

WP/18/122

IMF Working Paper

On the Macroeconomic Consequences of Over-Optimism

by Paul Beaudry and Tim Willems

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

IMF Working Paper

Strategy, Policy, and Review Department

On the Macroeconomic Consequences of Over-Optimism**Prepared by Paul Beaudry and Tim Willems**

Authorized for distribution by Mark Flanagan

May 2018

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

Abstract

Is over-optimism about a country's future growth perspective good for an economy, or does over-optimism also come with costs? In this paper we provide evidence that recessions, fiscal problems, as well as Balance of Payment-difficulties are more likely to arise in countries where past growth expectations have been overly optimistic. To examine this question, we look at the medium-run effects of instances of over-optimism or caution in IMF forecasts. To isolate the causal effect of over-optimism we take an instrumental variables approach, where we exploit variation provided by the allocation of IMF Mission Chiefs across countries. As a necessary first step, we document that IMF Mission Chiefs tend to systematically differ in their individual degrees of forecast-optimism or caution. The mechanism that transforms over-optimism into a later recession seems to run through higher debt accumulation, both public and private. Our findings illustrate the potency of unjustified optimism and underline the importance of basing economic forecasts upon realistic medium-term prospects.

JEL-Classification Numbers: E22, E32.

Keywords: noise shocks, sentiment, forecasts, recessions

Author's E-Mail Address: paul.beaudry@ubc.ca and twillems@imf.org.

On the Macroeconomic Consequences of Over-Optimism*

Paul Beaudry[†] and Tim Willems[‡]

May 2018

Abstract

Is over-optimism about a country's future growth perspective good for an economy, or does over-optimism also come with costs? In this paper we provide evidence that recessions, fiscal problems, as well as Balance of Payment-difficulties are more likely to arise in countries where past growth expectations have been overly optimistic. To examine this question, we look at the medium-run effects of instances of over-optimism or caution in IMF forecasts. To isolate the causal effect of over-optimism we take an instrumental variables approach, where we exploit variation provided by the allocation of IMF Mission Chiefs across countries. As a necessary first step, we document that IMF Mission Chiefs tend to systematically differ in their individual degrees of forecast-optimism or caution. The mechanism that transforms over-optimism into a later recession seems to run through higher debt accumulation, both public and private. Our findings illustrate the potency of unjustified optimism and underline the importance of basing economic forecasts upon realistic medium-term prospects.

JEL-classification: E22, E32.

Key words: noise shocks, sentiment, forecasts, recessions.

*We thank Matthieu Bellon, Paolo Cavallino, Mark Flanagan, Klaus Hellwig, Jean-Paul L'Huillier, Galen Sher, as well as seminar participants at the IMF and the University of Toronto for useful comments. Jelle Barkema, Chengyu Huang, and Roxana Pedraglio provided excellent research assistance. The views expressed in this paper are those of the authors and should not be attributed to the International Monetary Fund, its Executive Board, or its management.

[†]University of British Columbia, paul.beaudry@ubc.ca.

[‡]International Monetary Fund, twillems@imf.org.

1 Introduction

Recent years have brought a revival of an old idea by Arthur Pigou (1927), that macroeconomic fluctuations can arise due to difficulties that economic agents encounter in forecasting future economic developments. According to this view, over-optimism about the future may favor a short-term boom, but at the cost of a later reckoning period which often takes the form of a recession.¹ At the other end of the spectrum is the possibility that optimism – even if unfounded – can be good for an economy as it may induce a switch to a higher equilibrium output path in environments with strong complementarities (see e.g. Cooper and John (1988)).

In this paper, we examine these ideas by analyzing the impact of past forecast errors on future macroeconomic outcomes. We base our empirical analysis on forecasts produced by the International Monetary Fund (IMF). These data have broad coverage as they are available for all IMF member countries, enabling a panel data analysis. In addition, many member countries and private agents use IMF forecasts for spending decisions which implies that they can affect macroeconomic outcomes. We however believe that our empirical results are not specific to the IMF, but may apply to any forecast that affects aggregate economic decisions.

Our main finding is that past over-optimism with respect to growth makes recessions, fiscal problems, and Balance of Payment-difficulties more likely. The mechanism seems to run through higher debt accumulation (both public and private), which leads to a build-up of fragilities in the system and harms future growth.

By focusing on the effects of past growth forecast errors on future growth outcomes, our analysis is subject to an endogeneity problem as the dependent and independent variable of our regression specification are mechanically related. To address this issue, an instrumental variables approach is needed. To operationalize this in the IMF context, we first document the existence of IMF Mission Chief fixed effects which we use to identify instances of forecast-optimism/caution. In particular, we show that IMF forecasts for a country tend to become more optimistic (cautious) if that country is assigned a Mission Chief who, on average, produced more optimistic (cautious) forecasts when assigned to other countries. A subsequent instrumental variable regression suggests that there is a causal link from growth over-estimations, to future recessions. Results are robust to the inclusion of an additional instrument (capturing the degree of optimism in the wider IMF

¹Beaudry and Portier (2004) and Blanchard, L’Huillier, and Lorenzoni (2013) provide a formalization of this idea, showing how noisy news about the future can give rise to business-cycle type fluctuations with over-optimistic news giving rise to a delayed recession.

team responsible for the forecast), as well as to the addition of various covariates.

To the best of our knowledge, our empirical analysis on the medium-term macroeconomic consequences of forecast errors is a new contribution to the literature. The spirit of our paper is related to that of Baqir, Ramcharan, and Sahay (2005), who analyze whether the level of ambition embodied in IMF programs affects growth. For a large sample of IMF program cases, they compare outcomes with monetary and fiscal targets set in IMF programs and document that more ambitious fiscal targets tend to improve growth performance, while more ambitious monetary targets typically reduce growth.

Other related contributions include Bachmann and Elstner (2015, who analyze welfare losses stemming from misallocation due to over-optimism/pessimism at the firm-level), Arezki, Ramey, and Sheng (2017, who study the effects of news shocks in open economies using information from oil discoveries), and Oh and Waldman (1990). The latter paper started a literature with a shorter-run focus which analyzes the impact of expectational shocks (as measured by data revisions), with the work by Rodríguez Mora and Schulstad (2007) being a key contribution. The main take-away from this literature is that “false announcements” regarding the U.S. economy (of the type where date- t GDP is announced to equal 100 at date t , but this number being revised down to 99 at $t + 1$) have a positive effect on subsequent U.S. economic activity. Blanchard, Lorenzoni, and L’Huillier (2017) and Enders, Kleemann, and Müller (2017) report similar findings using different methods on a more recent vintage of U.S. data, while Di Bella and Grigoli (2018) confirm this result for a panel of countries.

Putting results of this existing literature and our contribution together points to a narrative where an over-estimation of the future rate of economic growth could provide a short-run boost to the economy, but it also increases the subsequent probability of a recession and other economic difficulties. In this sense, we find that adopting a sufficiently long time horizon when analyzing the effects of optimism and pessimism is key for understanding its overall impact.

2 Data

Our empirical analysis is based upon forecasts made by the IMF’s World Economic Outlook (WEO) publication. Those forecasts are published twice a year (in April and in October) and combine judgement by IMF staff with information from country authorities, economic and statistical models, as well as from other forecasters (see IEO (2014) for a more detailed discussion).

The IMF WEO publication yields one of the few comprehensive forecasting databases, covering all 189 IMF member states. WEO forecasts furthermore play a central role in policy circles, with 88 percent of country authorities (strongly) agreeing with the statement that they “consider the WEO’s projections to be the benchmark for assessing economic prospects” (IEO, 2006). Genberg and Martinez (2014b) moreover report that 64 percent of country authorities (strongly) agreed with the statement that they “use WEO forecasts to check the accuracy of [their] own forecasts”, while 75 percent (strongly) agreed that “WEO forecasts are valuable inputs to the economic policy process in [their] country”. As for example documented in the Annex to Genberg, Martinez, and Salemi (2014), WEO forecasts are also taken into account in projections made by private sector forecasters.

This paper uses the October-vintage of each WEO round (available annually since 1990). Calculating one-year ahead forecast errors (by taking the October WEO vintage of year T and comparing its prediction for real GDP growth in year $T+1$ with the subsequent year $T+1$ realization²), enables us to uncover several stylized facts that apply to the 1990-2016 sample period:³

1. *On average, forecasts for real GDP growth have been overly optimistic.*
2. *The IMF’s WEO has forecasted less recessions⁴ than actually occurred: 456 forecasted recessions, compared with 1,094 actual ones.*
3. *Of the 1,094 actual recessions that did occur over the sample period, the IMF’s WEO forecasted 258 (24 percent) one year ahead.*
4. *Of the 456 recessions forecasted one year ahead, 198 (43 percent) never materialized.*

Over our sample period, WEO forecasts for real GDP growth have shown a tendency towards over-optimism: on average, the prediction for next year’s growth rate has been 0.58 percentage points higher than the subsequent realization. Such an upward bias is also documented in IEO (2014), while Loungani (2001) shows that it is present in private sector forecasts as well. This observation rhymes with Stylized Facts #2 and #3, which suggest that there is a reluctance to predict recessions.⁵ The ratio of the number of actual recessions to one-year ahead predicted ones, stands at 2.4 and nearly 4 in every

²Taken to be the value that is present in the WEO database-vintage three years later. Comparing with more recent vintages makes growth numbers incomparable over time due to definitional changes, statistical revisions, and so forth.

³Similar statements also apply to different forecast horizons, such as three- and five-year ahead ones.

⁴Here defined as any year in which real GDP per capita falls.

⁵For an institution like the IMF, this might stem from a hesitation to “rock the boat”; see IEO (2014).

5 recessions arrived without being forecasted. At the same time, however, the IMF did sometimes predict recessions that subsequently didn't materialize (Stylized Fact #4).

The next section of this paper examines the effects of such forecast errors.

3 Forecast Errors and Recessions

This section contains our main empirical analysis: investigating whether elevated growth expectations affect the probability of a recession occurring (where recessions are defined as episodes of negative real GDP per capita growth).⁶

The most important input to our regressions is the IMF's WEO-implied forecast error for real GDP growth. We construct this forecast error variable as follows. Since many macro-relevant economic decisions (such as those related to investment and debt accumulation) are taken with a multi-year horizon in mind, we look beyond one-year ahead forecasts. In particular, for each year T , we start by averaging the forecasted real GDP growth rate over the years $T+1$, $T+2$, and $T+3$ (thereby thus simultaneously incorporating the one-, two-, and three-year forecast horizon). Next, we calculate the average realized growth rate over those years. The forecast error then follows by subtracting the latter from the former. Hence, for any arbitrary horizon $h > 0$, we can define:

$$F_{h,t} \equiv \frac{1}{h} \sum_{j=1}^h g_{t+j|t}^f - \frac{1}{h} \sum_{j=1}^h g_{t+j}, \quad (1)$$

where $g_{t+j|t}^f$ is the year t forecast for real GDP growth in year $t+j$, while g_{t+j} is the realized growth rate for year $t+j$. We focus on the case where $h = 3$, which strikes a balance between taking a medium-run perspective (reflecting that many macroeconomic decisions are based upon expectations for multi-year horizons) and not losing too many observations at the beginning of our sample (since we don't have pre-1990 forecasts, we can only start calculating forecast errors from the year $1990+h$ onwards). Our findings are however robust to using other horizons (noting that we cannot go higher than $h = 5$ since the IMF WEO does not contain forecasts beyond the five-year horizon). Figure 1 depicts the distribution of our main forecast error variable $F_{3,t}$, confirming the aforementioned bias towards growth over-estimation (the median forecast error equals 0.47).⁷

⁶Results are robust to using a different cut-off – such as real per capita growth below 1 or 2 percent.

⁷We drop observations that are more than two standard deviations (8.7 percentage points) from the mean since analysis of such large forecast errors shows that they often have reasons beyond economics (statistical revisions, natural disasters, wars, etc.). This shrinks our sample by about 3 percent.

[Insert Figure 1]

We begin by examining what impact a forecast error stemming from year T (for years $T + 1, T + 2, \dots, T + h$) has on the prevalence of recessions in year $T + h$. In this case, part of the forecast error is determined contemporaneously with the actual growth outcome.⁸ This mechanical link between the dependent and independent variable of our regression gives rise to a simultaneity problem. To enable causal inference regarding the effect of over-optimism on later recessions, we adopt an instrumental variable (“IV”) procedure. With this objective in mind, we start by building an instrument which exploits the possibility that certain IMF Mission Chiefs consistently display too much optimism or caution in their forecasts.⁹

Our sample includes all 189 nations that are a member of the IMF and the origin of included covariates is described in Appendix A. Since IMF forecasts carry a lower weight in advanced economies, we often report results for both our global sample as well as for a restricted sample, which features only emerging and developing economies.

3.1 Results When Exploiting Mission Chief Fixed Effects

As discussed above, endogeneity is a challenge in regressions where the dependent variable is the recession dummy and the explanatory variable is a forecast error (especially since forecasters often fail to predict recessions – introducing a mechanical tendency for positive forecast errors and recessions to coincide). To address this concern, we take an instrumental variables approach.

IMF country teams are led by a Mission Chief (henceforth “MC”). Together with other members of the team, the MC determines the growth forecast that is submitted to the IMF’s WEO exercise. Forecasting styles differ across individuals, both in terms of methods applied (ranging from econometric techniques to the use of judgement; IEO (2014)) as well as underlying philosophy: some individuals tend to have a more pessimistic outlook, while others will have a more optimistic outlook. Yet others might see strategic reasons to be overly rosy or cautious in their forecasts.

As part of the IMF’s policies, MCs are frequently rotated to different country assignments. While this re-allocation process has exogenous, random elements – such as MCs

⁸For example: when $h = 3$, the analysis looks at the impact of 1997 forecast errors (for the years 1998, 1999, and 2000) on the probability of a recession occurring in the year 2000. In this case, both the forecast error variable and the dependent variable use outcomes from the year 2000.

⁹In Appendix B we outline the type of setup in which our instrumental variable strategy is able to identify whether over-optimism causes eventual downturns in economic activity.

asking for re-assignment due to personal reasons – it is not completely arbitrary.¹⁰ To serve as a valid instrument in our paper, it is important that the MC allocation process does not assign more optimistic MCs to countries that are more likely to end up in difficulties. Apart from the fact that we are not aware of instances where optimism-considerations have affected MC allocations, it seems that we can rule out this scenario on informational grounds: it not only requires IMF decision makers to have detailed knowledge about MC-specific degrees of optimism/caution, but also about the recession-proneness of countries (while reality is that recessions are difficult to predict, as for example documented in Section 2 and the references therein). Consequently, we proceed by assuming:

ASSUMPTION 1. There is no systematic relationship in the IMF’s MC allocation process between a MC’s degree of forecast optimism/caution and economic prospects for a country.

In addition, we require that the MC’s degree of over-optimism does not have a direct effect on economic outcomes (beyond the forecast-channel). This exclusion restriction seems plausible, as it is hard to imagine that the personality of a MC has a direct effect on macroeconomic developments in a country.

In this section, we use MC re-allocations to estimate the degree of optimism for each and every MC in our sample. In order to implement this idea, we have constructed a new, comprehensive dataset on country assignments for IMF MCs since 1990. We managed to create our dataset using IMF staff reports (mainly Article IV documents): they indicate the name of the MC, typically on the first page of the Introduction. Given the large data volume, this information was extracted using text mining (exploiting the fact that the MC is always introduced by a standard sentence like “The mission was led by...”).

Hereby, we have been able to compile a dataset which identifies the IMF MC for about 80% of the country-year pairs since 1990. Our database contains 705 unique MCs, who have led IMF teams for 2.7 different countries on average; 475 MCs have led teams for multiple countries (which is necessary in order to identify MC fixed effects, as we will see later on). On average, the latter have led IMF teams for 3.5 different countries.

We estimate MC fixed effects using the approach of Bertrand and Schoar (2003, where they identify manager fixed effects in a corporate context) and start by estimating:

$$F_{1,it} = \alpha_t + \gamma_i + \lambda W_{it} + \mu^{(k)} + \epsilon_{it}. \quad (2)$$

Here, $F_{1,it}$ denotes the one-year ahead forecast error on real GDP growth (forecast

¹⁰For example, high-profile program countries tend to be covered by more experienced MCs.

minus realization), made for country i in year t , α_t are the year fixed effects, γ_i are country fixed effects, W_{it} is a vector of time-varying country level covariates, and ϵ_{it} is the error term. The main object of interest in (2) is $\mu^{(k)}$, MC k 's fixed effect (where MCs of an overly-optimistic nature will have $\mu > 0$, while MCs who tend to produce cautious forecasts will have $\mu < 0$). Observe from equation (2) that we are only able to estimate MC fixed effects for MCs that have led teams for multiple countries.

We subsequently use the MC fixed effects μ as an instrument for forecast errors in a regression where the LHS variable y_{it} is a dummy which takes the value 1 if country i is in recession in year t .¹¹ Using X_{it} to denote a vector of time-varying country level covariates, we thus estimate first-stage regression (3) and second-stage regression (4):

$$F_{3,it-3}^{(k)} = a_t + g_i + b\mu^{(k)} + cX_{it} + u_{it}, \quad (3)$$

$$y_{it} = \alpha_t + \gamma_i + \beta\hat{F}_{3,it-3}^{(k)} + \zeta X_{it} + \varepsilon_{it}. \quad (4)$$

To further mitigate any mechanical relation that could generate problems, we follow a “jackknife”-approach (Angrist, Imbens, and Krueger, 1999) by inferring MC fixed effects from the MC’s forecasts for countries *other than the country to which the fixed effect will be applied in regression (3)*. So if we consider a MC who has headed IMF teams for Albania, Rwanda, and Thailand, the fixed effect that is used as an instrument in regression (3) for Rwandan forecast errors, is only based upon the MC’s forecasts for Albania and Thailand. Similarly, the fixed effect used in (3) to instrument for Albanian forecast errors, is only based upon the MC’s forecasts for Rwanda and Thailand. In this way, no information from a country’s own history (or future) is applied back to that very same country – minimizing endogeneity of the instrument with respect to country outcomes.

When estimating MC fixed effects through equation (2), we include two covariates in the vector W : (i) a dummy variable equaling 1 if country i was under an IMF program in year t (as it has been argued that IMF forecasts are particularly optimistic in IMF programs; GAO (2003)) and (ii) PPP-adjusted real GDP per capita (as the level of development might play a role; Genberg and Martinez (2014a)).

Table 1 shows results for the regression used to estimate MC fixed effects. The average estimated MC fixed effect amounts to 0.34, while its standard deviation equals 2.6 (Figure 2 shows μ 's distribution¹²).

¹¹Appendix C shows that our results continue to hold if one takes y_{it} to be the growth rate of real GDP (suggesting the existence of a more general negative impact of over-optimism on future growth).

¹²In line with footnote 7, we drop estimates that are more than two standard deviations apart from the mean.

[Insert Figure 2]

A joint F -test that all MC fixed effects are equal to 0, is rejected at the 1% significance level (and inclusion of the MC fixed effects increases the adjusted R^2 from 0.06 to 0.13). Results furthermore suggest that, once country and time fixed effects are accounted for, there is no strong evidence for excessive optimism during IMF programs – a result that is in line with the recent study of Luna (2014). The level of real GDP per capita does not seem to play a major role either.¹³

Table 1: OLS regression (2), used to estimate Mission Chief fixed effects

Dependent variable: one-year ahead forecast error $F_{1,it}$		
	(1: without MC fixed effects)	(2: with MC fixed effects)
IMF program dummy	0.148049 (0.46)	0.3606151 (1.26)
real GDP per capita	0.00000454 (0.15)	-0.0000414 (-1.17)
Mission Chief fixed effects	no	yes
time and country fixed effects	yes	yes
countries	171	170
obs	4,251	3,318

Note: t-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level.

Table 2b contains results when the MC fixed effects from (2) are used to instrument forecast errors. The underlying first-stage regressions are displayed in Table 2a. They suggest that our instrument is relevant.¹⁴ In particular, the significantly positive estimates for the coefficients on μ are a strong finding since these fixed effects were obtained through a jackknife-procedure, solely relying upon forecasts for “other countries” (recall the discussion under equation (4)). So in a sense, this is an out-of-sample forecasting

¹³When country fixed effects are dropped the coefficient on the IMF program dummy turns significantly positive (0.43, p-value of 0.03), while that for real GDP per capita takes on a significantly negative value (-0.000015, p-value of 0.003) – pointing to larger forecast errors in less-developed countries.

¹⁴The Stock-Yogo (2005) 20% (15%) maximal bias critical value for the Cragg-Donald statistic equals 6.66 (8.96). These hurdles are consistently passed in our full sample (as is the rule-of-thumb of 10), even when using heteroskedasticity-robust Kleibergen-Paap statistics (which are slightly lower than Cragg-Donald statistics). The critical values are not always passed in the smaller sample of emerging and developing countries – pointing to a weaker instrument for these cases. Findings are however similar for both samples, so our conclusions continue to apply when basing inference solely upon the full sample.

exercise – showing that MCs who have been overly optimistic when assigned to countries other than A, also tend to produce overly-optimistic forecasts when assigned to country A. This is evidence in favor of the existence of “systematically optimistic/cautious” MCs.

Turning back to the question central to our paper, the second-stage results in Table 2b suggest that over-optimism *increases* the probability of a recession occurring 3 years later. This finding is robust to the inclusion of various covariates that are believed to affect the incidence of recessions (with endogenous variables being lagged by one period, see Columns (2) and (4)). It furthermore shows up in both our global sample (Columns (1) and (2)), as well as in the restricted sample featuring only emerging and developing economies (Columns (3) and (4)). Although not reported in Table 2, we find that the result no longer holds when only focusing on the subsample of advanced countries (in that case we obtain a z-statistic of 0.78).¹⁵ This is consistent with IMF forecasts (and any variation resulting from IMF MC rotations) playing a less important role in economies for which alternative forecasts are readily available.

3.2 Incorporating Team Sentiment

In addition to the MC fixed effects, we can rely on another instrument – exploiting the notion that bullish forecasts typically tend to be overly optimistic *across all forecast horizons*. This observation may be related to team dynamics: although the MC is a powerful force when it comes to determining the growth forecast, he or she by no means acts in isolation. A MC heads a wider IMF team which covers the country on a day-to-day basis for the institution and all team members are in principle able to affect the growth forecast through analytical input and team discussions. Within team settings in general, certain group dynamics may arise: being exposed to similar information and ideas, teams can become overly “bullish” about a country’s economic prospects, while an overly “bearish” dynamic is possible as well.

This enables us to build another instrument through which we can sever the mechanical link between our dependent and independent variable: we can take the one-year ahead forecast error and use it as an additional instrument for the average forecast error *at the two- and three-year horizon*.

¹⁵The same goes for reduced form estimations (which regress y directly on μ): they detect a significant impact in the full and eme/dev sample, but not in the advanced economy sample. As explained in Nunn and Wantchekon (2011), this type of falsification test supports the exclusion restriction (as it detects no direct impact from MC fixed effects on outcomes in countries where the forecast-channel matters less).

Table 2a: first-stage regressions

Dependent variable: forecast error on growth " $F_{3,t-3}$ "				
	(1: full)	(2: full)	(3: eme/dev)	(4: eme/dev)
Mission Chief fixed effect " μ "	0.1000104*** (3.57)	0.1067803*** (3.30)	0.0802024** (2.56)	0.092637** (2.52)
public debt/GDP (t-1)		-0.0011443 (-0.63)		-0.002104 (-1.10)
inflation (t-1)		0.0058827*** (3.69)		0.0053907*** (3.12)
avg growth in trading partners		-0.1779058*** (-3.26)		-0.1506727*** (-3.00)
pct change in terms-of-trade		-0.0025733 (-0.60)		-0.0017607 (-0.40)
time and country fixed effects	yes	yes	yes	yes
Cragg-Donald statistic	13.62	12.77	6.65	7.04
countries	182	165	147	131
obs	2,694	2,237	2,095	1,704

Table 2b: instrumental variable regressions

Dependent variable: recession dummy (1 = presence of recession)				
	(1: full)	(2: full)	(3: eme/dev)	(4: eme/dev)
forecast error on growth " $F_{3,t-3}$ "	0.1712131*** (3.03)	0.1389873** (2.51)	0.1754304** (2.24)	0.131479* (1.91)
public debt/GDP (t-1)		0.003652 (1.39)		0.0003613 (1.21)
inflation (t-1)		-0.0012023*** (-2.82)		-0.001184** (-2.57)
avg growth in trading partners		0.0089896 (0.76)		0.0122687 (1.02)
pct change in terms-of-trade		-0.0002605 (-0.32)		-0.0004421 (-0.53)
time and country fixed effects	yes	yes	yes	yes
countries	182	149	147	131
obs	2,694	2,282	2,095	1,704

Note: t (z)-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level.

The idea is that a rosy one-year ahead forecast indicates an optimistic IMF team and that this translates into optimistic forecasts at longer horizons too. Importantly, we exclude the one-year horizon from the forecast error variable that is instrumented – ensuring

that there is a separation between instrumental and instrumented variable (otherwise the mechanical link between our regression’s LHS and RHS would remain intact).¹⁶ The forecast error is now being calculated as $F(2)_{3,t} \equiv \frac{1}{2} \sum_{j=2}^3 g_{t+j|t}^f - \frac{1}{2} \sum_{j=2}^3 g_{t+j}$, the difference with equation (1) being that the summation here only starts at $j = 2$ (instead of $j = 1$).

At this stage, one might worry about the exclusion restriction. Along these lines, the earlier literature started by Oh and Waldman (1990) has shown that the isolated effect of short-term growth over-optimism is to stimulate the economy somewhat. Consequently, a possible violation of the exclusion restriction is of secondary concern, as it would only bias estimates *against* our finding that over-optimism increases the risk of a later recession.

Table 3 contains results when including this instrument next to the MC fixed effects. We continue to find that over-optimism increases the odds of a future recession occurring. As before, this result emerges in both samples and is robust to the inclusion of covariates. In this case, our estimation is overidentified and we can test whether our instruments are valid. As shown in Table 3b, Hansen’s J-statistic suggests that our instruments are uncorrelated with the error term – building further trust in the validity of our instruments.

Resulting estimates suggest sizeable effects from forecast errors: over-estimating average growth over the two/three-year horizon by 1 percentage point (with the over-estimation of growth being larger than this in 40 percent of all country-year pairs in our sample), increases the probability with which a recession arrives 3 years later by about 8 percentage points. Given that the unconditional recession probability in our sample is 25.2 percent, this reflects a 30 percent increase from the mean.

At the same time, our results suggest that cautious forecasts can help to prevent recessions. Frankel (2011) highlights the example of Chile in the late 2000s: he argues that Chile was able to moderate the Global Financial Crisis-induced downturn because they had been basing their budget decisions upon a conservative estimate for the copper price (an important Chilean export product). This enabled the Chilean government to implement countercyclical fiscal policy when external conditions deteriorated.

In Appendix D we show that our findings are robust to several perturbations. In particular, we report results based upon logit and OLS specifications, where the endogeneity problem is mitigated by lagging the forecast error variable (as to break the mechanical link between the dependent and independent variable in our regression).

¹⁶The mechanical link would still not be severed if economies systematically started losing momentum in years prior to a recession. However, as detailed in Appendix D, this does not seem to be the case: most recessions arrive rather suddenly – a point also made by Acemoglu and Scott (1997).

Table 3a: first-stage regressions

Dependent variable: forecast error on growth " $F(2)_{3,t-3}$ "				
	(1: full)	(2: full)	(3: eme/dev)	(4: eme/dev)
Mission Chief fixed effect " μ "	0.0766581** (2.31)	0.0795566** (2.11)	0.0482091 (1.31)	0.0553565 (1.30)
One-year ahead forecast error $F_{1,t-3}$	0.0851782*** (3.41)	0.0965994*** (3.66)	0.0864669*** (3.16)	0.0970479*** (3.33)
public debt/GDP (t-1)		-0.0008716 (-0.44)		-0.0013472 (-0.64)
inflation (t-1)		0.0058296 (0.94)		0.0037463 (0.59)
avg growth in trading partners		-0.211631*** (-3.11)		-0.16392*** (-2.85)
pct change in terms-of-trade		-0.0086653 (-1.57)		-0.0086058 (-1.51)
time and country fixed effects	yes	yes	yes	yes
Cragg-Donald statistic	13.42	13.86	9.46	9.67
countries	181	165	146	131
obs	2,555	2,132	1,978	1,617

Table 3b: instrumental variable regressions

Dependent variable: recession dummy (1 = presence of recession)				
	(1: full)	(2: full)	(3: eme/dev)	(4: eme/dev)
forecast error on growth " $F(2)_{3,t-3}$ "	0.0989018*** (3.20)	0.0872071*** (2.86)	0.0823811** (2.44)	0.0734848** (2.20)
public debt/GDP (t-1)		0.0004917** (2.27)		0.0003918* (1.71)
inflation (t-1)		-0.001161 (-1.59)		-0.0012383* (-1.87)
avg growth in trading partners		0.0027508 (0.28)		0.004145 (0.45)
pct change in terms-of-trade		0.0003208 (0.39)		0.0001141 (0.14)
time and country fixed effects	yes	yes	yes	yes
Hansen J-statistic (p-value)	4.349 (0.0370)	1.846 (0.1742)	3.424 (0.0643)	1.509 (0.2194)
countries	181	165	146	148
obs	2,555	2,132	1,978	2,220

Note: t (z)-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level. The instrument set includes both MC fixed effects and one-year ahead forecast errors.

Finally, one might wonder just how costly the recession actually is. A future recession does not necessarily harm GDP in the longer run if growth during the interim years was very high. To analyze this, we construct a “depression dummy” which we define as a variable that equals 1 if real GDP per capita in year T , is lower than real GDP per capita in year $(T - 4)$. As Table 4 shows, our IV estimates suggest that positive forecast errors make sustained drops in real income more likely – both in our global sample as well as in the restricted sample consisting of emerging and developing countries only.

Table 4: instrumental variable regressions

Dependent variable: depression dummy (1 = presence of depression)				
	(1: full)	(2: full)	(3: eme/dev)	(4: eme/dev)
forecast error on growth “ $F(2)_{3,t-3}$ ”	0.2693736*** (4.29)	0.2335627*** (4.39)	0.2733438*** (3.66)	0.230579*** (3.81)
public debt/GDP (t-1)		0.0017053*** (4.64)		0.0012623*** (3.04)
inflation (t-1)		-0.0004395 (-0.27)		0.000063 (0.04)
avg growth in trading partners		0.0450846*** (2.71)		0.0351513** (2.47)
pct change in terms-of-trade		0.0026219** (1.96)		0.0024401* (1.75)
time and country fixed effects	yes	yes	yes	yes
Hansen J-statistic (p-value)	2.164 (0.1413)	3.469 (0.0625)	1.157 (0.2822)	2.283 (0.1308)
countries	181	165	146	131
obs	2,548	2,126	1,972	1,612

Note: z-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level. The instrument set includes both MC fixed effects and one-year ahead forecast errors.

4 Forecast Errors and Crises

In addition to analyzing the impact of past forecast errors on recessions, it is also interesting to see whether they have had an effect on the occurrence of crises of various sorts. Besides the fact that these regressions are informative in their own right, they also serve as a further robustness check given that the LHS variable in these regressions is no longer mechanically related to the growth-forecast error variable on the RHS (reducing the challenge faced by our IV-strategy).¹⁷

¹⁷In Appendix D we show that these results are robust to OLS and logit estimations.

Along these lines, Table 5 shows that past over-optimism makes the arrival of fiscal crises more likely.¹⁸ Column (1) illustrates this for the baseline regressions (featuring only time and country fixed effects), while Column (2) shows that this finding is robust to the inclusion of our standard covariates. Columns (3) and (4) suggest that past over-optimism also increases the incidence of IMF programs. As the presence of the latter indicates a Balance of Payments (“BoP”) need, this suggests that elevated expectations also increase the likelihood of a BoP crisis occurring in the future.

Table 5: instrumental variable regressions

Dependent variable: crisis dummy (1 = presence of crisis)				
	(1: fiscal crisis)	(2: fiscal crisis)	(3: BoP crisis)	(4: BoP crisis)
forecast error on growth “ $F(2)_{3,t-3}$ ”	0.0890415** (2.36)	0.0806261** (2.36)	0.0904298*** (2.69)	0.0852909** (2.53)
public debt/GDP (t-1)		0.0019071*** (6.19)		0.0013382*** (4.62)
inflation (t-1)		0.0016078* (1.87)		-0.0004021 (-0.49)
avg growth in trading partners		0.0147679 (1.43)		0.0153645 (1.47)
pct change in terms-of-trade		0.0008022 (0.85)		0.0019667** (2.33)
time and country fixed effects	yes	yes	yes	yes
Hansen J-statistic (p-value)	0.947 (0.3305)	5.252 (0.0219)	0.666 (0.4144)	0.660 (0.4166)
countries	181	165	182	166
obs	2,411	2,134	2,566	2,137

Note: z-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level. The instrument set includes both MC fixed effects and one-year ahead forecast errors.

5 Mechanism

The previous sections provided evidence that overly-optimistic growth expectations make countries more prone to ending up in a recession, a fiscal crisis, as well as a BoP crisis. This begs the question what drives this result.

To obtain a better understanding, it is insightful to analyze the dynamics of certain variables of interest in response to “forecast shocks”. We do this by following Jorda (2005) in generating Impulse-Response Functions (“IRFs”) through Local Projections

¹⁸The fiscal crisis dummies (capturing sovereign defaults, episodes of public arrears accumulation, etc.) are taken from Gerling et al. (2017).

(“LPs”). Relative to VAR-methods, LPs have the advantage that they are more robust to misspecification of the data-generating process. Whereas VAR-based IRFs are generated iteratively (thereby compounding misspecification error as the time-horizon increases), LPs proceed by directly making multi-step forecasts.

We follow Jorda, Schularick, and Taylor (2013: 17-18) in extending the LP-IRF method to a panel data specification and in first-differencing our LHS variables (as all variables that are to be analyzed are potentially non-stationary). In particular, we start by estimating a series of regressions for each horizon j :¹⁹

$$\Delta y_{i,t+j} = \alpha_t^{(j)} + \gamma_i^{(j)} + \beta^{(j)} \bar{g}_{it}^f + \rho_1^{(j)} \Delta y_{i,t-1} + \rho_2^{(j)} \Delta y_{i,t-2} + u_{i,t+j}, \quad (5)$$

where $\Delta y_{i,t+j}$ is the first-difference of a variable y in country i at date $t + j$, and \bar{g}_{it}^f is the three-year-ahead average forecast for the rate of real economic growth in country i , made in year t (so the forecast averages over years $t + 1$, $t + 2$, and $t + 3$ with $\bar{g}_{it}^f \equiv (g_{t+1|t}^f + g_{t+2|t}^f + g_{t+3|t}^f)/3$); α_t (γ_i) is the time (country) fixed effect, and $u_{i,t+j}$ is the error term. By subsequently tracing out $\beta^{(j)}$ for $j = 1, \dots, N$ we obtain an estimate for the response of our LHS variable of interest ($\Delta y_{i,t+j}$) to an impulse in the growth forecast.

Note that the nature of the exercise has changed, as the key RHS variable is now a forecast as opposed to a forecast *error*. As a result, the dependent variable of the regression no longer mechanically correlates with the RHS. This has certain advantages (it for example removes the challenge addressed by our IV-strategy to isolate the effects of over-optimism), but also introduces the potential problem of interpreting results which now contain the effects of *both founded optimism and over-optimism* about the future.²⁰ However, we will argue below that this exercise will serve as a conservative robustness check to our previous finding that over-optimism causes later recessions.

All results that are to follow are robust to a slightly different specification where the RHS features $\Delta y_{i,t}$ and $\Delta y_{i,t-1}$ (instead of $\Delta y_{i,t-1}$ and $\Delta y_{i,t-2}$).²¹ In all cases, standard errors are calculated via Newey-West. As Jorda (2005) has shown that the disturbance term has an moving average component of order h (the forecast horizon), we set the

¹⁹As pointed out in Nickell (1981), dynamic panel data estimations with fixed effects are inconsistent if the panel has a short time dimension, T . Given that the resulting bias will be of order $(1/T)$, where in our case $T = 27$, we follow the recommendation of Judson and Owen (1999) and proceed by employing the standard fixed-effects estimator (also see Arezki, Ramey, and Sheng (2017) on this).

²⁰In the context of this paper, “founded optimism” (or pessimism) is defined as a predicted acceleration (deceleration) in growth which, on average, materializes; “over-optimism” is that part of the forecast which, on average, does not materialize (if it is negative, the forecast was overly cautious).

²¹Equation (5) assumes that the forecast \bar{g}_{it}^f may be affecting the contemporaneous growth rate $\Delta y_{i,t}$, while the alternative specification (which controls for $\Delta y_{i,t}$ on the RHS) assumes that the forecast is not affecting growth contemporaneously.

truncation lag for the Newey-West estimator equal to h in each regression.

Let us start by taking y to be the logarithm of real GDP, so that $\Delta y_{i,t+j}$ represents real GDP growth during year $t + j$. As Figure 3 shows, a forecast shock leads to a short-run boost in economic growth. In this paper, however, we are particularly interested in the response at longer horizons. There, we see that the stimulating effect of a forecast shock reverses as time goes by: the boost lasts for only two years, while the forecast shock actually tends to *reduce* growth from four years after the shock onwards. This result holds for both our full sample of countries, as well as for the restricted sample consisting of emerging and developing countries only. This pattern is consistent with our earlier findings in Sections 3 and 4, which indicated that past over-optimism tends to bring later recessions (and even crises of various sorts). Those sections were however aimed at isolating the effect of *over*-optimism, while Figure 3 shows that the negative effect at longer horizons continues to emerge even if one does not remove the founded optimism component. As retaining founded optimism (which could for example consist of knowledge that future oil production will rise due to the activation of a new field) stacks the cards *against* finding a negative impact on growth, this exercise can be interpreted as a robustness check of our earlier results. The observation that Figure 3 nevertheless continues to display a negative effect on growth at longer horizons, can thus be seen as a corroborating result.

[Insert Figure 3]

Next, we consider what happens to the government debt-to-GDP ratio when growth forecasts become rosier.²² In particular, let y represent the debt-to-GDP ratio (and $\Delta y_{i,t+j}$ its change in percentage points during year $t + j$). Here, we fail to find a significant effect on impact, but from the second year onwards we see that a forecast shock tends to lead to faster accumulation of public debt. Again, this result emerges in both samples (see Figure 4). It provides support for a popular narrative where optimism about future growth induces countries to borrow against a prospective income stream – creating debt vulnerabilities and potentially endangering the country’s future growth potential (given the evidence that public debt may hamper long-run growth; cf. Reinhart and Rogoff (2010) and Eberhardt and Presbitero (2015)).

²²Again, we follow a conservative approach by focusing on the combined response to founded optimism and over-optimism: since GDP at longer horizons would be lower following pure over-optimism (relative to the case where founded optimism is present as well; cf. the definitions in footnote 20), the depicted responses of the debt-to-GDP ratio are likely to be biased towards zero.

[Insert Figure 4]

Similarly, private credit is found to accumulate in response to a forecast shock, as evidenced by Figure 5 (where y is the private-credit-to-GDP ratio, and $\Delta y_{i,t+j}$ its change in percentage points during year $t + j$).²³ This response is notably quicker than that of government debt, which sits well with the notion that the public sector is relatively sluggish in responding to economic developments due to delays in the political process. Interestingly, the response of the private-credit-to-GDP ratio is much weaker in our restricted sample of emerging and developing countries. As private agents in less developed countries have weaker access to credit, this gives credence to the approach.

[Insert Figure 5]

Finally, both the global and emerging/developing sample suggest that the investment-to-GDP ratio falls after a forecast shock (Figure 6). While this does not necessarily mean that the level of investment falls, it does suggest that investment falls *relative to GDP* (which tends to rise following a forecast shock, recall Figure 3). Agents seem to celebrate positive news about the future by mainly increasing consumption relative to GDP.

[Insert Figure 6]

Together, these results point to a narrative that seems both plausible and familiar: in response to a rosier forecast about the future, both the public and private sector respond by accumulating more debt. At the same time, agents seem to favor consumption over investment. These developments lead to a build-up of fragilities in the system – setting the stage for recessions to occur and increasing the incidence of fiscal- and BoP-crises.

6 Conclusion

Using a broad sample of 189 countries, this paper has documented that recessions, fiscal problems, as well as Balance of Payments-difficulties are more likely to arise in economies

²³Here, proceeding by looking at private sector responses to a forecast shock (as opposed to a forecast *error* shock) assumes that the relevant agents are not able to distinguish between founded optimism and over-optimism. If the private sector were to have superior information relative to official sector forecasts, the response to over-optimism might be weaker than Figure 5 indicates. But given that private sector forecasts correlate highly with, and display similar biases to, forecasts produced by governments and the IMF (cf. Loungani (2001)), there is no strong evidence that this is the case. In addition, the downward bias described in footnote 22 remains present.

for which past growth forecasts have been overly optimistic. We arrived at this conclusion by taking an instrumental variables approach, which exploited the observation that IMF Mission Chiefs appear to differ in their degree of forecast optimism/caution. The fact that Mission Chiefs are periodically re-assigned to different countries in a pseudo-random fashion, gives us quasi-experimental evidence on the effect of varying the degree of optimism/caution at the country level. Even though Mission Chief-induced variation in forecast optimism/pessimism is small (as are its consequences), we show that it can be used as an effective instrument to analyze the causal impact of actual forecast errors (not just the minor part related to IMF Mission Chiefs).

Our finding that over-optimism brings economic damage in later years is highly robust: to the addition of controls, to different estimation methods (IV, OLS, probit, logit, Local Projections), to different dependent variables (growth rates, recession dummies, crisis dummies), as well as subsamples (with our results continuing to hold for the subsample containing only emerging and developing economies, where IMF forecasts carry the largest weight).²⁴

The mechanism which transforms over-optimism into future crises seems to run through higher debt accumulation: both the public and private sector seem to “celebrate” positive news about the future by borrowing more. If the expected rise in income subsequently fails to materialize, the amount of debt accumulated turns out to be excessive and negative dynamics set in.²⁵

Our results illustrate the potency of (non-materializing) optimism shocks and underline the importance of basing policy upon realistic (or even cautious) medium-term macroeconomic forecasts. Specifically, our finding regarding the impact of over-optimism on the incidence of future recessions provides support for existing models in the news/noise-tradition, but we are not aware of contributions which model the particular transmission channel that our results point at. Developing such a theoretical model could be an important avenue for future research.

²⁴Of note, we fail to detect a causal impact from forecast errors to future recessions when the sample is restricted to advanced economies only. This comforts us, as it is hard to imagine that IMF forecasts matter enough to materially affect an advanced economy’s course. Things are different for less developed countries, where the IMF is often a rare provider of external forecasts.

²⁵In line with this narrative, Zhou Xiaochuan (Governor of the People’s Bank of China) stated around the 2017 Communist Party Congress that “if we are too optimistic when things go smoothly, tensions build up, which could lead to a sharp correction, what we call a ‘Minsky Moment’. That’s what we should particularly defend against.”

7 Appendix A: origin of covariates

Table A1: origin of covariates

Variable	Origin	Remarks
forecast error on growth	IMF WEO	Constructed as described in Section 3.2
public debt/GDP	IMF FAD	
inflation	IMF WEO	CPI-based
govt spending/GDP	IMF WEO	
current account balance/GDP	IMF WEO	
FDI/GDP	IMF WEO	
population growth	Penn World Table	
ICRG civil war	ICRG	Integer ranging from 0 (high risk) to 4 (low risk)
exchange rate overvaluation	IMF WEO	Residuals of panel regression of price level on real GDP per capita (at PPP) and time dummies
investment/GDP	Penn World Table	
avg growth in trading partners	IMF WEO	Trade-weighted real GDP growth of partners
pct change in terms-of-trade	IMF WEO	

8 Appendix B: identifying the effects of unfounded optimism (noise)

When can one identify the effect of a boost of optimism on the economy? To tackle this problem, it is useful to consider a framework similar to that used in the news-noise literature, where we interpret optimism as a positive, but false, signal about the economy's future fundamentals. To clarify ideas, consider the two following simple environments. They are chosen so that one setup has optimism being unambiguously good as it only creates a boom, while in the other optimism also carries costs as it generates a recession after a temporary boom. The relevant question is whether our econometric approach would correctly identify which is which.

Environment (1) is a situation in which the *level* of GDP (possibly around a trend) is stationary and determined by:

$$y_t = \alpha E_t y_{t+1} + Z_t, \quad 0 < \alpha < 1$$

Environment (2) is a situation in which the *growth rate* of GDP is stationary:

$$g_t = \alpha E_t g_t + Z_t, \quad 0 < \alpha < 1$$

Here, y_t represents a country's (log) GDP and g_t represents its growth rate ($g_t \equiv y_t - y_{t-1}$), E_t is the expectations operator based on information available at time t and Z_t is an exogenous driving force.

To consider an environment where agents are always unsure about the fundamentals and optimism can be well-defined, let us assume that Z_t is the sum of a persistent process $\theta_t = \rho\theta_{t+1} + \nu_t$ ($0 < \rho < 1$), plus a white noise process ϵ_t (where all stochastic elements are assumed to be Gaussian). Agents are assumed to observe Z_t but not its components. Moreover, let us assume that at time t agents get a noisy signal s_t regarding ν_{t+1} , such that $s_t = \nu_{t+1} + \omega_t$. Given this setup, the expectation of θ_{t+1} conditional on information at time t can be expressed as:

$$\begin{aligned} E_t \theta_{t+1} &= \rho \lambda E_{t-1} \theta_t + (1 - \lambda)(\theta_t + \epsilon_t) + \psi(\nu_{t+1} + \omega_t) \\ &= \rho(1 - \lambda) \sum_{i=0}^{\infty} (\rho \lambda)^i (\theta_{t-i} + \epsilon_{t-i}) + \psi \sum_{i=0}^{\infty} (\rho \lambda)^i (\nu_{t+1} + \omega_t), \end{aligned}$$

where λ and ψ reflect the signal extraction problems faced by the agents and where these two parameters can be shown to lie between 0 and 1 and depend on the variances of

the driving forces. For each of these two environments we can derive the impulse-response of the *growth* in GDP induced by a shock ω_t , which will be referred to as a noise shock.

In environment (1), the impulse-response is given by an initial positive effect on growth of size $\frac{\alpha\psi}{1-\alpha\rho}$, followed by a contractionary phase where the effects on growth are negative and given by $-\frac{\alpha\psi}{1-\alpha\rho}(\rho\lambda)^i$ for i periods after the receipt of the signal ($i \geq 1$). Hence, in environment (1), noise (unfounded optimism) causes a boom followed by a recessionary period.

In environment (2), the impulse-response of GDP growth is given by a declining series of positive effect of size $\frac{\alpha\psi}{1-\alpha\rho}(\rho\lambda)^i$ for periods $i \geq 0$, hence in this case there is no induced recession.

While in both environments a noise shock leads to an initial period of positive growth, the first environment leads to a subsequent recession while the second environment simply exhibits a petering out boom. How could we tell these two environments apart? Obviously, if we could observe ω_t this would be rather easy. However, this is generally not the case. Instead, let us suppose that we can observe agents' forecasts of next period's growth rate, as denoted by $E_t g_{t+1}$ and we have an instrument I_t that is positively correlated with ω_t and uncorrelated with ν_{t-i} and ϵ_{t-i} (for all i). Our Mission Chief fixed effects are a good candidate for such an instrument as Mission Chiefs get re-allocated across countries in a pseudo-random fashion, while econometric estimates indicate that a higher fixed effect is associated with unfounded forecast optimism (i.e.: a positive noise shock ω_t). However, because Mission Chiefs typically serve on each country for several years, it is more than plausible that this positive correlation persists over time. Such positive correlation is also likely to be present for our second instrument (one-year ahead forecast errors determined by a "bullish" or "bearish" atmosphere in the IMF country team). The remaining question therefore becomes: can we still use such instruments to differentiate between these two environments when I_t is correlated with future noise shocks ω_{t+1} (with a particular interest in the most plausible case of positive correlation)?

We claim that our instrument can still do so, provided that its correlation with future noise shocks is not negative. To see this, consider an IV regression for the growth rate at time $t + 1$ (or a recession dummy) on the forecast error for growth at $t + 1$ based on information given at time t (i.e.: $E_t g_{t+1} - g_{t+1}$), where we instrument the forecast error using the instrument I_t . If the result is a negative coefficient, we claim that we are in the first environment where the noise-driven forecast error is creating a recession, while if we get a positive coefficient on the instrumented forecast error we are likely in the second environment (but can't rule out the first environment).

To verify the above claim, one can write the growth rate at $t + 1$ as follows: in environment (1) we obtain:

$$g_{t+1} = -\frac{\alpha(1-\lambda\rho)}{1-\alpha\rho\lambda}(E_t g_{t+1} - g_{t+1}) + \frac{\alpha\psi(1-\alpha)}{(1-\alpha\rho)(1-\alpha\rho\lambda)}\omega_{t+1} + \xi_{t+1}, \quad (6)$$

where ξ_{t+1} is a complicated sum of ν s and ϵ s that does not include any terms ω_{t-i} for any i . So if we estimate (6) by instrumental variables using I_t to instrument the forecast error ($E_t g_{t+1} - g_{t+1}$), then we should get a consistent estimate of $-\frac{\alpha(1-\lambda\rho)}{1-\alpha\rho\lambda}$ if I_t is uncorrelated with ω_{t+1} and ξ_{t+1} . If I_t is positively correlated with ω_{t+1} , then the IV estimate will be biasing this coefficient towards a positive number. If we nonetheless obtain a negative estimate for the effect of the forecast error on growth, then we can conclude that we are in case (1) and that optimism causes a delayed recession. A positive estimate, on the other hand, would be less conclusive as that could simply stem from a positive correlation between the instrument and ω_{t+1} .

Things are very similar in environment (2). There we obtain:

$$g_{t+1} = \frac{\alpha\lambda\rho}{1-\alpha\rho\lambda}(E_t g_{t+1} - g_{t+1}) + \frac{\alpha\psi}{(1-\alpha\rho)(1-\alpha\rho\lambda)}\omega_{t+1} + \xi_{t+1}, \quad (7)$$

where ξ_{t+1} again does not include ω_{t-i} for any i . So if we estimate (7) by instrumental variables using I_t to instrument the forecast error ($E_t g_{t+1} - g_{t+1}$), then we get a consistent estimate of $\frac{\alpha\rho}{1-\alpha\rho\lambda}$ if I_t is also uncorrelated with ω_{t+1} . If I_t is positively correlated with ω_{t+1} , then the IV estimate will be biasing the coefficient upwards (more positive). Again, obtaining a negative estimate for the effect of the forecast error on growth is a strong finding given the likely presence of positive correlation between I_t and ω_{t+1} .

If, on the other hand, our instrument were negatively correlated with ω_{t+1} , a symmetric argument to the above implies that negative estimates could arise purely due to the associated negative bias. However, this case seems implausible as a negative correlation between I_t and ω_{t+1} requires a “flipping” allocation of IMF Mission Chiefs, with overly optimistic types consistently being replaced by overly cautious ones at the annual frequency. Apart from the fact that this requires IMF decision makers to know the exact type of each Mission Chief, such an alternating allocation policy is not in line with IMF practices. Similarly, with respect to our second instrument (one-year ahead forecast errors), this would require the team’s perspective on a country’s economic future to systematically alternate between overly-positive and overly-negative from one year to the next, which seems rather unlikely.

9 Appendix C: results when the dependent variable is real GDP growth

Here, we display the equivalents of our key tables for the global sample when the dependent variable is the growth rate of real GDP (instead of a recession dummy). Table numbers corresponding to tables in the main text are in parentheses.

Table C1 (~2b): instrumental variable regressions

Dependent variable: real GDP growth rate				
	(1: full)	(2: full)	(3: eme/dev)	(4: eme/dev)
forecast error on growth " $F_{3,t-3}$ "	-1.388208*** (-3.18)	-1.146396*** (-2.70)	-1.402359** (-2.30)	-1.079749** (-1.99)
public debt/GDP (t-1)		-0.0022497 (-0.95)		-0.0017742 (-0.65)
inflation (t-1)		0.134745** (2.52)		0.0135519** (2.43)
avg growth in trading partners		0.0125092 (0.12)		-0.0270904 (-0.26)
pct change in terms-of-trade		0.0125762* (1.90)		0.0144283** (2.11)
time and country fixed effects	yes	yes	yes	yes
countries	183	149	166	132
obs	2,704	2,282	2,243	1,710

Note: z-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level.

Table C2 ($\sim 3b$): instrumental variable regressions

Dependent variable: real GDP growth rate				
	(1: full)	(2: full)	(3: eme/dev)	(4: eme/dev)
forecast error on growth " $F(2)_{3,t-3}$ "	-1.397817*** (-4.82)	-1.551172*** (-4.96)	-1.496284*** (-4.07)	-1.69447*** (-4.29)
public debt/GDP (t-1)		-0.0027589 (-1.15)		-0.0023162 (-0.83)
inflation (t-1)		0.0285066 (1.63)		0.028489 (1.62)
avg growth in trading partners		-0.1311117 (-1.60)		-0.163496** (-1.97)
pct change in terms-of-trade		0.001386 (0.16)		0.0016217 (0.17)
time and country fixed effects	yes	yes	yes	yes
Hansen J-statistic (p-value)	0.243 (0.6222)	0.086 (0.7688)	0.110 (0.7397)	0.112 (0.7374)
countries	182	166	147	132
obs	2,563	2,136	1,986	1,621

Note: z-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level. The instrument set includes both MC fixed effects and one-year ahead forecast errors.

10 Appendix D: logit and OLS estimation results

Results from the main text indicate that past over-estimations of growth increase the odds with which a country ends up in recession. These results were established by addressing the endogeneity problem through an IV approach. However, as already emphasized in the seminal IV-contribution of Sargan (1958) and reiterated in a recent paper by Young (2017), IV estimates are inefficient if the endogeneity problem is not severe and better information can sometimes be obtained by analyzing OLS, logit, or probit results.²⁶ This is potentially true in the present setup, where the endogeneity problem can be mitigated by lagging the forecast error variable instead of using an IV approach. While some lagging was already present in the IV-regressions for economic reasons (many of the relevant decisions are based upon multi-year horizons), introducing further lags may allow us to proceed without IV.

Let's indicate these further lags by $k > 0$. In that case, we can analyze the impact of year T forecast errors (for years $T + 1, T + 2, \dots, T + h$) on the prevalence of recessions in year $T + h + k$. We then thus ask whether a recession in, say, the year 2001 (for $k = 1$), 2002 ($k = 2$), or 2003 ($k = 3$) was more likely to arise if growth expectations in 1997 for the years 1998, 1999, and 2000 were overly optimistic. Having $k > 0$ cuts the mechanical relationship between the growth outcome (which now lies in the year 2001 or beyond) and previous forecast errors (constructed by only using data from *before* the year 2001). This opens the door to an analysis which does not rely on IV. However, such a strategy can still be subject to endogeneity issues if economies systematically started losing momentum in years prior to a recession (as that would mechanically generate positive forecast errors in years before a downturn). To look at this possibility, we examined the growth behavior of all IMF member states over the period 1990-2016 prior to recessions. As can be seen in Table D1, the average per capita growth rate remains fairly stable in the run-up to a recession, suggesting that recessions arrive rather suddenly (a feature also described in the literature on business cycle asymmetries, see e.g. Acemoglu and Scott (1997)). While this does not guarantee that there isn't some third factor – distinct from optimism – systematically affecting past forecast errors as well as current recessions, it does suggest that such a problem may be somewhat minor.

²⁶Probit or logit specifications are often applied in the “early warning” literature on crisis-determinants (see e.g. Eichengreen, Rose, and Wyplosz (1995) and Gourinchas and Obstfeld (2012)).

Table D1: average annual growth rate of real per capita GDP in years prior to recessions

avg in non-recession years	avg $(T - 4), (T - 3), (T - 2)$	avg $(T - 3), (T - 2), (T - 1)$
3.8%	3.6%	3.9%

Note: in these calculations, year T always denotes the first year of a recession. All years $(T-4), \dots, (T-1)$ are recession-free.

Proceeding with logit gives rise to the following estimation results for the global sample (results are fully robust to using probit or focusing on the restricted sample of emerging and developing countries only). As Table D2 illustrates, our main result survives when considering k up to 2 lags (in which case we are analyzing the impact of 1995 forecast errors for the years $\{1996, 1997, 1998\}$, on the probability of a recession occurring in the year 2000 – a temporal distance that seems long enough to mitigate endogeneity-concerns).

Table D2: logit regressions when forecast error is lagged $(k + 3)$ periods, global sample

Dependent variable: recession dummy (1 = presence of recession)			
	(1: $k = 3$)	(2: $k = 2$)	(3: $k = 1$)
forecast error on growth " $F_{3,t-3-k}$ "	0.0180404 (0.93)	0.0508912*** (2.68)	0.1014209*** (5.48)
time and country fixed effects	yes	yes	yes
countries	167	168	171
obs	3,214	3,396	3,601

Note: t-statistics in parentheses. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level.

Table D3 shows that the significant coefficient on the forecast error variable survives the inclusion of our standard set of covariates that are commonly believed to affect the probability of a recession occurring.

Table D3: logit regressions with covariates when forecast error is lagged ($k + 3$) periods, global sample

Dependent variable: recession dummy (1 = presence of recession)			
	(1: $k = 3$)	(2: $k = 2$)	(3: $k = 1$)
forecast error on growth " $F_{3,t-3-k}$ "	0.0161239 (0.73)	0.0470071** (2.20)	0.1081991*** (5.11)
public debt/GDP (t-1)	0.0038168*** (2.74)	0.0040461** (3.06)	0.0041014*** (3.30)
inflation (t-1)	0.0011684 (0.72)	-0.0000383 (-0.04)	-0.000004 (-0.00)
avg growth in trading partners	-0.1963731*** (-4.03)	-0.1621607*** (-3.44)	-0.1523854*** (-3.38)
pct change in terms-of-trade	-0.0055524 (-1.47)	-0.004661 (-1.28)	-0.0043153 (-1.23)
time and country fixed effects	yes	yes	yes
countries	150	151	154
obs	2,657	2,797	2,969

Note: t-statistics in parentheses. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level. Endogenous variables are lagged by one period.

It is also interesting to analyze logit-results for the regressions in Section 4, where crisis dummies act as the dependent variable. In that case, the LHS is no longer mechanically related to the growth-forecast error variable on the RHS (reducing the need to rely on IV or introduce additional lags). In particular, by including a dummy variable which controls for recessions, one can analyze whether past over-optimism makes the arrival of certain crises more likely *whilst controlling for the state of the economy*.

Along these lines, Table D4 confirms our IV-results by showing that past over-optimism²⁷ makes the arrival of fiscal- and BoP crises more likely.

²⁷As measured by $F_{3,t-4}$, but results typically survive up to $F_{3,t-6}$ (or beyond).

Table D4: logit regressions

Dependent variable: crisis dummy (1 = presence of crisis)				
	(1:fiscal crisis)	(2: fiscal crisis)	(3:IMF program)	(4: IMF program)
forecast error on growth " $F_{3,t-4}$ "	0.1266774*** (6.60)	0.1409052*** (6.14)	0.0865626*** (4.29)	0.1275916*** (5.28)
public debt/GDP (t-1)		0.0130631*** (6.75)		0.0040253*** (2.75)
inflation (t-1)		0.0091687** (2.47)		-0.0027348* (-1.79)
avg growth in trading partners		-0.0229725 (-0.62)		-0.0222658 (-0.45)
pct change in terms-of-trade		0.0033591 (0.97)		0.0017901 (0.50)
recession dummy		0.8597959*** (6.25)		0.1980642 (1.34)
time and country fixed effects	yes	yes	yes	yes
countries	132	119	112	103
obs	2,679	2,280	2,384	2,003

Note: t-statistics in parentheses. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level.

Tables D5, D6, and D7 show that an OLS-approach yields similar results.

Table D5: OLS regressions when forecast error is lagged ($k + 3$) periods, global sample

Dependent variable: recession dummy (1 = presence of recession)			
	(1: $k = 3$)	(2: $k = 2$)	(3: $k = 1$)
forecast error on growth " $F_{3,t-3-k}$ "	0.0016635 (0.44)	0.0070361* (1.93)	0.0156576*** (4.13)
time and country fixed effects	yes	yes	yes
countries	184	185	186
obs	3,527	3,704	3,878

Note: t-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level.

Table D6: OLS regressions with covariates when forecast error is lagged ($k + 3$) periods, global sample

Dependent variable: recession dummy (1 = presence of recession)			
	(1: $k = 3$)	(2: $k = 2$)	(3: $k = 1$)
forecast error on growth " $F_{3,t-3-k}$ "	0.0011849 (0.29)	0.0060964 (1.54)	0.0161568*** (3.86)
public debt/GDP (t-1)	0.000627** (2.18)	0.0006551** (2.55)	0.0006602*** (2.83)
inflation (t-1)	0.0001653 (0.46)	-0.0000224 (-0.14)	-0.0000119 (-0.07)
avg growth in trading partners	-0.0171324 (-1.04)	-0.0143511 (-0.97)	-0.0136017 (-0.96)
pct change in terms-of-trade	-0.0007727 (-1.44)	-0.0006899 (-1.31)	-0.0006162 (-1.25)
time and country fixed effects	yes	yes	yes
countries	167	167	167
obs	2,917	3,053	3,190

Note: t-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level. Endogenous variables are lagged by one period.

Table D7: OLS regressions

Dependent variable: crisis dummy (1 = presence of crisis)				
	(1:fiscal crisis)	(2: fiscal crisis)	(3:IMF program)	(4: IMF program)
forecast error on growth " $F_{3,t-4}$ "	0.0222807*** (4.95)	0.0242053*** (5.46)	0.012702*** (2.96)	0.0184258*** (3.55)
public debt/GDP (t-1)		0.0019895*** (5.39)		0.0007892 (1.42)
inflation (t-1)		0.0004139*** (3.75)		-0.0002349 (-1.62)
avg growth in trading partners		-0.0046344 (-1.08)		-0.0011767 (-0.31)
pct change in terms-of-trade		0.0003677 (0.74)		0.0003258 (0.77)
recession dummy		0.1221439*** (6.72)		0.0355325 (1.62)
time and country fixed effects	yes	yes	yes	yes
countries	184	165	187	167
obs	3,690	3,166	3,913	3,190

Note: t-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level.

Given that there is uncertainty over which exact covariates to include in the regression, Table D8 shows that the estimate for the coefficient on the forecast error variable continues to be significantly positive in a Bayesian Model Averaging (BMA) exercise ($\hat{\beta} = 0.018$, t-statistic > 3), using the estimator developed in Magnus, Powell, and Prüfer (2010). It estimates models featuring all possible combinations for a large number of explanatory variables, given the presence of our main focus variable of interest (the average forecast error on growth). So every time one of the non-focus, “auxiliary” variables is dropped, a different model arises. A priori, all of these models are assigned equal prior probability, while the averaged parameter estimates (presented in Table D8) come about by weighing according to posterior model probabilities. As Column 2 of Table D8 shows, the coefficient on the average forecast error on growth remains significant if we do not force its presence and treat it as an auxiliary variable instead. It is striking to note from Table D8 that there are not many variables on which estimated coefficients are robustly significant. In fact, *the forecast error variable is the only one to do so* – suggesting that it is a relevant determinant of future recessions.

Table D8: BMA regressions explaining the occurrence of recessions, global sample

Dependent variable: recession dummy (1 = yes)		
	(1: $F_{3,t-4}$ always present)	(2: $F_{3,t-4}$ as auxiliary variable)
forecast error on growth " $F_{3,t-4}$ "	0.0183287*** (3.84)	0.0177432*** (3.13)
public debt/GDP (t-1)	0.0000749 (0.32)	0.0000743 (0.32)
inflation (t-1)	-0.0000106 (-0.07)	-0.0000104 (-0.07)
govt spending/GDP (t-1)	-0.0000967 (-0.13)	-0.0000949 (-0.13)
current account balance/GDP (t-1)	-0.0017875 (-0.61)	-0.0019068 (-0.63)
FDI/GDP (t-1)	0.0005146 (0.56)	0.0005219 (0.57)
population growth (t-1)	0.0187978 (0.79)	0.0189366 (0.80)
ICRG civil war (t-1)	0.0004478 (0.09)	0.0004467 (0.09)
exchange rate overvaluation (t-1)	0.0000556 (0.20)	0.0000573 (0.20)
investment/GDP (t-1)	-0.0022391 (-0.61)	-0.0024415 (-0.64)
avg growth in trading partners	-0.0001758 (-0.13)	-0.0001784 (-0.13)
pct change in terms-of-trade	-0.0000268 (-0.13)	-0.000027 (-0.13)
time and country fixed effects	yes	yes
obs	1,602	1,602

Note: t (z)-statistics in parentheses, calculated using robust standard errors. * denotes significance at the 10% level, ** implies significance at the 5% level, *** indicates significance at the 1% level.

11 References

Acemoglu, Daron and Andrew Scott (1997), “Asymmetric Business Cycles: Theory and Time-series Evidence”, *Journal of Monetary Economics*, 40 (3), pp. 501-533.

Angrist, Joshua, Guido W. Imbens, and Alan Krueger (1999), “Jackknife Instrumental Variables Estimation”, *Journal of Applied Econometrics*, 14 (1), pp. 57-67.

Arezki, Rabah, Valerie A. Ramey, and Liugang Sheng (2017), “News Shocks in Open Economies: Evidence from Giant Oil Discoveries”, *Quarterly Journal of Economics*, 132 (1), pp. 103-155.

Bachmann, Rüdiger and Steffen Elstner (2015), “Firm Optimism and Pessimism”, *European Economic Review*, 79, pp. 297-325.

Baqir, Reza, Rodney Ramcharan, and Ratna Sahay (2005), “IMF Programs and Growth: Is Optimism Defensible?”, *IMF Staff Papers*, 52 (2), pp. 260-286.

Bertrand, Marianne and Antoinette Schoar (2003), “Managing with Style: The Effect of Managers on Firm Policies”, *Quarterly Journal of Economics*, 118 (4), pp. 1169-1206.

Beaudry, Paul, and Franck Portier (2004), “An exploration into Pigou’s theory of cycles”, *Journal of Monetary Economics*, 51 (6), pp. 1183-1216.

Blanchard, Olivier J., Jean-Paul L’Huillier, and Guido Lorenzoni (2013), “News, Noise, and Fluctuations: An Empirical Exploration”, *American Economic Review*, 103 (7), pp. 3045-3070.

Blanchard, Olivier J., Guido Lorenzoni, and Jean-Paul L’Huillier (2017), “Short-Run Effects of Lower Productivity Growth. A Twist on the Secular Stagnation Hypothesis”, *Journal of Policy Modeling*, 39 (4), pp. 639-649.

Cooper, Russell and Andrew John (1988), “Coordinating Coordination Failures in Keynesian Models”, *Quarterly Journal of Economics*, 103 (3), pp. 441-463.

Di Bella, Gabriel and Francesco Grigoli (2018), “Optimism, Pessimism, and Short-Term Fluctuations”, IMF Working Paper No 18/1.

Eberhardt, Markus and Andrea F. Presbitero (2015), “Public Debt and Growth: Heterogeneity and Non-Linearity”, *Journal of International Economics*, 97 (1), pp. 45-58.

Eichengreen, Barry, Andrew Rose, and Charles Wyplosz (1995), “Exchange Market Mayhem: The Antecedents and Aftermath of Speculative Attacks”, *Economic Policy*, 0 (21), pp. 249–96.

Enders, Zeno, Michael Kleemann, and Gernot Müller (2017), “Growth Expectations, Undue Optimism, and Short-Run Fluctuations”, Deutsche BundesBank Discussion paper No 11/2017.

IEO (2006), “An Evaluation of the IMF’s Multilateral Surveillance – Background Doc-

uments”, Washington, D.C.: Independent Evaluation Office of the International Monetary Fund.

IEO (2014), “IMF Forecasts: Process, Quality, and Country Perspectives”, Washington, D.C.: Independent Evaluation Office of the International Monetary Fund.

Frankel, Jeffrey (2011), “Over-optimism in Forecasts by Official Budget Agencies and Its Implications”, *Oxford Review of Economic Policy*, 27 (4), pp. 536-562.

GAO (2003), “International Financial Crises: Challenges Remain in IMF’s Ability to Prevent and Resolve Financial Crises”, GAO-03-734, Washington, D.C.: U.S. Government Accountability Office.

Genberg, Hans and Andrew Martinez (2014a), “On the Accuracy and Efficiency of IMF Forecasts: A Survey and Some Extensions”, IEO Background Document No. BD/14/04, Washington, D.C.: Independent Evaluation Office of the International Monetary Fund.

Genberg, Hans and Andrew Martinez (2014b), “User Perspectives on IMF Forecasts: Survey Methodology and Results”, IEO Background Document No. BD/14/01, Washington, D.C.: Independent Evaluation Office of the International Monetary Fund.

Genberg, Hans, Andrew Martinez, and Michael Salemi (2014), “The IMF/WEO Forecast Process”, IEO Background Document No. BD/14/03, Washington, D.C.: Independent Evaluation Office of the International Monetary Fund.

Gerling, Kerstin, Paulo A. Medas, Tigran Poghosyan, Juan Pablo Farah Yacoub, and Yizhi Xu (2017), “Fiscal Crises”, IMF Working Paper No. 17/86, Washington, D.C.: International Monetary Fund.

Gourinchas, Pierre-Olivier and Maurice Obstfeld (2012), “Stories of the Twentieth Century for the Twenty-First”, *American Economic Journal: Macroeconomics*, 4 (1), pp. 226-265.

Jorda, Oscar (2005), “Estimation and Inference of Impulse Responses by Local Projections”, *American Economic Review*, 95 (1), pp. 161-182.

Jorda, Oscar, Moritz H.P. Schularick, and Alan M. Taylor (2013), “When Credit Bites Back”, *Journal of Money, Credit, and Banking*, 45 (2), pp. 3-28.

Judson, Ruth A. and Ann L. Owen (1999), “Estimating Dynamic Panel Data Models: A Guide for Macroeconomists”, *Economics Letters*, 65 (1), pp. 9-15.

Loungani, Prakash (2001), “How Accurate Are Private Sector Forecasts? Cross-Country Evidence from Consensus Forecasts of Output Growth”, *International Journal of Forecasting*, 17 (3), pp. 419-432.

Luna, Francesco (2014), “IMF Forecasts in the Context of Program Countries”, IEO Background Document No. BD/14/05, Washington, D.C.: Independent Evaluation Office

of the International Monetary Fund.

Magnus, Jan R., Owen Powell, and Patricia Prüfer (2010), “A Comparison of Two Model Averaging Techniques with an Application to Growth Empirics”, *Journal of Econometrics*, 154 (2), pp. 139-153.

Nickell, Stephen J. (1981), “Biases in Dynamic Models with Fixed Effects”, *Econometrica*, 49 (6), pp. 1417-1426.

Nunn, Nathan and Leonard Wantchekon (2011), “The Slave Trade and the Origins of Mistrust in Africa”, *American Economic Review*, 101 (7), pp. 3221-3252.

Oh, Seonghwan and Michael Waldman (1990), “The Macroeconomic Effects of False Announcements”, *Quarterly Journal of Economics*, 105 (4), pp. 1017-1034.

Pigou, Arthur C. (1927), *Industrial Fluctuations*, London: MacMillan.

Rodríguez Mora, José V. and Paul Schulstad (2007), “The Effect of GNP Announcements on Fluctuations of GNP Growth”, *European Economic Review*, 51 (8), pp. 1922-1940.

Reinhart, Carmen M. and Kenneth S. Rogoff (2010), “Growth in a Time of Debt”, *American Economic Review Papers and Proceedings*, 100 (2), pp. 573-578.

Sargan, J. Denis (1958), “The Estimation of Economic Relationships Using Instrumental Variables”, *Econometrica*, 26 (3), pp. 393-415.

Stock, James H. and Motohiro Yogo (2005), “Testing for Weak Instruments in Linear IV Regression,” in: Donald W.K. Andrews and James H. Stock (eds), *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge, UK: Cambridge University Press.

Young, Alwyn (2017), “Consistency without Inference: Instrumental Variables in Practical Application”, mimeo, London School of Economics.

12 Figures

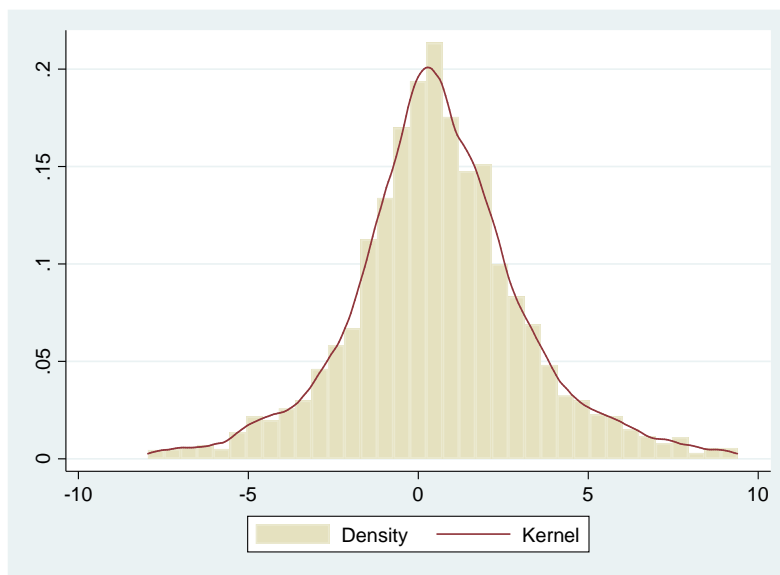


Figure 1: Kernel density plot for forecast error variable F_3 .

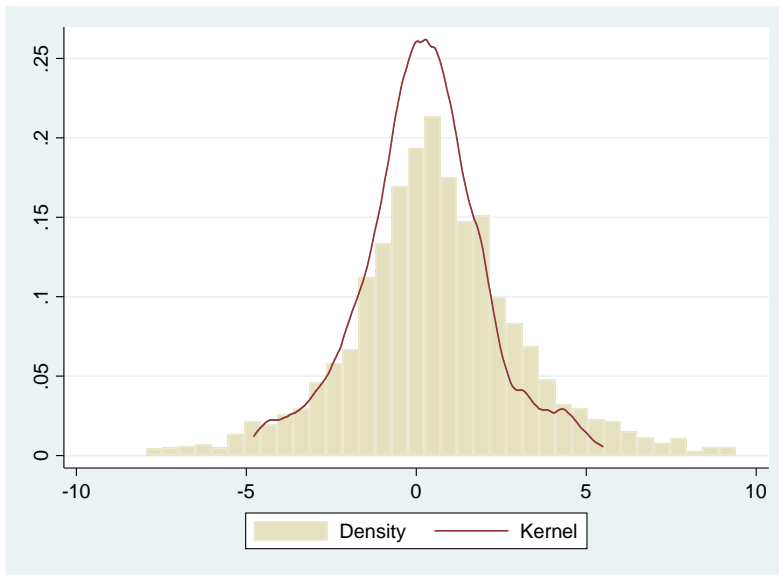


Figure 2: Kernel density plot for Mission Chief fixed effects μ .

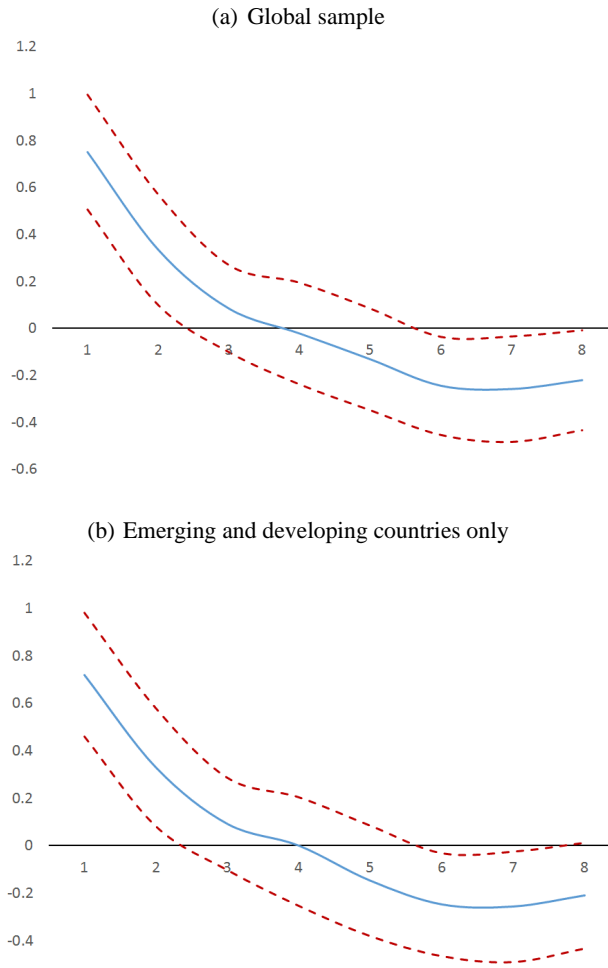


Figure 3: Response of real GDP growth to a forecast shock. Dashed lines indicate the 95% confidence interval.

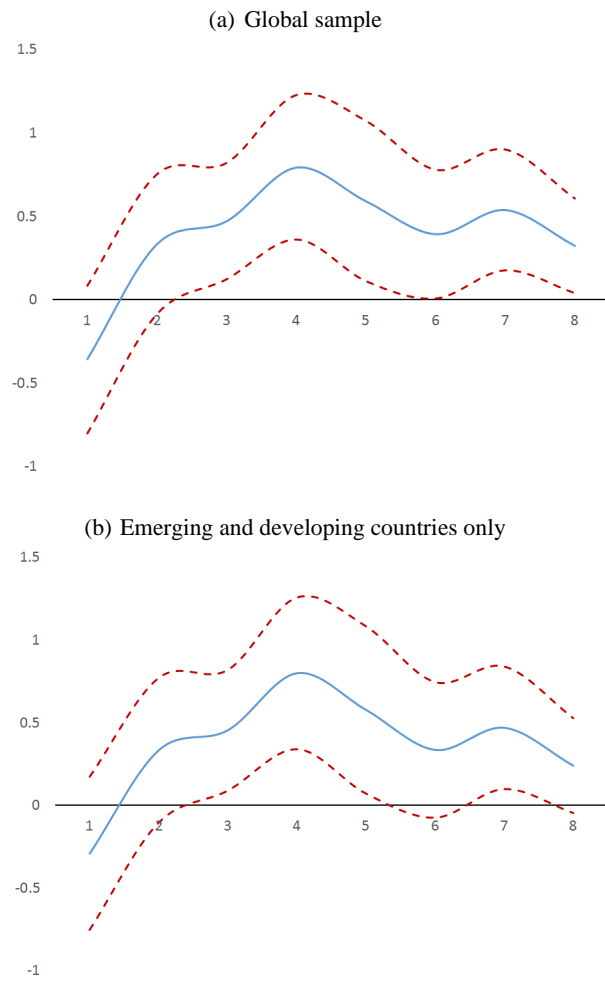


Figure 4: Response of the change in the debt-to-GDP ratio to a forecast shock. Dashed lines indicate the 95% confidence interval.

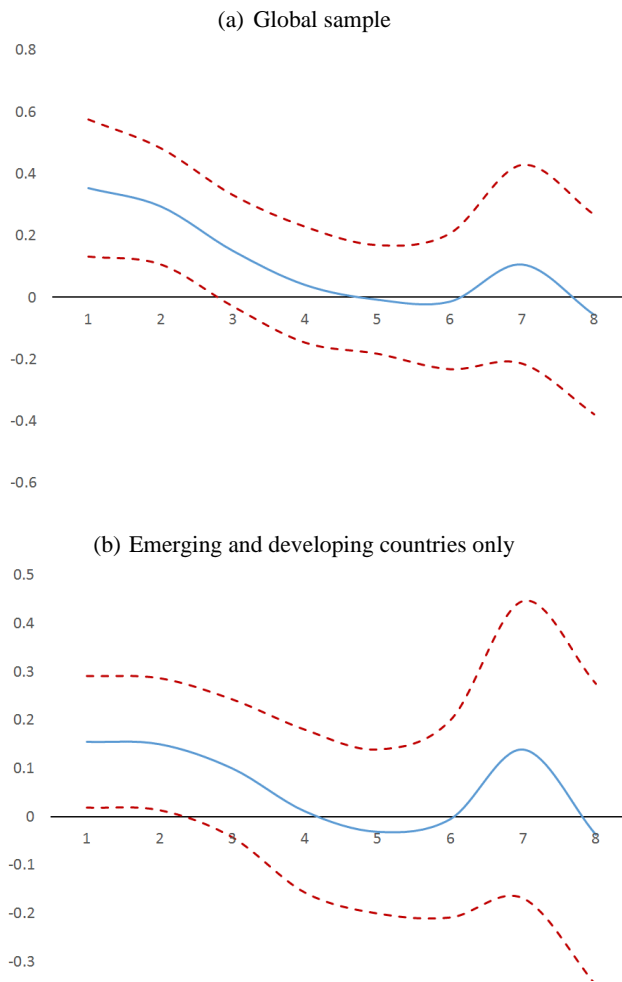


Figure 5: Response of the change in the private credit-to-GDP ratio to a forecast shock. Dashed lines indicate the 95% confidence interval.

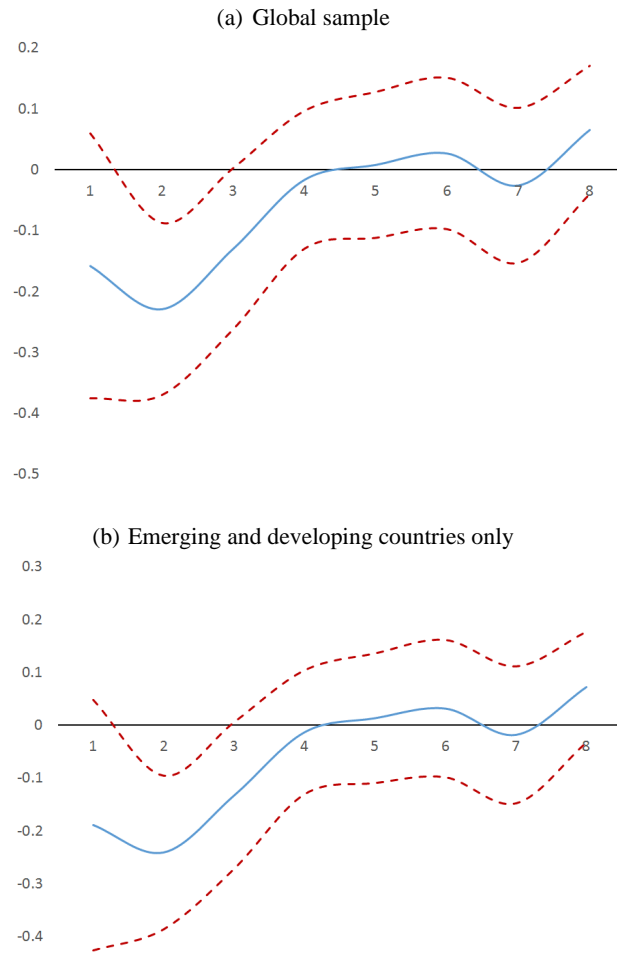


Figure 6: Response of the change in the investment-to-GDP ratio to a forecast shock. Dashed lines indicate the 95% confidence interval.