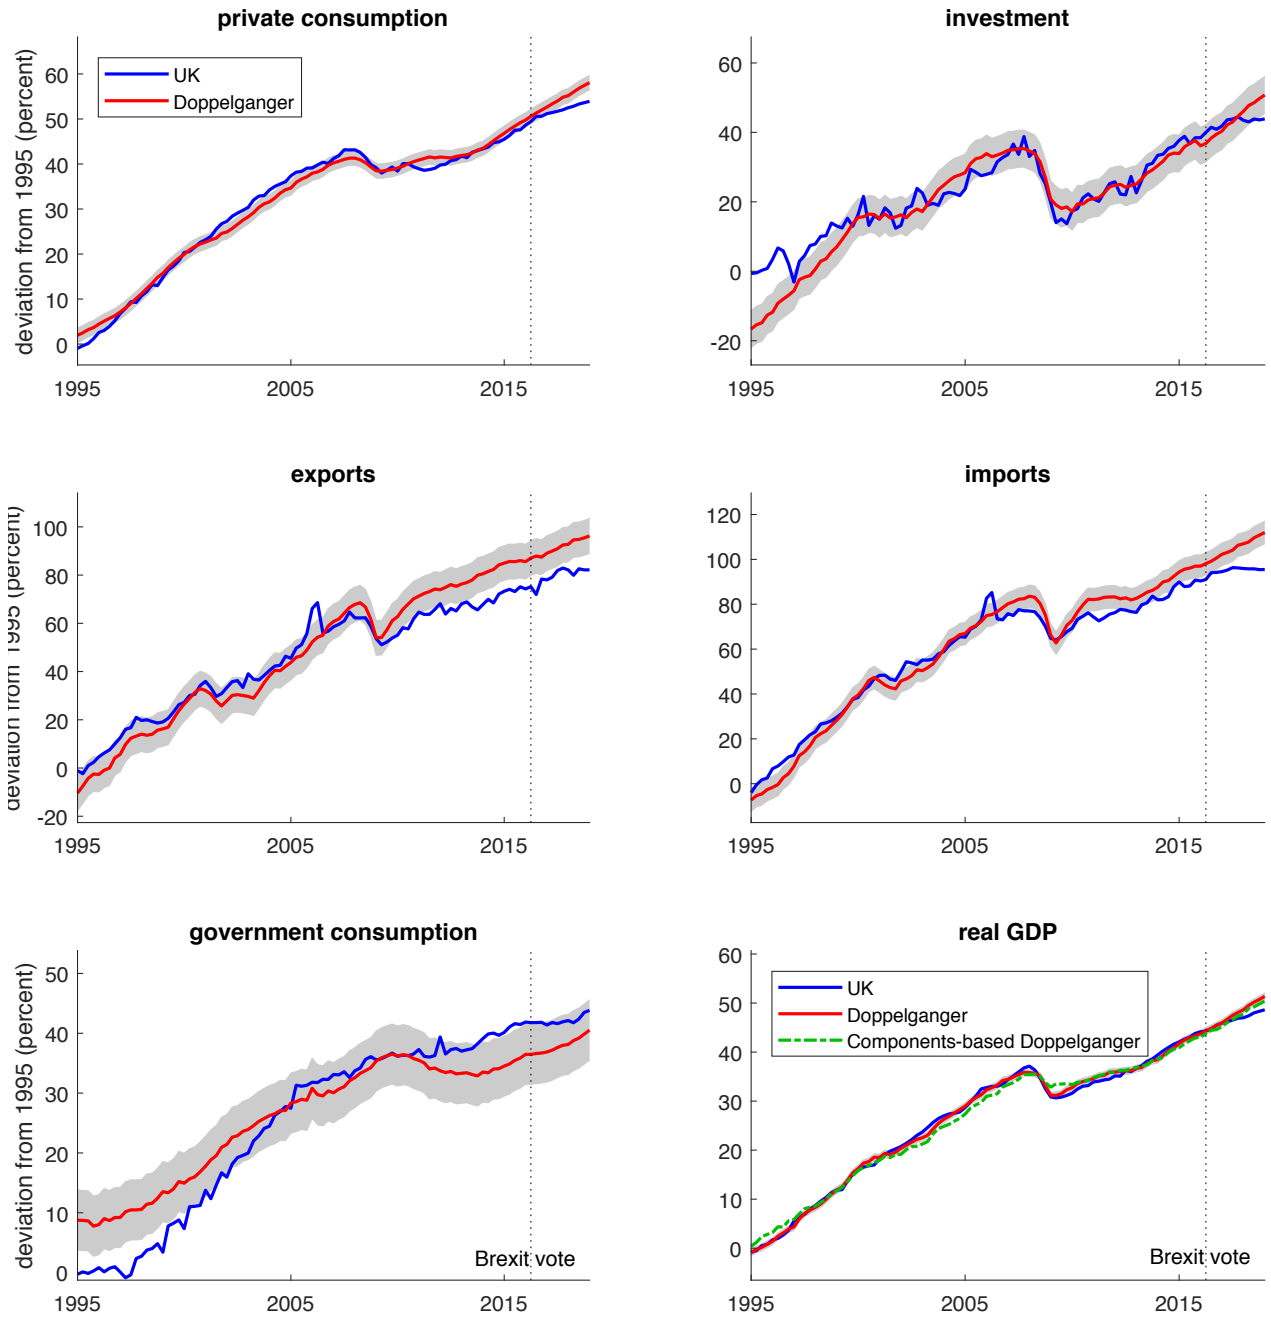


Figure 7: GDP components: UK (blue) and doppelgänger (red). *Note:* shaded area is one standard deviation of difference prior to Brexit vote. Data source: OECD Economic Outlook.

After all, it is clear that Brexit will amount to a bundle of policy measures which will result in economic disintegration between the UK and the European Union. Whether this is because of higher tariffs, non-tariff barriers or both, it is likely to bring about a reduction of living standards which, in turn, may rationalize reduced investment and consumption expenditures: not only in the future, but—because of anticipation effects—already today.

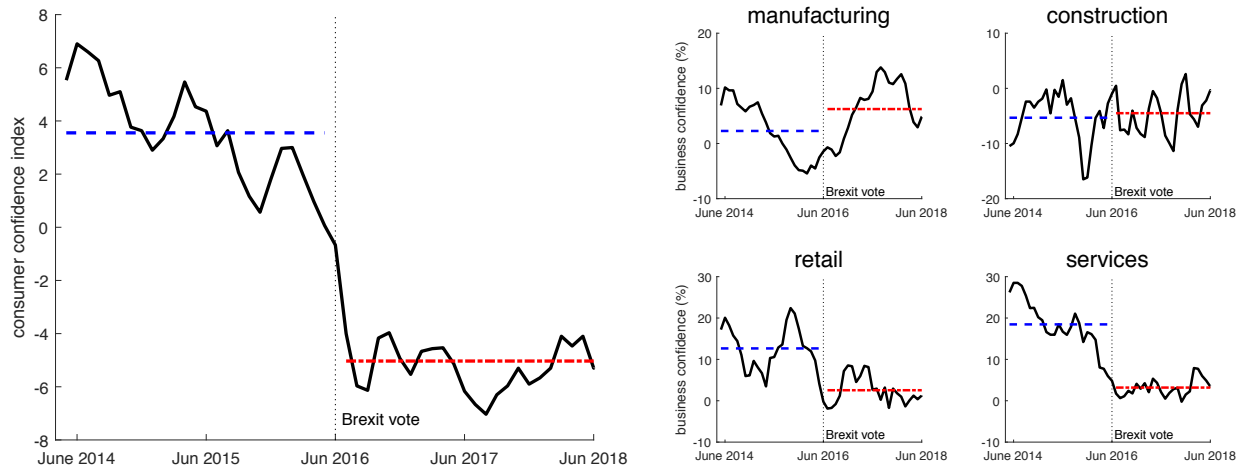


Figure 8: Consumer and business confidence. *Note:* left panel shows consumer confidence and the right panel shows business confidence taken from the OECD Economic Outlook database. All are expressed as balances in percent.

This notion is supported by data on consumer and business confidence taken from the same Economic Outlook database, see Figure 8. The left panel shows that consumer confidence dropped strongly around the Brexit vote and, more importantly, that it remained low ever since. The right panel shows the same for business confidence, but here the tendencies are mixed across sectors. While manufacturing sentiment increased somewhat, possibly driven by the devaluation, construction confidence was, by and large, unaffected. Retail and service industry mimic the more gloomy outlook of consumers.

3.2 The role of uncertainty and anticipation effects

The Brexit vote has led households and firms to reduce their expenditures. This may reflect “anticipation effects” because households and firms expect Brexit to lower prosperity eventually. However, in addition to a possible downgrade in the average economic outlook, the Brexit vote also increased economic uncertainty considerably—not least because the details of Brexit are still unclear. Higher economic uncertainty is likely to take its toll on investment and consumption expenditures, quite independently of any anticipation effects. In fact, even if the economic outlook were unchanged on average, an increase of uncertainty will hamper economic activity, as established in a seminal contribution by Bloom (2009).²⁰

In what follows, we explicitly quantify the extent to which the doppelgänger gap identified above is due to (i) anticipation effects of reduced future prosperity, and (ii) more dispersed

²⁰For a simultaneous analysis of anticipation and uncertainty shocks, see e.g. Cascaldi-Garcia and Galvao (2018); Forni *et al.* (2017); Song and Tang (2018).

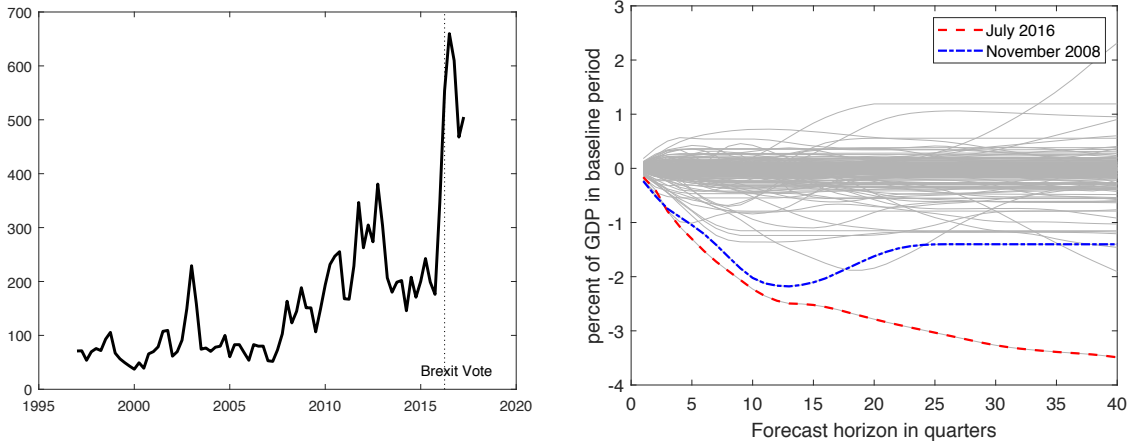


Figure 9: Increase of uncertainty and downgrade of output expectations after the Brexit vote. *Note:* left panel shows the index of economic policy uncertainty (EPU, source: www.policyuncertainty.com); right panel shows cumulated month-to-month changes of output growth forecasts by *Oxford Economics* with the downgrades between July and June 2016 (dashed red line) and November and October 2008 (dash-dotted blue line) highlighted.

expectations, that is, uncertainty effects. Specifically, we estimate a structural EVAR and identify shocks to uncertainty and expectations. Once again using the notion that the Brexit vote is a well-defined natural experiment allows us to single out uncertainty and expectations shocks occurring in 2016Q3 as those *caused* by the referendum. The estimated model, together with these identified “Brexit shocks”, enables us to quantify to what extent the doppelganger gap is caused by anticipation or uncertainty effects.

3.2.1 Uncertainty and expectations data

Given that the Brexit vote has primarily uncertain consequences for future policies, we measure uncertainty using the Economic Policy Uncertainty (EPU) index. The index is based on a (standardized) count of newspaper articles containing the terms uncertain or uncertainty, economic or economy, and one or more policy-relevant terms (see Baker *et al.*, 2016). The left panel of Figure 9 shows that the EPU index increased dramatically around the Brexit referendum. For our application, however, it is especially important that it captures *mean-preserving* changes in policy uncertainty. Baker *et al.* (2016, Table IV) show that controlling for various proxies of future expectations changes little of their results.²¹

To capture anticipation effects we rely on proprietary data of the professional forecasting firm *Oxford Economics* which provides growth forecasts for the UK up until the year 2050.²²

²¹In the online appendix, we consider an additional measure of macroeconomic uncertainty proposed by Jurado *et al.* (2015) and computed for the UK by Redl (2017).

²²Clearly, the forecasts are dependent of the particular model used by Oxford Economics and may not reflect “true” expectations in the economy. Ideally, we would use forecasts from a variety of

Indeed, there is little doubt that the Brexit vote induced market participants to reduce their long-term income expectations. This is exemplified in the right panel of Figure 9, where we display month-to-month changes of *Oxford Economics*' output growth forecasts throughout our available sample, cumulated over a ten year forecast horizon. The figure shows that the forecast revisions in response to the Brexit vote, that is, the difference between growth forecasts in July and June 2016, were unprecedented in size *and persistence* even when compared to the Great Recession period.²³ Sampson (2017) surveys studies which quantify the per capita income loss due to Brexit and finds plausible estimates range between -1 and -10 percent for a forecast horizon of 10 or more years after Brexit. The downgrade of output growth by *Oxford Economics* after the referendum is consistent with these estimates.

3.2.2 Estimation

In order to quantify the uncertainty and anticipation effects that are manifest in the doppelganger gap, we estimate an EVAR on quarterly time series. The VAR features news regarding future output growth in addition to conventional variables. Specifically, letting $x_{t+h,t}$ denote the h -quarter ahead output growth forecast in period t , and $x_{t+h,t-1}$ the output growth forecast for the same period made one quarter before, we define $news_{t+h,t} \equiv x_{t+h,t} - x_{t+h,t-1}$. Formally, we use \mathbf{y}_t to denote the vector of endogenous variables of our EVAR

$$\mathbf{y}_t = \left[\begin{array}{cccccccc} EPU_t & news_{t+h_1,t} & news_{t+h_2,t} & news_{t+h_3,t} & r_t & y_t & \pi_t & s_t \end{array} \right]' . \quad (6)$$

It includes the log of the EPU, EPU_t , news which relate to three different forecasting horizons, as well as the bank rate of the Bank of England, r_t , and the log of real GDP, y_t , inflation π_t , and the log of the nominal effective exchange rate, s_t . We include inflation and the exchange rate into the EVAR to account for the “real squeeze” channel investigated by Breinlich *et al.* (2017) by which the depreciation of the pound affected consumer prices and in turn real consumption. In principle, this channel may have operated independently of uncertainty and

firms. However, *Oxford Economics* stands out in terms of forecasting horizon. Reassuringly, we find that forecast revisions by *Oxford Economics* for the short run are very similar to the average forecast revision by a large group of professional forecasters published by Her Majesty's Treasury (see <https://www.gov.uk/government/collections/data-forecasts>), see Figure 7 in the online appendix. See also Born *et al.* (forthcoming) for a discussion on the quality of the *Oxford Economics* forecasts. Interestingly, Figure 7 in the online appendix also shows that the drop in output growth expectations is mostly due to a decrease in expected consumption and investment.

²³Growth forecasts are more strongly downgraded in the short run. However, even long-horizon growth forecasts were substantially downgraded resulting in a persistent fall of cumulated output losses.

expectation shocks.²⁴ We then estimate the model:

$$\mathbf{y}_t = \mathbf{c} + \mathbf{A}(L)\mathbf{y}_{t-1} + \nu_t, \quad (7)$$

where \mathbf{c} is a constant term, $\mathbf{A}(L)$ is a lag polynomial, and $\nu_t \sim (0, \mathbf{\Omega})$ is a vector of white noise errors.

Model (7) is an expectations-augmented VAR (EVAR).²⁵ We use it in our analysis for two reasons. First, conventional VAR models face difficulties when it comes to recovering anticipation effects (Lippi and Reichlin, 1994; Fernández-Villaverde *et al.*, 2007; Leeper *et al.*, 2013).²⁶ Hence, several contributions, notably in the context of fiscal policy, have extended traditional VAR models in order to control directly for “foresight” of market participants, by including either narratively identified measures of anticipated shocks or data on expectations (Ramey, 2011; Mertens and Ravn, 2012). Leduc and Sill (2013) also include survey expectations for the unemployment rate in an otherwise conventional VAR model to assess the contribution of changes in expectations to economic fluctuations.

Second, the EVAR specification allows us to pursue a semi-structural identification strategy. As we identify the anticipation effects of Brexit, we capture the possibility that news regarding future output growth impact the economy. In doing so, we account for news which relate to a wide range of forecast horizons, but remain agnostic as to its specific causes. For instance, one may think of growth news as ultimately being due to expected changes in total factor productivity (see, for instance, Beaudry and Portier, 2006). Alternatively, output expectations may decline because of economic disintegration and reduced gains from trade. We do not take a stand in this regard because a broad perspective seems warranted in light of the multifaceted event that looms on the horizon.

We identify uncertainty and growth news shocks on the basis of a recursive identification scheme (or, equivalently, through a Choleski decomposition of Ω). This identification strategy has been commonly pursued in the literature on uncertainty shocks. As we order EPU_t first in (6), we allow uncertainty shocks to play the largest possible role: all variables may respond contemporaneously to an uncertainty shock. Growth news, in turn, are ordered second such that they may respond contemporaneously to uncertainty shocks—a likely scenario, notably

²⁴In the online appendix, we use the estimated EVAR to investigate to what extent the exchange rate depreciation was driven by the uncertainty and expectation shocks. The results suggest that the bulk of the exchange rate movement following the Brexit vote can indeed be accounted for by the identified shocks.

²⁵Perotti (2014) suggests the label “EVAR”.

²⁶The moving average representation of structural models with foresight is often non invertible, or non fundamental given the set of variables that are typically included in VAR models. This provides a rationale for estimating structural models using full-information econometric methods (e.g., Blanchard *et al.*, 2013; Schmitt-Grohé and Uribe, 2012). Mertens and Ravn (2010) also rely on a theoretical model to develop an augmented SVAR estimator that is able to identify fiscal shocks in the face of anticipation.

for short-term forecasts. Our specification also allows for a contemporaneous effect of growth news shocks on output, interest rates, inflation, and the exchange rate, but rules out an immediate effect on uncertainty. We do not provide a structural interpretation of the other shocks in the model such that the ordering of output, interest rates, inflation, and the exchange rate relative to each other is immaterial.²⁷

We estimate model (7) on observations for the period 1997Q1–2018Q2.²⁸ We allow for four lags and set the forecast horizon of news, h_1 , h_2 and h_3 to 8, 12, and 16 quarters, respectively. Our closest forecast horizon is 8 quarters because this allows us to capture growth news shocks in response to the Brexit vote that relate to the period just beyond the end of our sample: given $h_1 = 8$, Brexit news that potentially materialize in 2016Q3 pertain to growth in 2018Q3. In addition, we consider news regarding output growth that is expected to take place after another 1 and 2 years, respectively.²⁹

3.2.3 Quantifying uncertainty and anticipation effects

We are now ready to quantify to what extent uncertainty and anticipation effects drive the estimated doppelganger gap. Towards this end, we use the estimated model to construct a counterfactual GDP path. Specifically we switch off uncertainty and news shocks in 2016Q3, the quarter immediately after the Brexit vote. Our maintained assumption is that these shocks are caused by the referendum: we once more interpret the Brexit vote as a natural experiment, just like in our analysis in Section 2 above. As, by construction, all eight shocks taken together generate time-paths of the endogenous variables which exactly track the data, switching off the Brexit shocks provides us with the counterfactual GDP path which would have been observed in the absence of the Brexit vote.³⁰

Figure 10 shows the results together with the actual time path of GDP in the UK. The figure also reproduces the UK doppelganger which was computed independently in Section 2 above. The dashed line represents the implied path for GDP, had there been no uncertainty shock in 2016Q3. The distance between the counterfactual, “no uncertainty” GDP path

²⁷In the online appendix, we perform a number of robustness exercises which include the relaxation of our identification assumptions. There, we also display impulse response functions to an uncertainty shock which we find to be in line with what Baker *et al.* (2016) find for US data.

²⁸Further details on the data can be found in the online appendix.

²⁹Including news for longer horizons gives rise to unstable results. The likely reason is that changes of forecasts for the very long run are fairly volatile.

³⁰In the online appendix, we also consider an alternative identification approach based on dummy variables. Specifically, we include dummy variables into our baseline reduced-form EVAR which are equal to 1 in 2016Q3 and zero otherwise. In the same way as in our baseline exercise, we interpret the coefficients on these dummy variables as measuring the effect of the Brexit vote on the respective variables. Constructing a counterfactual time path for real GDP by “switching off” the effects of uncertainty and news measured by the dummy variables gives very similar results to our baseline exercise.

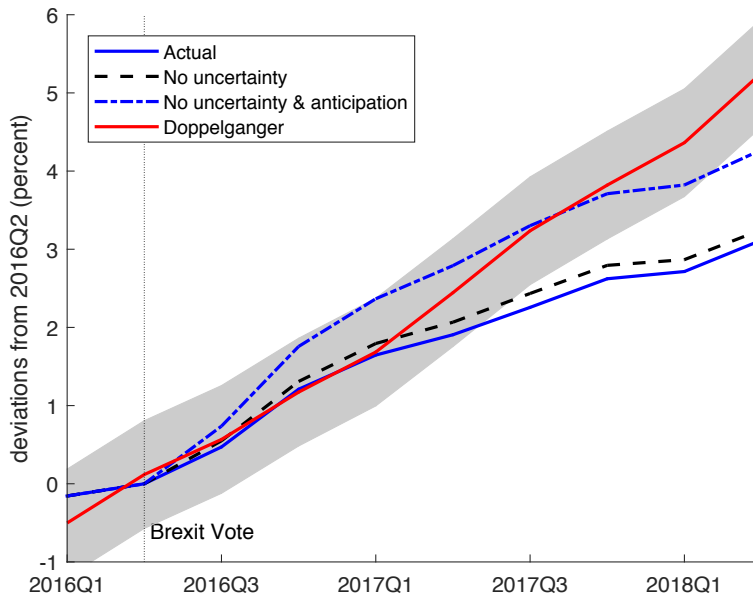


Figure 10: Output path in the UK. *Note:* actual GDP (blue solid line) vs doppelganger (red solid line) vs counterfactual GDP path when uncertainty shock in 2016Q3 is switched off (dashed line) and when both uncertainty and news shocks in 2016Q3 are switched off (dash-dotted line).

and actual GDP is rather small. We find that the increase of uncertainty due to the Brexit vote explains on average about 20 percent of the doppelganger gap—although the relative contribution declines towards the end of the sample.

The dashed-dotted line in Figure 10 shows the counterfactual GDP path that obtains when we switch off the news shocks in 2016Q3 in addition to the uncertainty shock (dashed-dotted line). This GDP path would have been observed in the absence of the Brexit vote according to our estimated VAR model. Notice that it aligns fairly well with the output path for the doppelganger. Taken together, uncertainty and anticipation effects, originating *only* in 2016Q3, account almost fully for the doppelganger gap. Initially, they over predict output relative to the doppelganger somewhat.³¹ Towards the end, they underpredict it. Throughout the contribution of news effects dominates the contribution of uncertainty effects.

In the online appendix, we report results of various robustness checks. Overall, we find that the role of uncertainty in accounting for the doppelganger gap is sizable, but it does not explain more than 30 percent. This is noteworthy, because our identification scheme allows uncertainty shocks the largest possibility to impact economic activity. Despite this—heightened uncertainty explains only a modest share of the overall losses. We find that the remaining gap is accounted for by anticipation effects.

³¹In a robustness check, we use a simple AR model and relate the doppelganger gap obtained on the basis of the synthetic control approach directly to the identified uncertainty and anticipation shocks. Here too, we find that both shocks can almost completely account for the gap (results are available upon request).

4 Conclusion

In this paper, we exploit the natural experiment of the Brexit vote to quantify the costs of economic disintegration. Natural experiments in macroeconomics are rare, but when they occur they offer unique insights into causal mechanisms and the validity of major assumptions underlying macroeconomic models. The unexpected outcome of the Brexit referendum in June 2016 offers such a window on causal relationships.

Our first main result is that the Brexit vote has already impacted economic activity well before any policy change has occurred. We show that by the end of 2018 the Brexit vote has caused a reduction of GDP by 1.7-2.5 percent. Zooming in on the behavior of the private sector, we find that households and firms have adjusted their behavior in anticipation of Brexit, as macroeconomic theory predicts. We observe a considerable decline of consumption and investment in response to the Brexit vote.

However, while the Brexit referendum shapes Britain's economic present in addition to its future it also raised economic uncertainty. The binary choice question "Should the United Kingdom remain a member of the European Union or leave the European Union?" left important issues open. And while the direction of future economic policies seems clear, the exact extent of Britain's economic disintegration from Europe remains unclear.

Still, we find that a wider dispersion of future economic outcomes, that is, heightened economic policy and macroeconomic uncertainty, accounts for only 20 percent of the observed effects. Stripping the overall output loss due to the Brexit vote of the effect of heightened uncertainty—which is arguably temporary—leaves us with anticipation effects of households and firms, which have downgraded their expectations about future incomes. These effects are not only large, but to the extent that they reflect long-run outcomes, they are also there to stay.

References

- Abadie, A., Diamond, A. and Hainmueller, J. (2010). “Synthetic control methods for comparative case studies: estimating the effect of California’s tobacco control program”, *Journal of the American Statistical Association*, vol. 105(490), pp. 493–505.
- Abadie, A., Diamond, A. and Hainmueller, J. (2015). “Comparative politics and the synthetic control method”, *American Journal of Political Science*, vol. 59(2), pp. 495–510.
- Abadie, A. and Gardeazabal, J. (2003). “The economic costs of conflict: A case study of the Basque country”, *American Economic Review*, vol. 93(1), pp. 113–132.
- Alabrese, E., Becker, S.O., Fetzer, T. and Novy, D. (2019). “Who voted for Brexit? Individual and regional data combined”, *European Journal of Political Economy*, vol. 56, pp. 132 – 150.
- Alesina, A. and Fuchs-Schündeln, N. (2007). “Goodbye Lenin (or not?): The effect of communism on people”, *American Economic Review*, vol. 97(4), pp. 1507–1528.
- Andrews, D.W.K. (2003). “End-of-sample instability tests”, *Econometrica*, vol. 71(6), pp. 1661–1694.
- Baker, S.R., Bloom, N. and Davis, S.J. (2016). “Measuring economic policy uncertainty”, *Quarterly Journal of Economics*, vol. 131(4), pp. 1593–1636.
- Barsky, R.B. and Sims, E.R. (2011). “News shocks and business cycles”, *Journal of Monetary Economics*, vol. 58(3), pp. 273 – 289.
- Barsky, R.B. and Sims, E.R. (2012). “Information, animal spirits, and the meaning of innovations in consumer confidence”, *American Economic Review*, vol. 102(4), pp. 1343–77.
- Beaudry, P. and Portier, F. (2006). “Stock prices, news, and economic fluctuations”, *American Economic Review*, vol. 96(4), pp. 1293–1307.
- Becker, S.O., Fetzer, T. and Novy, D. (2017). “Who voted for Brexit? A comprehensive district-level analysis”, *Economic Policy*, vol. 32(92), pp. 601–650.
- Berg, T., Saunders, A., Schäfer, L. and Steffen, S. (2017). ““Brexit” and the contraction of syndicated lending”, Mimeo, Frankfurt School of Finance & Management.
- Billmeier, A. and Nannicini, T. (2013). “Assessing economic liberalization episodes: A synthetic control approach”, *The Review of Economics and Statistics*, vol. 95(3), pp. 983–1001.

- Blanchard, O., L’Huillier, J.P. and Lorenzoni, G. (2013). “News, noise, and fluctuations: An empirical investigation”, *American Economic Review*, vol. 103, pp. 3045–70.
- Bloom, N. (2009). “The impact of uncertainty shocks”, *Econometrica*, vol. 77(3), pp. 623–685.
- Born, B., Müller, G.J. and Pfeifer, J. (forthcoming). “Does austerity pay off?”, *Review of Economics and Statistics*.
- Born, B. and Pfeifer, J. (2014). “Policy risk and the business cycle”, *Journal of Monetary Economics*, vol. 68, pp. 68–85.
- Breinlich, H., Leromain, E., Novy, D. and Sampson, T. (2017). “The consequences of the Brexit vote for UK inflation and living standards: First evidence”, Mimeo, London School of Economics.
- Campos, N.F., Coricelli, F. and Moretti, L. (forthcoming). “Institutional integration and economic growth in Europe”, *Journal of Monetary Economics*.
- Cascaldi-Garcia, D. and Galvao, A.B. (2018). “News and uncertainty shocks”, International Finance Discussion Papers 1240.
- Davies, R.B. and Studnicka, Z. (2018). “The heterogeneous impact of Brexit: Early indications from the FTSE”, *European Economic Review*, vol. 110, pp. 1 – 17.
- Dhingra, S., Huang, H., Ottaviano, G., Pessoa, J.P., Sampson, T. and Reenen, J.V. (2017). “The costs and benefits of leaving the EU: trade effects”, *Economic Policy*, vol. 32(92), pp. 651–705.
- Eaton, J. and Kortum, S. (2002). “Technology, geography, and trade”, *Econometrica*, vol. 70(5), pp. 1741–1779.
- Ferman, B. and Pinto, C. (forthcoming). “Inference in differences-in-differences with few treated groups and heteroskedasticity”, *The Review of Economics and Statistics*.
- Fernández-Villaverde, J., Guerrón-Quintana, P.A., Kuester, K. and Rubio-Ramírez, J.F. (2015). “Fiscal volatility shocks and economic activity”, *American Economic Review*, vol. 105(11), pp. 3352–3384.
- Fernández-Villaverde, J., Rubio-Ramírez, J.F., Sargent, T.J. and Watson, M.W. (2007). “ABCs (and Ds) of understanding VARs”, *American Economic Review*, vol. 97(3), pp. 1021–1026.

- Fetzer, T. (2018). “Did austerity cause Brexit?”, CAGE Working Paper No. 381.
- Forni, M., Gambetti, L. and Sala, L. (2017). “News, uncertainty and economic fluctuations (no news is good news)”, Center for Economic Research Working Paper 132.
- Fuchs-Schündeln, N. and Hassan, T. (2016). “Natural experiments in macroeconomics”, in (J. B. Taylor and H. Uhlig, eds.), *The Handbook of Macroeconomics*, pp. 923–1012, vol. 2, Elsevier.
- Goldberg, P. and Pavcnik, N. (2016). “The effects of trade policy”, in (K. Bagwell and R. Staiger, eds.), *Handbook of Commercial Policy*, pp. 161–206, North Holland.
- Hahn, J. and Shi, R. (2017). “Synthetic control and inference”, *Econometrics*, vol. 5(4), p. 52.
- Handley, K. and Limão, N. (2015). “Trade and investment under policy uncertainty: Theory and firm evidence”, *American Economic Journal: Economic Policy*, vol. 7(4), pp. 189–222.
- Handley, K. and Limão, N. (2017). “Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States”, *American Economic Review*, vol. 107(9), pp. 2731–2783.
- Jurado, K., Ludvigson, S.C. and Ng, S. (2015). “Measuring uncertainty”, *American Economic Review*, vol. 105(3), pp. 1177–1216.
- Kaul, A., Klößner, S., Pfeifer, G. and Schieler, M. (2018). “Synthetic control methods: Never use all pre-intervention outcomes together with covariates”, Mimeo, University of Hohenheim.
- Leduc, S. and Sill, K. (2013). “Expectations and economic fluctuations: An analysis using survey data”, *The Review of Economics and Statistics*, vol. 95(4), pp. 1352–1367.
- Leeper, E.M., Walker, T.B. and Yang, S.S. (2013). “Fiscal foresight and information flows”, *Econometrica*, vol. 81(3), pp. 1115–1145.
- Limão, N. and Maggi, G. (2015). “Uncertainty and trade agreements”, *American Economic Journal: Microeconomics*, vol. 7(4), pp. 1–42.
- Lippi, M. and Reichlin, L. (1994). “VAR analysis, nonfundamental representations, Blaschke matrices”, *Journal of Econometrics*, vol. 63(1), pp. 307–325.
- Melitz, M.J. (2003). “The impact of trade on intra-industry reallocations and aggregate industry productivity”, *Econometrica*, vol. 71(6), pp. 1695–1725.

- Mertens, K. and Ravn, M.O. (2010). “Measuring the impact of fiscal policy in the face of anticipation: A structural VAR approach”, *The Economic Journal*, vol. 120(544), pp. 393–413.
- Mertens, K. and Ravn, M.O. (2011). “Understanding the aggregate effects of anticipated and unanticipated tax policy shocks”, *Review of Economic Dynamics*, vol. 14(1), pp. 27–54.
- Mertens, K. and Ravn, M.O. (2012). “Empirical evidence on the aggregate effects of anticipated and unanticipated US tax policy shocks”, *American Economic Journal: Economic Policy*, vol. 4(2), pp. 145–81.
- Nakamura, E. and Steinsson, J. (2014). “Fiscal stimulus in a monetary union: Evidence from US regions”, *American Economic Review*, vol. 104(3), pp. 753–792.
- Nakamura, E. and Steinsson, J. (2018). “Identification in macroeconomics”, *Journal of Economic Perspectives*, vol. 32(3), pp. 59–86.
- Novy, D. and Taylor, A.M. (2014). “Trade and uncertainty”, NBER Working Paper 19941.
- Parker, J.A., Souleles, N.S., Johnson, D.S. and McClelland, R. (2013). “Consumer spending and the economic stimulus payments of 2008”, *American Economic Review*, vol. 103(6), pp. 2530–2553.
- Perotti, R. (2014). “Defense government spending is contractionary, civilian government spending is expansionary”, NBER Working Papers 20179.
- Ramey, V.A. (2011). “Identifying government spending shocks: It’s all in the timing”, *Quarterly Journal of Economics*, vol. 126(1), pp. 1–50.
- Ramiah, V., Pham, H. and Moosa, I. (2016). “The sectoral effects of Brexit on the British economy: early evidence from the reaction of the stock market”, *Applied Economics*, pp. 1–7.
- Redl, C. (2017). “The impact of uncertainty shocks in the United Kingdom”, Bank of England Working Paper No. 695.
- Saia, A. (2017). “Choosing the open sea: The cost to the UK of staying out of the euro”, *Journal of International Economics*, vol. 108, pp. 82 – 98.
- Sampson, T. (2017). “Brexit: The economics of international disintegration”, *Journal of Economic Perspectives*, vol. 31(4), pp. 163–184.

Schmitt-Grohé, S. and Uribe, M. (2012). “What’s news in business cycles”, *Econometrica*, vol. 80(6), pp. 2733–2764.

Song, D. and Tang, J. (2018). “News-driven uncertainty fluctuations”, Federal Reserve Bank of Boston, Working Paper 18-3.

Online Appendix

The Costs of Economic Nationalism: Evidence from the Brexit Experiment*

Benjamin Born

Gernot J. Müller
Petr Sedláček

Moritz Schularick

May 2, 2019

A Data

The donor group contains all OECD countries for which we were able to obtain contiguous real GDP data starting in 1995Q1 and pre-Brexit-vote averages of the additional covariates – allowing for missing values (see Table 2 for a list). If not otherwise noted, the data comes from OECD Economic Outlook database (November 2018 edition). For the baseline, we use real GDP (gross domestic product, volume, market prices), consumption (private final consumption expenditure, volume), investment (gross fixed capital formation, total, volume), real export (exports of goods and services, volume) and import (imports of goods and services, volume), employment (total employment, labour force survey basis), and annual population (working-age population, age 15-74). The latter is linearly interpolated to the quarterly frequency. For the decomposition exercise, we additionally use government consumption (government final consumption expenditure, volume). For the VAR exercises, we also use inflation (change in consumer price index, harmonized, index 2015) and the nominal exchange rate (nominal effective exchange rate, chain-linked, overall weights).

For the uncertainty analysis, we use the Baker *et al.* (2016) Economic Policy Uncertainty index available at www.policyuncertainty.com. The index is based on a (standardized) count of newspaper articles containing the terms uncertain or uncertainty, economic or economy, and one or more policy-relevant terms. The macroeconomic uncertainty index, based on Jurado *et al.* (2015), has been computed for the UK in Redl (2017) and has been made available to us by Chris Redl. Stock market returns are based on the Datastream total

*Born: Frankfurt School of Finance & Management and CEPR, b.born@fs.de, Müller: University of Tübingen and CEPR, gernot.mueller@uni-tuebingen.de, Schularick: University of Bonn and CEPR, schularick@uni-bonn.de, Sedláček: University of Oxford and CEPR, petr.sedlacek@economics.ox.ac.uk.

market total returns index for the UK (TOTMKUK). The policy rate is the Bank of England bank rate.

The proprietary real output growth forecasts come from Oxford Economics (<https://www.oxfordeconomics.com>). We aggregate the monthly growth forecasts to quarterly frequency by taking the average over the three months pertaining to a quarter.

B SMC: details and robustness

B.1 End-of-sample instability test

Andrews’ (2003) end-of-sample instability test has recently been suggested in the context of synthetic control approaches by Hahn and Shi (2017) and shown to have good size properties. While the test is technically based on stationary data, Andrews (2003) notes (p. 1681, comment 4), that his test can be shown to be asymptotically valid under stationary errors. To conduct the test, we run the SCM over the *whole* (pre- and post-Brexit vote) sample and then base the test statistic on the sum-of-squares of the post-Brexit vote errors. Following Andrews (2003), the distribution of the test statistic is computed using a subsampling scheme. Specifically, we conduct the matching on the sample $1, \dots, T_0$, where observations $j, \dots, j + \lceil m/2 \rceil - 1$ are excluded. Here, m is the number of post-Brexit vote observations, T_0 is the time of the treatment, and we resample for $j = 1, \dots, T_0 - m + 1$. For each iteration, the resampled test statistic is based on the matching errors from j to $j + m - 1$.

B.2 Additional country placebo tests

The main text conducted placebo tests on the relevant subset of the donor pool, i.e. the countries which received a non-zero weight in the baseline UK doppelganger. First, we report the standard deviations of the differences between the respective doppelgangers and actual time paths of GDP in each of these countries. This is done in Table 1 below.

Table 1: Country placebos: fit

| | | | | | | | |
|---------|------|-------------|------|---------------|------|----------------|------|
| Germany | 0.02 | Hungary | 0.03 | Iceland | 0.03 | Ireland | 0.08 |
| Italy | 0.03 | New Zealand | 0.02 | United States | 0.03 | United Kingdom | 0.01 |

Standard deviation of the difference between the doppelganger and actual real GDP paths in the respective donor pool countries prior to the Brexit vote. The UK (our baseline estimate) is shown for comparability.

Next, this appendix shows that similar results are obtained for the rest of the donor

pool countries. Figure 1 shows the country placebo time paths and Figure 2 depicts the two statistics of relative pre- and post-treatment fit. Again, the UK economy stands out as the one with the largest post-treatment deviation.

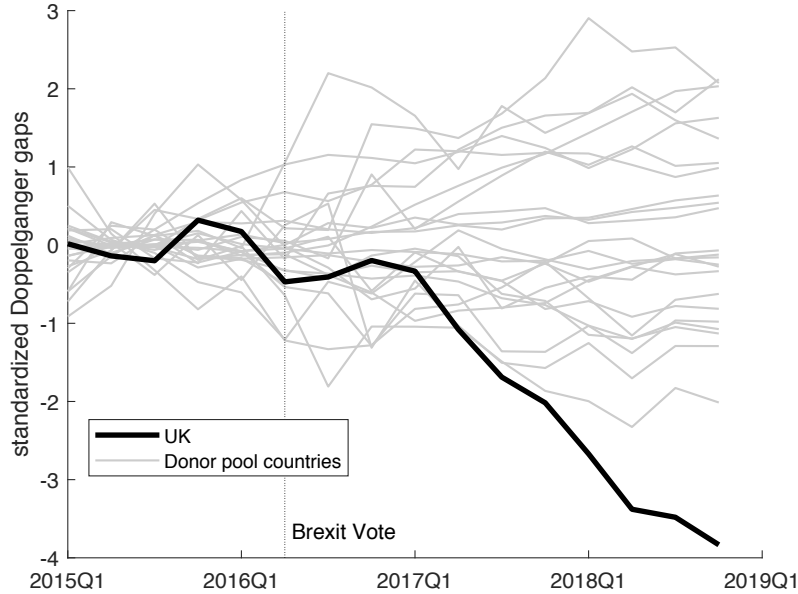


Figure 1: Country placebo tests *Note:* UK doppelganger *gap* (thick black line), with gray lines representing country placebo doppelganger gaps estimated by considering fictitious Brexit votes in all donor pool economies. For comparability, all doppelganger gaps are normalized by their respective pre-Brexit standard deviations and centred around their 2015 means.

B.3 Western donor pool robustness

This section provides an additional robustness check with respect to the donor pool. In particular, we restrict the donor pool to include countries that are likely most similar to the UK and exclude all east and south European economies (Hungary, Italy, Portugal, Slovak Republic, Spain) as well as Luxembourg. Nevertheless, using this *a priori* restricted donor pool gives similar results as suggested by our donor pool exercises in the main text. The cost of the Brexit vote at the end of 2018 is 1.7 percent of GDP. While this is lower than the 2.4 percent in our baseline specification, it is nevertheless still sizeable.




Figure 2: Relative measures of the pre- and post-treatment doppelganger gaps. *Note:* left panel shows the relative maximum absolute prediction error ρ_2 , the right panel shows the relative root means squared prediction error ρ_1 .

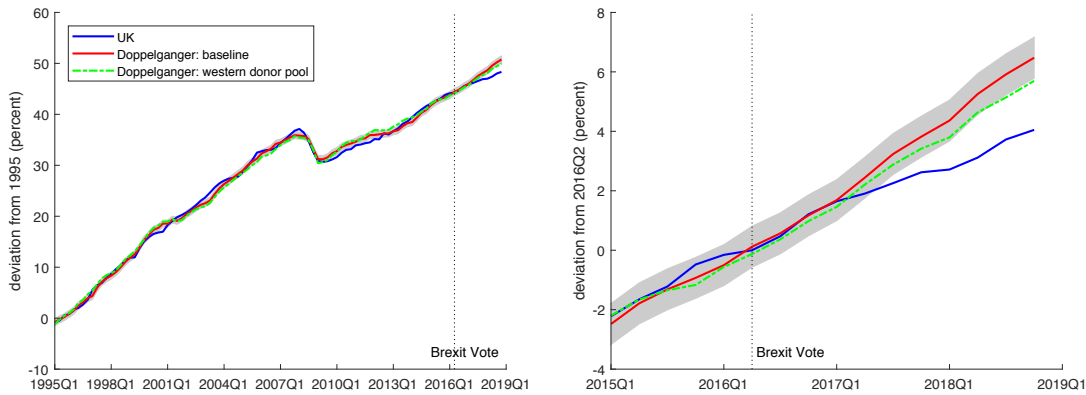


Figure 3: Real GDP of the UK. Actual data (blue line), baseline doppelganger (red line), doppelganger based on restricted donor pool (green line). *Note:* shaded area is one standard deviation of difference prior to Brexit vote. Data source: OECD Economic Outlook.

B.4 Alternative covariate specification

In the baseline, we match on a number of covariates where we average over the sample 1995Q1 to 2016Q2. In this robustness check, we average the covariates instead over the year prior to the Brexit vote (2015Q2 to 2016Q2). The results shown below indicate that while this affects the country weights (the US become even more important), the results remain unchanged with regard to the output effect of the Brexit vote.

Table 2: Matching of covariates

| | UK | Doppelganger |
|---------------------------|-------|--------------|
| Consumption / GDP | 65.09 | 63.68 |
| Investment / GDP | 16.73 | 20.99 |
| Exports / GDP | 28.42 | 28.45 |
| Imports / GDP | 29.34 | 25.67 |
| Labor productivity growth | -0.08 | -0.07 |
| Employment share | 65.16 | 61.90 |

Note: All numbers are in percent. Labor productivity growth is the log difference between quarterly real GDP and quarterly total employment; employment share is the ratio between total employment and the working age population.

Table 3: Composition of the doppelganger: country weights

| | | | | | | | |
|-----------|-------|-------------|-------|-----------------|-------|-------------|-------|
| Australia | <0.01 | Austria | <0.01 | Belgium | <0.01 | Canada | <0.01 |
| Finland | <0.01 | France | <0.01 | Germany | <0.01 | Hungary | 0.13 |
| Iceland | 0.02 | Ireland | <0.01 | Italy | 0.07 | Japan | 0.11 |
| Korea | <0.01 | Luxembourg | <0.01 | Netherlands | <0.01 | New Zealand | 0.04 |
| Norway | <0.01 | Portugal | <0.01 | Slovak Republic | <0.01 | Spain | <0.01 |
| Sweden | <0.01 | Switzerland | <0.01 | United States | 0.63 | | |

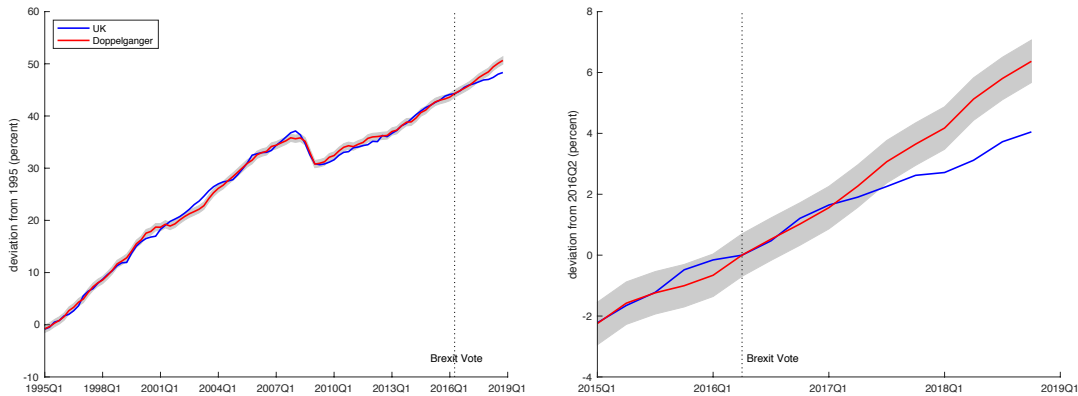


Figure 4: Real GDP of the UK. Actual data (blue line) vs doppelganger (red line). *Note:* shaded area is one standard deviation of difference prior to Brexit vote. Data source: OECD Economic Outlook.

C VAR: details and robustness

This part of the appendix provides additional results for our VAR analysis in the main text. First, we show that the aggregate dynamics following the identified uncertainty shocks are consistent with findings in the literature. Second, we perform a series of robustness checks on our baseline specification.

C.1 Impulse response functions

Figure 5 shows the impulse response functions of the VAR variables to a one standard deviation uncertainty shock. As has been extensively discussed in the literature, uncertainty increases have contractionary effects.

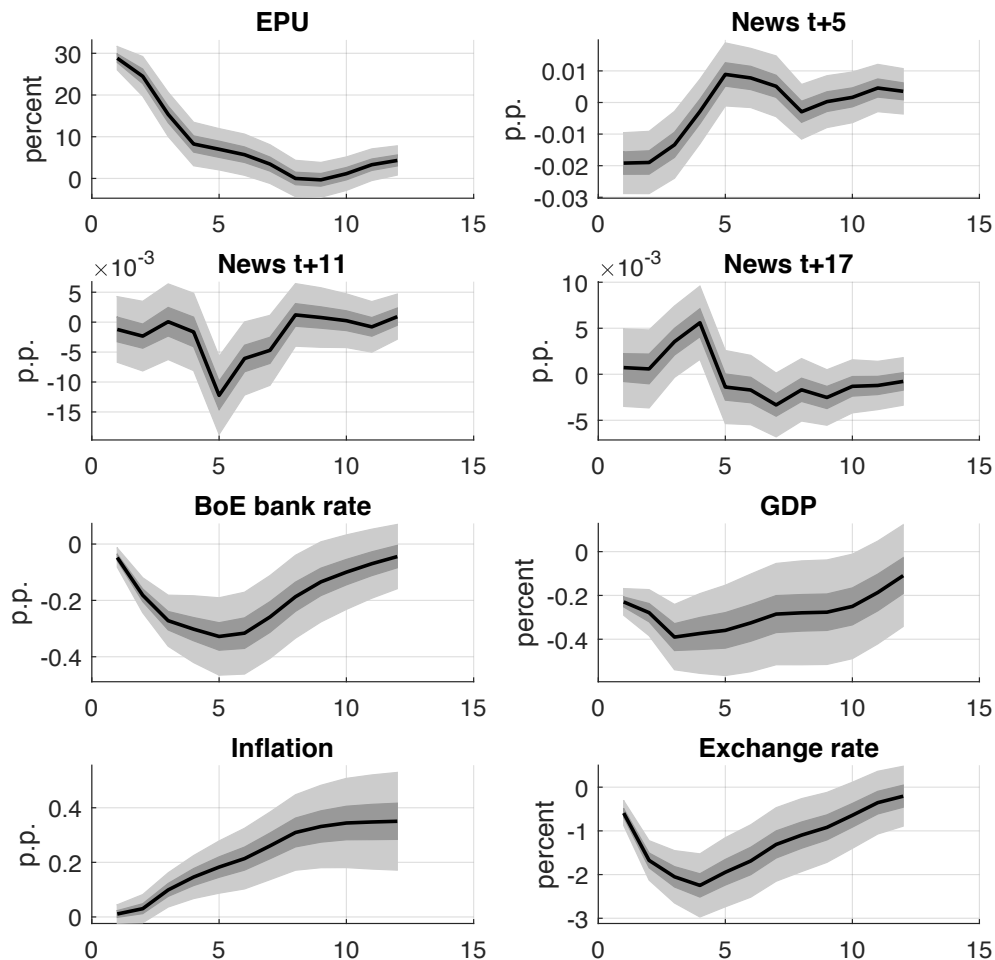


Figure 5: Impulse response functions: economic policy uncertainty. *Note:* impulse response functions to a one standard deviation increase in uncertainty. Light and dark shaded areas indicate one standard deviation and 90 percent bootstrapped confidence bands, respectively.

C.2 Response of the exchange rate to the Brexit vote

In this section, we investigate to what extent the exchange rate depreciation following the Brexit vote was driven by the identified uncertainty and expectation shocks. Towards this end, we use the estimated EVAR to first construct a “no-shock” time path for all the model variables. This is done by simply using the EVAR as a forecasting tool while setting all the shocks (uncertainty, expectations and the other, unidentified, disturbances) to zero from 2016Q3 onwards. The lagged structure of the EVAR then implies particular time-paths for all variables.

In the next step, we use this “no-shock” time path as the baseline to which we compare other time-paths implied by the EVAR depicted in Figure 6. The black line (“data”) simply shows how the exchange rate in the data evolved following the Brexit vote, in comparison to the no-shock scenario. We clearly see the strong depreciation. Second, the dashed red line (“uncertainty shocks”) is the same as the “no-shock” time path, except that we retain the identified uncertainty shock in 2016Q3. This response is very similar to the IRF of uncertainty shown in the previous section. The difference again stems from the endogenous propagation inherent to the EVAR. Finally, the blue line (“uncertainty and anticipation shocks”) retains both the uncertainty and expectation shocks identified in 2016Q3.

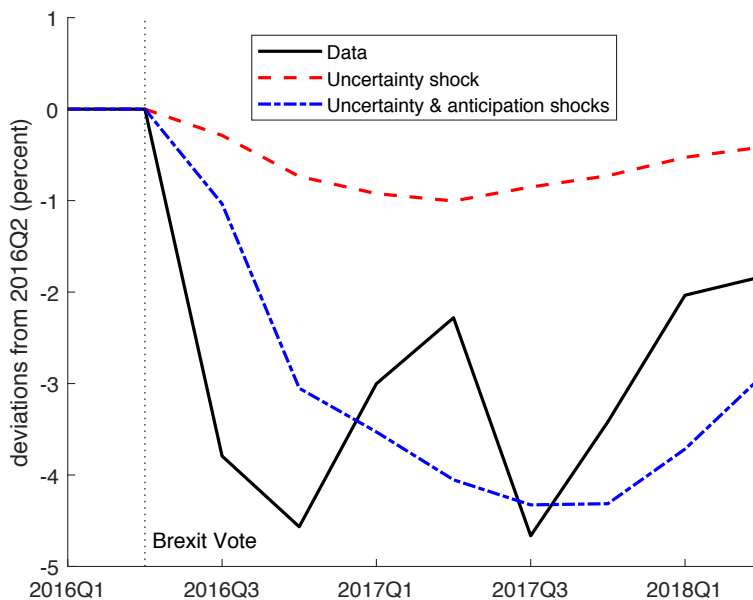


Figure 6: Impulse response functions in 2016Q3. *Note:* impulse response functions to the uncertainty shock (dashed red line) and the uncertainty and anticipation shocks (dash-dotted blue line) identified in 2016Q3. Data is shown in black.

Interestingly, the dash-dotted blue line, capturing the effect that uncertainty and anticipa-

tion shocks had on the exchange rate, is very close to the black line showing the time path of the exchange rate in the data. In other words, the bulk of the exchange rate dynamics following the Brexit vote can be accounted for by the identified uncertainty and anticipation shocks.

C.3 Robustness

We now investigate whether the contributions of uncertainty shocks and growth news shocks to the doppelganger gap depend on the specific VAR setup chosen. The left panel of Figure 7 presents the results for the contribution of uncertainty shocks alone, while the right panel shows the joint contribution of uncertainty and expectation changes. We estimate alternative VARs with one of the following items changed: 6 instead of 4 lags, news horizons of 8, 10, and 12 quarters, investment instead of the 12 quarter ahead news variable, and a linear trend in the VAR. One may also be concerned with the ordering of the variables in the VAR, so we change this as well by ordering the shocks last. We also conduct a check where we replace economic policy uncertainty by a Jurado *et al.* (2015)-style proxy of general macroeconomic uncertainty (available only up until 2017Q2), computed for the UK by Redl (2017). We also exclude either inflation or the exchange rate, order the exchange rate first instead of last and consider the real instead of the nominal exchange rate. The last check we conduct is the dummy EVAR specification described in footnote 30.

Overall, as Figure 7 shows, these changes do not fundamentally change our conclusion that uncertainty can explain at most half of the doppelganger gap and that negative expectations about the future are a major driver of the gap.

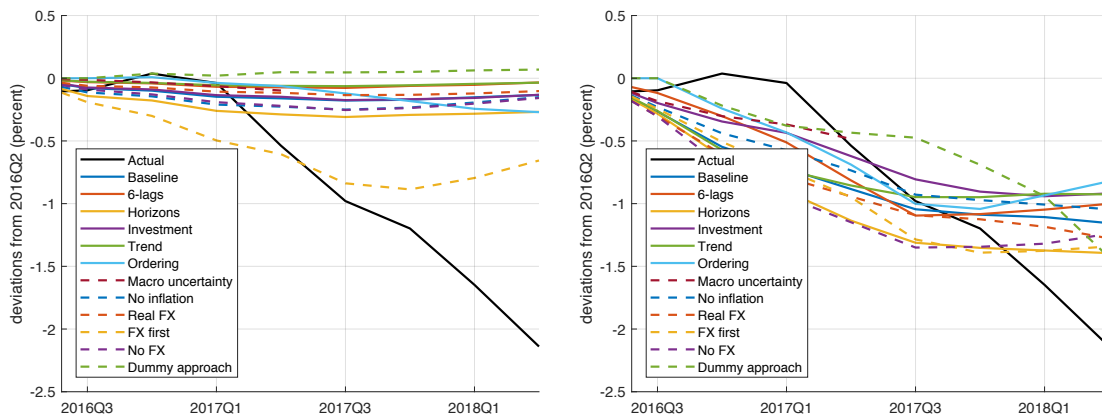


Figure 7: VAR results: robustness. *Note:* doppelganger gaps (difference between doppelganger and data) when uncertainty shock in 2016Q3 is switched off (left panel) and when both uncertainty shock and news shocks in 2016Q3 are switched off (right panel) for alternative specifications.

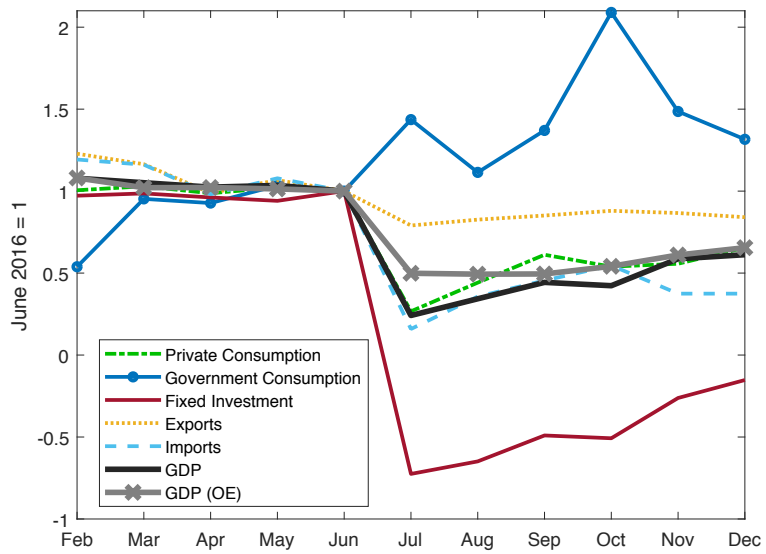


Figure 8: Expectations around the Brexit vote. *Note:* Averages of professional forecasts of 2017 GDP growth and its components published during the months of 2016.

References

- Andrews, D.W.K. (2003). “End-of-sample instability tests”, *Econometrica*, vol. 71(6), pp. 1661–1694.
- Baker, S.R., Bloom, N. and Davis, S.J. (2016). “Measuring economic policy uncertainty”, *Quarterly Journal of Economics*, vol. 131(4), pp. 1593–1636.
- Hahn, J. and Shi, R. (2017). “Synthetic control and inference”, *Econometrics*, vol. 5(4), p. 52.
- Jurado, K., Ludvigson, S.C. and Ng, S. (2015). “Measuring uncertainty”, *American Economic Review*, vol. 105(3), pp. 1177–1216.
- Redl, C. (2017). “The impact of uncertainty shocks in the United Kingdom”, Bank of England Working Paper No. 695.