

Teulings, Coen N.; Zubanov, Nick

Working Paper

Is Economic Recovery a Myth? Robust Estimation of Impulse Responses

Tinbergen Institute Discussion Paper, No. 10-040/3

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Teulings, Coen N.; Zubanov, Nick (2010) : Is Economic Recovery a Myth? Robust Estimation of Impulse Responses, Tinbergen Institute Discussion Paper, No. 10-040/3, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/86695>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TI 2010-040/3

Tinbergen Institute Discussion Paper



Is Economic Recovery a Myth? Robust Estimation of Impulse Responses

Coen N. Teulings^{1,2}
Nick Zubanov²

*1 University of Amsterdam, and Tinbergen Institute;
2 CPB, The Hague.*

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Is Economic Recovery a Myth? Robust Estimation of Impulse Responses

Coen N. Teulings & Nick Zubanov

Abstract

We estimate the impulse response function (IRF) of GDP to a banking crisis, applying an extension of the local projections method developed in Jordà (2005). This method is shown to be more robust to misspecification than calculating IRFs analytically. However, it suffers from a hitherto unnoticed systematic bias which increases with the forecast horizon. We propose a simple correction to this bias, which our Monte Carlo simulations show works well. Applying our corrected local projections estimator to a panel of 99 countries observed between 1974-2001, we find that an average banking crisis yields a long-term GDP loss of around 10 percent with little sign of recovery within 10 years. GDP losses to banking crises are even more severe in African countries. Like the original Jordà's (2005) method, our extension of it is quite widely applicable.

JEL codes: G01; E27; C53

Key words: banking crisis; impulse response; panel data

The authors thank Adam Elbourne, Wouter den Haan, Oscar Jordà, Bert Smid, and the participants of 2010 Canadian Economics Association conference for comments on previous versions of this paper.

1 Introduction

The demise of Lehman Brothers left the world economy in a state of disarray. GDP in most OECD countries has declined by an order of magnitude of 5 percent, a number unseen since World War II and the Great Depression. Will output recover from this shock in the next five to ten years, or will (part of) the loss be permanent? The Council of Economic Advisors has a clear view on this issue: “*A key fact is that recessions are followed by rebound. Indeed, if periods of lower-than-normal growth were not followed by periods of higher-than-normal growth, the unemployment rate would never return to normal.*” (Cited in Greg Mankiw’s blog of March 3, 2009).

This view, implying no long-term GDP loss after a recession, is not shared by all economists. For example, Campbell and Mankiw (1987) show that there is little mean reversion in output. This does not automatically imply that the effect of a banking crisis will be permanent. GDP can be mixture of random processes, some of which have a unit root, while others do not. Banking crises may affect the non-unit root components of GDP, an idea paraphrased by Paul Krugman: “*I always thought the unit root thing involved a bit of deliberate obtuseness - it involved pretending that you didn’t know the difference between, say, low GDP growth due to a productivity slowdown like the one that happened from the 1973 to 1995, on one side, and low GDP growth due to a severe recession.*” (Krugman’s blog of March 3, 2009.) For Krugman, the conclusion that shocks in GDP have permanent effects is implied by the fact that long-run productivity growth follows a random walk. Contrary to productivity growth, short-run fluctuations in the business cycle have largely temporary effects.

Empirical evidence presented in a recent paper by Valerie Cerra and Sweta Saxena (2008), which goes under an ominous title “The Myth of Economic Recovery”, seems to run contrary to Krugman’s view. The authors estimate a dynamic model of GDP growth as a function of lagged growth rates and a dummy for the occurrence of a banking crisis. They use their regression estimates for a recursive calculation of the impulse-response function (IRF) of GDP to a banking crisis event. We refer to this approach as the *analytical estimator* of the IRF, since it calculates the IRF analytically from the assumed data generating process which parameters are estimated only once. As the title of their paper suggests, the authors find strong persistence of the initial negative effect of banking crises. However, the analytical estimator is sensitive to misspecifications of the data generating process for GDP, which may have dramatic consequences for IRF estimates as the forecast horizon increases. This is a serious constraint for the assessment of the long

run effects of banking crises on GDP. We give some simple examples illustrating the severity of this problem of the analytical IRF estimates.

Oscar Jorda (2005) proposed a method to reduce the dependence of the IRF estimates on the specification of the data generating process which we will refer to as the *local projections estimator*. Essentially, a local projections estimate of the IRF of variable y to a shock in a regressor x k periods after the shock is the coefficient on x in the regression of y_{t+k} , $k > 1$, on the regressors measured at time t , without including their intermediate values realised between t and $t+k$. This method is far more robust to misspecifications because a new regression is estimated for each k , instead of using the same set of coefficients from the assumed, and potentially incorrect, autoregressive specification for y in calculating the IRF analytically. This advantage comes at the cost of some loss in the IRF estimates' efficiency. We provide a formula for this efficiency loss for a simple case.

However, in the panel data context, Jorda's (2005) local projections method is shown to be prone to a hitherto unnoticed bias similar in origin to the bias in the dynamic panel regression estimates demonstrated by Nickell (1981). Intuitively, fixed effects soak up part of the dynamic interactions, so that the variation in the fixed effect is overestimated and the dynamic interactions are underestimated. We show that in our application where banking crises happen infrequently (28 out of 99 countries in our sample had no banking crises at all, and 56 had it only once) but have persistent effects, this problem can be serious, in particular for long forecasting horizons. We derive a formula for the size of the bias in a simple case and propose a solution for this problem – the inclusion of banking crises happening between t and $t+k$ in the local projections regression. Our Monte Carlo simulations show that our simple solution produces estimates which are fairly close to the true parameters and are more robust to misspecification than the analytical estimator of the IRF.

Turning to our empirical results, we find that an average GDP loss linked to a banking crisis is about 10 percent and persists for years after the crisis. There is appreciable heterogeneity in the effect of the crisis, with African countries suffering much greater losses after the crisis – 14 percent compared to 6 percent for Non-Africa. The rest of this paper is organized as follows. Section 2 discusses the local projections estimator. In section 2.1, we compare and contrast it to the analytical IRF estimator in terms of robustness to misspecifications and efficiency loss. Its robustness notwithstanding, section 2.2 shows that the local projections estimator is biased for longer forecast horizons. Therein we also present a method to eliminate the local projections bias, followed

by the results of numerical simulations (section 2.3) illustrating the bias performance of the corrected local projections estimator. In section 3 we apply the corrected local projections and other estimators to the data on banking crises in order to estimate their effect on GDP. Section 4 concludes with a brief summary of our findings.

2 IRF estimation in models with fixed effects

2.1 Analytical and local projections estimators

Consider the following stationary AR(R) panel data model with fixed effects:

$$y_t = \alpha_{0i} + \alpha_0^* t + \sum_{r=1}^R \alpha_{1r} y_{t-r} + \sum_{l=1}^L \alpha_{2l} d_{t-l} + u_t, \quad (1)$$

where y_t is log GDP per capita for country i in year t , d_t is a dummy variable that takes the value of 1 at the start of a banking crisis in country i and 0 otherwise, and the error term u_t is an i.i.d. random variable for country i , $u_t \sim N(0, \sigma^2)$. Since all our data are country specific, we omit the suffix i for all variables for the sake of convenience, except when necessary. Hence, α_{0i} denotes country fixed effect, while α_0^* refers to time trend common to all countries. Stationarity implies $\left| \sum_{r=1}^R \alpha_{1r} \right| < 1$. We assume that banking crises arrive randomly at a country specific rate λ_i , so that $E[d_t - \lambda_i] = 0$ (see Cerra and Saxena (2008) for some evidence on this issue). We further assume that banking crises happen independently of all forward and backward realizations of the error term, that is, $E[u_s d_t] = 0, \forall s, t$.

Define the IRF of GDP to banking crisis k years after its start as

$$\begin{aligned} \text{IRF}(k) &\equiv E[y_{t+k-1} | d_{t-1} = 1, y_s, d_{s-1}, s < t] \\ &\quad - E[y_{t+k-1} | d_{t-1} = 0, y_s, d_{s-1}, s < t]. \end{aligned} \quad (2)$$

Based on this definition, one can calculate the IRF analytically for each k by expressing the conditional expectation of y_{t+k} as a function of equation (1)'s parameter estimates. For example, if $R = L = 1$ $\text{IRF}(k) = \alpha_{21} \alpha_{11}^{k-1}$. We will refer to this method of calculating IRFs as the *analytical estimator*. There are two problems with the analytical estimator. First, it relies heavily on the correct specification of the underlying model (1). As the model includes more lags of y_t and d_t , and as the length of the forecast k increases, $\text{IRF}(k)$ becomes a more and more complex expression that is more and more sensitive to even slight specification errors. Since we are interested in the long run effect of a banking crisis, this is a serious problem. Second, the growing complexity of $\text{IRF}(k)$ will make the calculation of its standard errors increasingly

cumbersome and the distribution of the parameter increasingly fat-tailed since it raises the underlying coefficient to higher powers (e.g. a_{11}^{k-1} in the example above).

An alternative to calculating the IRF analytically which circumvents both these problems is known as the *local projections estimator* (Jorda, 2005). This estimator, initially developed for the case of a VAR model, is rather generally applicable and has recently been extended to the case of a nonstationary VAR (Chong, Jorda and Taylor, 2011). Here we consider its application for a one-equation model. The local projections method estimates $\text{IRF}(k)$ directly from the forecast equation for GDP k periods ahead, as we now explain. Since the process is stationary, the lagged values of y_t can be eliminated from the model by recursive substitution of lagged versions of equation (1) to get

$$y_t = \gamma_{0i} + \gamma_0^* t + \sum_{l=1}^{\infty} \gamma_{2l} d_{t-l} + \sum_{m=1}^{\infty} \gamma_{3m} u_{t-m} + u_t, \quad (3)$$

where the γ -parameters are functions of the α -parameters. For example, for $R = L = 1$:

$$\begin{aligned} \gamma_{0i} &= \alpha_{0i} / (1 - \alpha_{11}) - \alpha_0^* \alpha_{11} / (1 - \alpha_{11})^2, \\ \gamma_0^* &= \alpha_0^* / (1 - \alpha_{11}), \\ \gamma_{2l} &= \alpha_{21} \alpha_{11}^{l-1}, \\ \gamma_{3m} &= \alpha_{11}^m. \end{aligned}$$

Therefore, y_t can be expressed as a linear function of all past shocks d_t and u_t . Let y_{tk} denote a k -period ahead forecast conditional on the information available at $t - 1$:

$$y_{tk} \equiv \text{E}[y_{t+k-1} | y_s, d_s, s < t] = \text{E}[y_{t+k-1} | u_s, d_s, s < t]$$

The second equality uses equation (3) to eliminate the lags of y_t from equation (1). Since $\text{E}[d_s] = \lambda_i$ and $\text{E}[u_s] = 0$, y_{tk} satisfies:

$$\begin{aligned} y_{tk} &= \gamma_{0ik} + \gamma_0^* (t + k - 1) + \sum_{l=k}^{\infty} \gamma_{2l} d_{t+k-1-l} + \sum_{m=k}^{\infty} \gamma_{3m} u_{t+k-1-m}, \quad (4) \\ \gamma_{0ik} &\equiv \gamma_{0i} + \sum_{l=1}^{k-1} \gamma_{2l} \lambda_i. \end{aligned}$$

The shocks hitting y_t between time t and time $t + k$ do not affect the forecast y_{tk} since their expected value conditional on the information available at the moment that the forecast is made, is zero. For $k = 1$, we

obtain $y_{t1} = y_t - u_t$, that is, the one-period ahead forecast is the actual value of y_t minus the error term. By the definition of IRF (equation (2)), $\text{IRF}(k) = \gamma_{2k}$. This substitution procedure can be reversed, by solving equation (1) for u_t and using that solution to eliminate u_t and its lags from equation (4):

$$y_{tk} = \delta_{0ik} + \delta_{0k}^* t + \sum_{r=1}^R \delta_{1rk} y_{t-r} + \sum_{l=1}^L \delta_{2lk} d_{t-l}, \quad (5)$$

where the δ -parameters can be expressed as functions of the γ -parameters in the original forecast equation; in particular, $\delta_{21k} = \gamma_{2k} = \text{IRF}(k)$.

The local projections method involves estimating equation (5) with an added error term v_{tk} :

$$y_{t+k-1} = \delta_{0ik} + \delta_{0k}^* t + \sum_{r=1}^R \delta_{1rk} y_{t-r} + \sum_{l=1}^L \delta_{2lk} d_{t-l} + v_{tk}, \quad (6)$$

where

$$v_{tk} \equiv \sum_{l=1}^{k-1} \gamma_{2l} (d_{t+k-1-l} - \lambda_i) + \sum_{m=1}^{k-1} \gamma_{3m} u_{t+k-1-m} + u_{t+k-1} \quad (7)$$

This method can be interpreted as estimating a reduced form model, where all “endogenous” variables - the realizations of y_s between the moment t at which the forecast is made and the moment $t+k$ for which y_s is forecasted - have been eliminated. Clearly, the forecast error is serially correlated, following a moving average process of order k . Fully efficient estimation of the IRF by the local projections method would therefore require an appropriate general least squares transformation of the original equation to eliminate serial correlation in the residuals. The simpler alternative, which we will follow, is to use OLS with serial correlation-robust standard errors, which of course involves some efficiency loss. To gauge the magnitude of this loss compared to the case of analytical IRF estimation, the following proposition is helpful:

Proposition 1

Consider model (1) with $\alpha_{0i} = \alpha_0^ = 0$, $R = 1$ (that is, $AR(1)$ without trend and intercept) and $L = 0$ (that is, no banking crises). The variance of the analytical estimator for the k -period ahead forecast, $y_{tk} = \hat{\alpha}_{11}^k$, satisfies:*

$$\text{plim} (N(T-1)\text{Var} [\hat{\alpha}_{11}^k]) = k^2 (1 - \alpha_{11}^2) \alpha_{11}^{2(k-1)}$$

where T is the number of observations per country and where N is the number of countries. The variance of the local projections estimator for

the same forecast, $y_{tk} = \hat{\delta}_{11k}$, satisfies

$$\text{plim} \left(N(T-k) \text{Var} \left[\hat{\delta}_{11k} \right] \right) = 1 - \alpha_{11}^{2k} + 2(k-1) - 2 \frac{\alpha_{11}^2 - \alpha_{11}^{2k}}{1 - \alpha_{11}^2}$$

The proof of this proposition is in the Appendix. By way of visualizing the differences between the analytic IRF and local projections forecast variances, Figure 1 plots the variances of $\hat{\alpha}_{11}^k$ and $\hat{\delta}_{11k}$ and their ratio for $\alpha_{11} = 0.95$. The efficiency loss of the local projections estimator compared with analytical IRF estimator is growing with k and can be quite large. However, the advantage of the local projections estimator is its greater robustness to misspecifications of the underlying model (1) than the analytic estimator of the IRF can deliver. We illustrate this point in the coming two paragraphs.

Consider a simple third-order autoregressive model (1) where $\alpha_{0i} = \alpha_0^* = 0$, $R = 3$, and $L = 1$. Suppose $\alpha_{12} = 0$ and α_{13} is so much smaller in magnitude than α_{11} that it turns out to be insignificant in the regression, and the econometrician decides therefore to proceed with the parsimonious AR(1) specification. This misspecification will result in a biased estimate for the AR(1) parameter α_{11} , since the probability limit of its OLS estimate is equal to $\frac{\alpha_{11}}{1 - \alpha_{13}(\alpha_{11} + \alpha_{13})}$.¹ Figure 2 plots the IRF for the true AR(3) model and the estimated AR(1) model for the case of $\alpha_{11} = 0.95$ and $\alpha_{13} = -0.10$. In this case, if one proceeds with the misspecified AR(1) process one gets a biased $\alpha_{11} = 0.88$, and consequently an IRF which underestimates the shock for relatively small k and overestimates it for larger k . On the other hand, applying the

¹Consider the general AR(3) model:

$$y_t = \alpha_{11}y_{t-1} + \alpha_{12}y_{t-2} + \alpha_{13}y_{t-3} + u_t, u \sim N(0, \sigma^2)$$

Rewrite this model in the three equivalent ways: first, as it stands, second, with the part $\alpha_{11}y_{t-1}$ brought to the left side, and third, with the part $\alpha_{11}y_{t-1} + \alpha_{12}y_{t-2}$ brought to the left side. Taking the variance of these three equations, one has a system of linear equations in the variance of y , denoted $\text{var}(y)$, and two autocovariances, $\text{cov}(y_t, y_{t-1})$ and $\text{cov}(y_t, y_{t-2})$:

- (i) $\text{var}(y) = (\alpha_{11}^2 + \alpha_{12}^2 + \alpha_{13}^2) \text{var}(y) + 2\alpha_{12}(\alpha_{11} + \alpha_{13}) \text{cov}(y_t, y_{t-1}) + 2\alpha_{11}\alpha_{13} \text{cov}(y_t, y_{t-2}) + \sigma^2$,
- (ii) $(1 + \alpha_{11}^2) \text{var}(y) - 2\alpha_{11} \text{cov}(y_t, y_{t-1}) = (\alpha_{12}^2 + \alpha_{13}^2) \text{var}(y) + 2\alpha_{12}\alpha_{13} \text{cov}(y_t, y_{t-1}) + \sigma^2$,
- (iii) $(1 + \alpha_{11}^2 + \alpha_{12}^2) \text{var}(y) - 2\alpha_{11}(1 - \alpha_{12}) \text{cov}(y_t, y_{t-1}) - 2\alpha_{12} \text{cov}(y_t, y_{t-2}) = \alpha_{13}^2 \text{var}(y) + \sigma^2$.

Solving this system and substituting the expressions for $\text{cov}(y_t, y_{t-1})$ and $\text{var}(y)$ in the OLS formula for α_{11} ,

$$\hat{\alpha}_{11} = \frac{\text{cov}(y_t, y_{t-1})}{\text{var}(y)},$$

one obtains the result in the text.

local projections method yields a consistent IRF, since δ_{21k} is estimated separately for each k so as to minimize the difference between the actual and predicted value of GDP for each k , irrespective of the assumed specification of the data generating process.

Another form of misspecification to which the local projections estimator is robust is the assumption that shocks in d_t and u_t have the same dynamics. This assumption will result in a bias in analytical estimator of the IRF in the presence of unobserved dynamic components in y_t . Suppose y_t has two components: an unobserved AR(1) process z_t , and the banking crisis d_t which effect dies out after a year:

$$\begin{aligned} y_{t+1} &= x_{t+1} + z_{t+1}, \\ z_{t+1} &= \gamma z_t + u_{t+1}, u \sim N(0, \sigma^2) \\ x_{t+1} &= \alpha_{21} (d_t - \lambda). \end{aligned}$$

The presence of z_t induces autocorrelation in y_t , and if z_t is unobserved, the AR(1) coefficient in the regression for y_t (model (1) under the simplifying restriction $\alpha_{0i} = \alpha_0^* = 0$) satisfies:²

$$\text{plim}(\hat{\alpha}_{11}) = \frac{\gamma}{1 + \lambda(1 - \lambda)\alpha_{21}^2(1 - \gamma^2)\sigma^{-2}}.$$

Consequently, a banking crisis will appear to have the same dynamic effect as any shock to z_t through the AR(1) term α_{11} , whereas in fact its effect is one-off.³ Analytical IRFs are clearly prone to this error. The local projections estimator, on the other hand, does not suffer from this problem because it estimates the IRF (the coefficient δ_{21k}) directly from the forecast equation (6) rather than relying on the (mis)estimated autoregression coefficients when calculating it analytically. Since the local projections estimator is robust to misspecification of the order of the AR process, albeit at the expense of lower estimation efficiency, this estimator gives a more reliable picture of the true IRF.

²Since

$$\begin{aligned} \text{var}[x_t] &= \lambda(1 - \lambda)\alpha_{21}^2, \\ \text{cov}[x_t, x_{t-1}] &= 0, \\ \text{var}[z_t] &= (1 - \gamma^2)^{-1}\sigma^2, \\ \text{cov}[z_t, z_{t-1}] &= \gamma(1 - \gamma^2)^{-1}\sigma^2, \\ \text{plim}(\hat{\alpha}_1) &= \frac{\text{cov}[x_t, x_{t-1}] + \text{cov}[z_t, z_{t-1}]}{\text{var}[x_t] + \text{var}[z_t]}. \end{aligned}$$

³Cai and Den Haan (2009) make the same point criticizing “one-type-shock” models of output which in their view tend to overestimate the impact of recessions.

2.2 The bias in the local projections estimator

Regrettably, the local projections estimator is subject to a bias which increases with the forecast horizon k . This bias has the same origins as the more general one, derived in Nickell (1981) and Alvarez and Arellano (2003), which applies to all dynamic panel models, with fixed or random effects. Namely, part of the persistent influence of the variables in the model is misattributed to the effects. The mechanisms through which the two biases come about, however, are different. The Nickell bias happens due to a negative correlation between the error term and lagged values of y_t induced by the fixed effects or first difference transformation, which correlation increases with the strength of autocorrelation in y (the coefficients α_{1r} in model (1)). The relative importance of this correlation fades as the panel grows longer, and consequently the Nickell bias decreases in magnitude as T becomes large relative to N (see Judson and Owen (1999) and further studied in Alvarez and Arellano (2003) for Monte Carlo evidence on this bias). In what follows, we will abstract from the Nickell bias.

The local projections bias occurs because the forecast equation (6) does not include banking crises happening within the forecast horizon, that is between the moment of forecasting at time t and the moment for which a forecast is made, $t + k$. The effects of these omitted crises enter equation (6)'s error term (see the first term of equation (7)) and consequently disturb the moment conditions identifying its regression coefficients. Unlike the Nickell bias, the local projections bias does not always decrease with T ; holding N fixed, it goes down only as T grows relative to the length of the forecast horizon k . In fact, as the example in the following paragraph shows, the local projections bias does not even require a dynamic specification to occur.

Suppose a country is hit by a banking crisis only once at, say, $t = 5$, so $d_5 = 1$ and $d_t = 0$ for all other t . Let us assume for the ease of explanation that $R = 0$ (no lagged dependent variables) and α_0^* (no trend). Consider what happens when the econometrician applies the local projections estimator to estimate IRF(4). He will regress y_{t+4} on d_t , a fixed effect and nothing else for all t . For $t = 1$, this is the right model, since $y_{t+4} = y_5$ is just equal to the country fixed effect (the banking crisis at $t = 5$ will start affecting GDP only a year later, see equation 1). However, for $t = 2$, $y_{t+4} = y_6$, which is affected by the banking crisis that will happen at $t = 5$, but that has not yet happened at $t = 2$. Therefore, y_6 is misspecified because in fact it is equal to the fixed effect plus IRF(1), not just the fixed effect that the local projections estimate. We have a similar misspecification again for $t = 3$, when y_{t+4} equals the fixed effect plus IRF(2), and for $t = 4$, when y_{t+4} equals fixed effect plus

IRF(3). From $t = 5$ onwards, we are back to the correct specification for y_{t+4} since the banking crisis happens at $t = 5$ and is therefore included in the model. Misspecification of y_{t+4} for some t 's results in an omitted variable bias to the estimate of the fixed effect, and this bias will spill over to the estimates for IRF(4) causing them to be attenuated. The following Proposition 2 gives insight into the magnitude of the biases to the fixed effect and IRF estimates for the case without lags of the dependent variable, $R = 0$.⁴

Proposition 2:

Consider equation (6). Let $R = \alpha_0^ = 0$ and let data be available for estimation for $t = 1, T$ (so $L + k - 1$ lags of the data are available) and suppose that there are no countries with a banking crisis happening in years $t = 1, k - 1$.⁵ The local projections estimates of the fixed effect (δ_{0ik}) and IRF(k) (δ_{2lk}), $\hat{\delta}_{0ik}$ and $\hat{\delta}_{2lk}$, respectively, satisfy:*

$$\begin{aligned} E \left[\hat{\delta}_{0ik} \right] &= \delta_{0ik} + (T - L)^{-1} \sum_{l=1}^{k-1} \gamma_{2l}, \\ E \left[\hat{\delta}_{2lk} \right] &= \delta_{2lk} - (T - L)^{-1} \sum_{l=1}^{k-1} \gamma_{2l}. \end{aligned}$$

This proposition, which proof is in the Appendix, implies that the IRF estimates δ_{2lk} are biased downward in absolute value. The bias is zero for $k = 1$, but it increases with the forecast horizon k , both because the cumulative effect of the omitted banking crisis dummies $\sum_{l=1}^{k-1} \gamma_{2l}$ becomes larger and because the effective number of observations becomes smaller since more lags of data are lost in the estimation when y_{t+k} is regressed on y_t . The bias is also larger for bigger L , because the estimated coefficient $\hat{\delta}_{2lk}$ will have to adjust to offset the bias in the estimated fixed effect $\hat{\delta}_{0ik}$ for L of the T observations available.

Our proposed solution to the bias in the local projections estimator is to augment equation (6) with the banking crises occurring between t and $t + k - 1$ and estimate the *corrected* local projections equation:

$$y_{t+k-1} = \delta_{0ik} + \delta_{0k}^* t + \sum_{r=1}^R \delta_{1rk} y_{t-r} + \sum_{l=1}^L \delta_{2lk} d_{t-l} \quad (8)$$

⁴Including lags of the dependent variable will not change the message of Proposition 2, that the local projections estimator brings an upward bias to the fixed effects and a downward bias to the IRF estimates. These lags, however, will have to be controlled away by means of orthogonal projections, which will make the maths less succinct without adding any extra insights.

⁵Having countries with banking crisis in these initial years complicates the expressions without changing the essentials of the argument.

$$\begin{aligned}
& + \sum_{l=1}^{k-1} \gamma_{2l} d_{t+k-1-l} + v_{tk}^*, \\
v_{tk}^* & = \sum_{m=1}^{k-1} \gamma_{3m} u_{t+k-1-m} + u_{t+k-1}.
\end{aligned}$$

where v_{tk}^* is the error term v_{tk} conditional on the information available at time t and the information on the occurrence of banking crises between t and $t + k - 1$ (compare it with the v_{tk} as defined by equation(7)). With the intermediate banking crisis observations included, solving the moment conditions will produce unbiased IRF estimates, since the error term no longer contains d_t .

The idea of augmenting the forecast equation with intermediate observations is not new and can be applied in a variety of settings beyond the immediate topic of this paper. A recent application of this idea is Faust and Wright (2011), showing that a forecast of excess bond and stock returns improves in accuracy when the forecast equation is augmented with the forecast errors observed *ex post* between t and $t+k$. Extending their argument to our case, adding intermediate banking crises will also reduce the local projections IRF's standard errors through a reduction in the variance of the original forecast equation (6)'s error term v_{tk} which contains the intermediate banking crises (see equation (7)). A further insight, specific to our panel data case, however, is that augmentation not only improves estimation efficiency but also reduces the estimation bias when the regressors in the forecast equation are correlated with its error term because of the presence of intermediate observations.

Interestingly, a slightly rewritten version of equation (8) yields estimates of all γ_{2l} for $l = 1, k$. This suggests an alternative IRF estimator, where all γ_{2k} 's are estimated at once from a version of equation (8) in which all lags of y_t are eliminated through recursive substitution (as in equation (4)):

$$y_{t+k-1} = \gamma_{0ik} + \gamma_{0!}t + \sum_{k=1}^{\infty} \gamma_{2k} d_{t-k} + v_t. \quad (9)$$

The advantage of estimating equation (9) rather than equation (8) is that (9) yields estimates of γ_{2k} free from the Nickell bias due to the lagged dependent variable. The disadvantage is that controlling for lags of y_t improves the efficiency of the estimation of γ_{2k} . Besides, there is a question about the number of lags of the banking crisis dummy to be included in the feasible version of equation (9), which ultimately is

an empirical question. In our regression analysis, we observe that the estimated γ_{2k} 's stabilize when the number of lags is 18 or more.

2.3 Monte Carlo simulations of the bias in the IRFs

This section presents some simulation results for the bias in the IRF estimates obtained from the three estimators discussed in the paper: (i) analytical estimator applied by Cerra and Saxena (2008); (ii) Jorda's (2005) local projections estimator based on equation (6) which does not include banking crises between years t and $t + k$; and (iii) the corrected local projections estimator based on equation (8) which includes these banking crises. We have generated 1000 counterfactual datasets based on the following AR(5) process:

$$\begin{aligned} y_t = & \alpha_{0i} + 0.25y_{t-1} + 0.8y_{t-2} + 0.4y_{t-3} - 0.1y_{t-4} - 0.5y_{t-5} \\ & - 0.035d_{t-1} - 0.045d_{t-2} - 0.03d_{t-3} - 0.01d_{t-4} - 0.01d_{t-5} \\ & + u_t, \end{aligned} \quad (10)$$

where the fixed effects $\alpha_{0i} \sim U(0, 3)$, d_t is equal to 1 if $\alpha_{0i}/5 + 3U(0, 1) < 0.45$, and zero otherwise, and the idiosyncratic error term u_t follows a standard normal distribution. For each of the 1000 datasets, one hundred observations of y_t were generated for $t = 1, 100$, after which the first 70 observations were discarded to avoid y_t being determined by the initial, randomly chosen, values. Hence, each dataset had dimensions corresponding to the real data at hand with $T = 30$ and $N = 100$ (see section 3.1). The parameters of equation (10) have been chosen so that the moments of the generated data are similar to those in the real data. Thus, the probability of a banking crisis is about 0.05, the correlation between the country fixed effect α_{0i} and the banking crisis dummy d_t is about -0.2 , and the analytical IRF that one would obtain in the perfect knowledge of the parameters of (10) is close to the one that we actually estimate on our data, see section 3.2.

We then estimated a selection of empirical specifications of equation (10) with different numbers of lags of y and d , as well as their corresponding forecast equations up to $k = 10$, on each resulting dataset using the fixed effects estimator with serial correlation-robust standard errors. Figure 3 plots for each of the three estimators the means of the IRFs estimated in each of the 1000 runs on different assumed specifications of the data generating process (10), as well as the true IRF implied by (10). The results for the analytical estimator are widely different from the true model. Parsimonious specifications produce much underestimated results on the entire length of the forecast horizon. Only for the number of lags of y and, especially, d close to the true 5, do the

analytical IRF estimates converge to the true IRF. The analytical estimator is particularly sensitive to the number of lags of d included in the estimated regression equation. The consequences of this sensitivity are important, because longer lags of the banking crisis dummy are increasingly likely to be excluded from the specification by an econometrician who strives for a parsimonious model. Thus, our simulations (results available on request) show that the estimated t -statistics of farther lags of d are progressively smaller and farther apart from their true levels as implied by the data generation process we specified.

The second method, Jorda's (2005) local projections estimator that does not include banking crisis observations within the forecast horizon, produces IRFs that are much more robust to misspecification and less biased. However, while the bias of the analytical estimator decreases as the model's specification begins to resemble correct, the bias in this version of the local projections estimator remains appreciable, and increases towards the end of the forecast horizon. The corrected local projections estimator, which includes intermediate banking crisis observations, produces the best results. Its estimated IRFs are robust to even serious misspecifications of the model, and its bias is smaller than that of the other two estimators. Hence, our simulation results corroborate our theoretical results.

All three estimators are sensitive to the standard Nickell bias which we have so far ignored. Although this bias can be large in finite samples, it disappears asymptotically. To gauge how it affects our estimates for different lengths of the available time period, we simulated the IRFs obtained through two methods, the analytical estimator and the corrected local projections estimator, applied to panels of different lengths $T = 20, 30, 40, 50$, and 100 . The results are plotted in Figure 4. For each estimator, we estimate one misspecified model (three instead of five lags for both y_t and d_t) and one correctly specified model. As expected, the bias is rather large for $T = 20$ and declines when T goes up. The estimated IRFs are tolerably close to the true IRF even for $T = 30$. The corrected local projections estimator performs better than analytical IRFs with respect to Nickell bias, too. First, the corrected local projections estimates are remarkably robust to model misspecifications. Second, as T increases they converge to the true parameter values quicker than the analytical IRF estimates do even when the model is correctly specified.

3 Empirical estimates of the IRFs

3.1 Data

Our dataset is compiled from several sources. The data on banking crises come from Gerard Caprio and Daniela Klingebiel's (2003) study with observations available from 1974 to 2001. Penn World Tables (Heston, Summers and Aten, 2006) provide GDP data from 1960 onwards. The complete data are available for a panel of 99 countries. The average length of a time series depends on the forecast length and ranges from 24 (one year after a banking crisis) to 15 observations (ten years after). There have been 89 banking crises within our sample. The majority of countries (56) had one banking crisis, 28 countries had no crisis, 13 countries had two, one country had three, and one had four banking crises. There is a negative correlation ($r = -0.35$) between the frequency of banking crises and a country's GDP level. Our assumption that the likelihood of a banking crisis is time-invariant is supported by the data: there is no significant time trend in the observed frequency of crises.

3.2 Regression results

In Table 1 we report estimates of the IRFs to a banking crisis from a variety of specifications both in levels of log GDP as well as GDP growth rates, with and without country fixed effects. The specification in growth rates is a constrained version of the model in levels, implying $\Sigma_r \alpha_{1r} = 1$. We estimate all specifications with the corrected and uncorrected versions of the local projections estimator. Every specification includes four lags of log GDP (or its growth rate) and of the banking crisis dummy ($R = L = 4$). We test these lag restrictions and find that our specifications pass them.

All specifications produce similar estimates of GDP loss within the first few years after the start of a banking crisis, but diverge thereafter. The differences between the estimates in different specifications of the medium- and long-run effects of banking crises can be explained by the two types of biases discussed above. The first, and most important, source of bias is the omission of banking crises happening within the forecast period (recall section 2.2). This omission yields a downward bias (in absolute value) of the estimated IRFs for longer forecast horizons. The difference in estimates between the specifications with and without banking crises leads increases with the length of the forecast period, as we showed in section 2.2, because the cumulative effect of those intermediate banking crises increases with k . Then there is the Nickell bias due to the combination of fixed effects and lags of the dependent variable. This bias can be appreciated by comparing the specifications in GDP levels

with and without fixed effects. It is not as important quantitatively as the local projections bias. Even in the growth rate specifications, where the Nickell bias is more pronounced, adding banking crises happening within the forecast horizon brings the estimates fairly close to those in our preferred specification (equation (8)). The last specification in Table 1, based on equation (9) where all lags of GDP have been recursively substituted in an attempt to eliminate the Nickell bias, gives only slightly larger, but less precise, IRF estimates.

Our estimation results imply that banking crises lead to large and prolonged GDP losses, and that recovery, if any, is a distant and uncertain prospect. Figure 5 plots the IRF point estimates and their the 95% confidence interval bounds for our preferred specification (equation (8)). GDP loss continues for seven years after the crisis with only a little recovery starting from the eighth year, far from sufficient to make up for the prolonged loss. The estimated cumulative GDP loss ten years after the crisis is around 9 percent. To put this estimate into perspective, consider that the median GDP growth rate observed in our sample is roughly 3% per year, so that the median economy would have grown by 34% over ten years. Thus, a single episode of a banking crisis would cost about a quarter of that economy's long-term growth potential.

As another illustration of the quantitative importance of banking crises, we calculate the share of residual variance in log GDP explained by these events, and report the results in Table 2. The first column in Table 2 reports the share in the current GDP's residual variance explained by one particular crisis which happened k years ago. The second column contains the shares in residual variance due to the entire history of crises between now and k years ago. Starting with 1% in the next year, a single banking crisis can explain up to about 2.5% of GDP's residual variance in the next several years. The history of banking crisis during the last 10 years explains just under 9% of the current (log) GDP's residual variance. For a relatively rare event such as a banking crisis, these shares are quite high.

To check the robustness of our results to a change in the identifying assumptions, we test whether a part of the observed negative effect of a banking crisis on GDP may in fact be a correction of excessive growth in the period of an expectations bubble typically preceding the crisis, an argument advanced in Jorda, Schularick and Taylor (2011). We do so by including in our regressions banking crises happening up to three years after the current GDP is observed. The results in Table 3 do not suggest that exuberant expectations play a significant part in shaping our IRF estimates, at least to the extent that these expectations cause surges in GDP. While the correlations between the current GDP and the

upcoming banking crisis event are indeed positive, they are not nearly significant.

Finally, as another extension to our results, we estimate IRFs on different parts of our sample to see how sensitive our results are to variations in local context. In particular, since banking crises sometimes coincide with war, social unrest, major political reform or other disturbing events which consequences for GDP are hard to isolate, our estimates may be vulnerable to a lack of appropriate controls. We therefore rerun our regressions (specification (8)) separately on subsamples of African countries known to be hard hit by these problems and the rest of the countries (excluding countries in transitions). There are 41 countries in the African subsample, and 54 countries in the non-Africa subsample (the remaining four countries, excluded from both subsamples, are transition economies). Table 4 and Figure 6, reporting the results, reveal that African countries suffer more profound GDP losses than other countries, but experience stronger recovery towards the end of the ten-year period. The IRFs for the African subsample are less precisely estimated, however, suggesting that the effects of banking crises for those countries are also more heterogeneous. GDP losses for the non-African sample are less severe, but the recovery is slower, too.

4 Concluding remarks

Solid empirical evidence is needed to inform the lively theoretical debate on the long-run impact of banking crises on GDP. Our study’s main contribution to the empirical literature on banking crises lies in the proposed improvement to Jordà’s (2005) local projections method of estimating the impulse response function of GDP to a banking crisis. This method is a good alternative to calculating the IRF analytically from an estimated dynamic model, since it relies far less on a particular specification of the dynamic model which may be incorrect, estimating instead a forecast equation in each year after the crisis. Yet, we have demonstrated, both theoretically and by stochastic simulations, the existence of a hitherto unknown bias in the IRF estimates to which the local projections method is vulnerable. Our proposed correction to this bias involves augmenting the GDP forecast equation with banking crises occurring within the forecast horizon. Our Monte Carlo simulation results for this correction show its effectiveness and exceptional robustness of its results even under severe misspecifications of the underlying dynamic model for GDP. One particularly interesting case of misspecification to which our method is robust is what Cai and Den Haan (2009) call “one-type-shock” model, whereby biased estimates of IRF are derived because of inability to distinguish between different components of GDP each of

which having its own dynamic properties (recall the last example at the end of section 2.1).

Applying our method to the data, we find, as Cerra and Saxena (2008) do, that GDP loss from a banking crisis is largely permanent. Our findings suggest that an average banking crisis may cause an GDP loss of around 9 percent over a period of ten years from its start. The consequences of banking crises vary by country. Thus, comparing the IRF estimation results on the two sub-samples of countries – Africa versus the other countries (excluding also 4 transition economies) – we find that in African countries banking crises are quite severe, consting 10-12 percent of GDP over the long run, whereas in other countries GDP loss is not as strong (6-7 percent). Even though the estimates vary by country subsample and are increasingly imprecise for more distant future, the upshot of our findings is that GDP is unlikely to return to its pre-crisis path.

5 References

- Alvarez, J., and Arellano, M. (2003), The Time Series and Cross-Section Asymptotics of Dynamic Panel Data Estimators, *Econometrica*, 71(4), 1121-1159.
- Campbell, J. Y., and Mankiw, N. G. (1987), Are output fluctuations transitory?, *Quarterly Journal of Economics*, 102(4), 857-880.
- Caprio, G., and Klingebiel, D. (2003), Episodes of Systemic and Borderline Financial Crises, <http://go.worldbank.org/5DYGICS7B0>.
- Cerra, V., and Saxena, S. C. (2008), Growth Dynamics: The Myth of Economic Recovery, *American Economic Review*, 98(1), 439-457.
- Chong, Y., Jorda, O., and Taylor, A. M. (2011), The Harrod-Balassa-Samuelson Hypothesis: Real Exchange Rates and their Long-Run Equilibrium, *International Economic Review*, forthcoming.
- Faust, J., and Wright, J.H. (2011), Efficient Prediction of Excess Returns, *The Review of Economics and Statistics*, 93(2), 647-659.
- Heston, A., Summers, R., and Aten, B. (2006), Penn World Table Version 6.2, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.
- Jorda, O. (2005), Estimation and Inference of Impulse Responses by Local Projections, *American Economic Review*, 95(1), 161- 182.
- Jorda, O., Schularick, M., and Taylor, A.M. (2011), The Harrod-Balassa-Samuelson Hypothesis: Real Exchange Rates and their Long-Run Equilibrium, *International Economic Review*, forthcoming.
- Judson, R. A., and Owen, A. L. (1999), Estimating dynamic panel data models: a guide for macroeconomists, *Economics Letters*, 65(1), 9- 15.
- Mood, A., Graybill, F., and Boes, D. (1974), *Introduction to the theory of statistics*, McGraw-Hill.
- Nickel, S. (1981), Biases in Dynamic Models with Fixed Effects. *Econometrica*, 49(6), 1417-26.

6 Appendix

Proof of proposition 1:

We have:

$$\begin{aligned}
y_t &= \alpha_{11}y_{t-1} + u_t, \\
\text{Var}(y_t) &\equiv \sigma_y^2 = \sigma^2 (1 - \alpha_{11}^2)^{-1}, \\
\text{Cov}(y_t, y_{t-1}) &= \sigma^2 \alpha_{11}^2 (1 - \alpha_{11}^2)^{-1}, \\
y_t &= \sum_{j=1}^{\infty} \alpha_{11}^{j-1} u_{t-j}, \\
v_{tk} &= \sum_{j=0}^{k-1} \alpha_{11}^{k-j-1} u_{t+j}.
\end{aligned}$$

The variance estimator for α_{11} reads:

$$\text{plim} [(T-1) \text{Var}(\hat{\alpha}_{11})] = \sigma^2 \sigma_y^{-2} = (1 - \alpha_{11}^2), \quad (11)$$

The variance of $\hat{\alpha}_{11}^k$ satisfies:

$$\begin{aligned}
\text{Var}(\hat{\alpha}_{11}^k) &= \text{E}(\hat{\alpha}_{11}^{2k}) - \text{E}^2(\hat{\alpha}_{11}^k); \\
\text{E}(\hat{\alpha}_{11}^k) &= \lim_{t \rightarrow 0} \frac{d^k}{dt^k} \exp \left[\alpha_{11}t + \frac{1}{2} \text{Var}(\hat{\alpha}_{11}) t^2 \right].
\end{aligned}$$

where we use the moment generating function of the normal distribution in the second line, see Mood, Graybill, and Boes (1974, p.540). We have:

$$\begin{array}{lll}
k & 1 & 2 & 3 \\
\text{E}(\hat{\alpha}_{11}^k) & \alpha_{11} & \text{Var}(\hat{\alpha}_{11}) + \alpha_{11}^2 & 3\text{Var}(\hat{\alpha}_{11})\alpha_{11} + \alpha_{11}^3 \\
\text{Var}(\hat{\alpha}_{11}^k) & \text{Var}(\hat{\alpha}_{11}) & 2\text{Var}(\hat{\alpha}_{11})^2 + & 15\text{Var}(\hat{\alpha}_{11})^3 + 36\text{Var}(\hat{\alpha}_{11})^2\alpha_{11}^2 + \\
& & 4\text{Var}(\hat{\alpha}_{11})\alpha_{11}^2 & 9\text{Var}(\hat{\alpha}_{11})\alpha_{11}^4
\end{array}$$

$$\begin{array}{lll}
k & 4 & 6 \\
\text{E}(\hat{\alpha}_{11}^k) & 3\text{Var}(\hat{\alpha}_{11})^2 + & 15\text{Var}(\hat{\alpha}_{11})^3 + 45\text{Var}(\hat{\alpha}_{11})^2\alpha_{11}^2 + \\
& 6\text{Var}(\hat{\alpha}_{11})\alpha_{11}^2 + \alpha_{11}^4 & 15\text{Var}(\hat{\alpha}_{11})\alpha_{11}^4 + \alpha_{11}^6
\end{array}$$

Substitution of the expression for $\text{Var}(\hat{\alpha}_{11})$ and taking limits yields:

$$\text{plim} [(T-1)\text{Var}(\hat{\alpha}_{11}^k)] = k^2 (1 - \alpha_{11}^2) \alpha_{11}^{2(k-1)}.$$

The variance of $\hat{\delta}_{11k}$ satisfies:

$$\begin{aligned}
&: \text{Var}(\hat{\delta}_{11k}) = (X'X)^{-1} X'VX (X'X)^{-1} \\
&= (T-k)^{-2} \sigma_y^{-4} \text{E} \left[(T-k) (y_t^2 v_{tk}^2) + 2 \sum_{j=1}^{k-1} (T-k-j) (y_t v_{tk} v_{t+j(k)} y_{t+j}) \right]
\end{aligned}$$

$$\begin{aligned}
&= \sigma_u^{-4} \frac{(1 - \alpha_{11}^2)^2}{T - k} \times \\
&\quad \mathbb{E} \left[\frac{(\sum_{r=-\infty}^{-1} \alpha_{11}^{-r-1} u_{t+r})^2 \left(\sum_{l=0}^{k-1} \alpha_{11}^{k-l-1} u_{t+l} \right)^2 +}{2 \sum_{j=1}^{k-1} \frac{T-k-j}{T-k} \sum_{r=-\infty}^{-1} \sum_{s=-\infty}^{j-1} \sum_{l=0}^{k-1} \sum_{m=j}^{k-1+j} \alpha_{11}^{2k-r-s-l-m-4} u_{t+r} u_{t+s} u_{t+l} u_{t+m}} \right] \\
&= \sigma_u^{-4} \frac{(1 - \alpha_{11}^2)^2}{T - k} \alpha_{11}^{2k} \times \\
&\quad \mathbb{E} \left[\sum_{r=-\infty}^{-1} \sum_{l=0}^{k-1} \alpha_{11}^{-2(r+l+2)} u_{t+r}^2 u_{t+l}^2 + 2 \sum_{j=1}^{k-1} \frac{T-k-j}{T-k} \sum_{r=-\infty}^{-1} \sum_{l=j}^{k-1} \alpha_{11}^{-2(r+l+2)} u_{t+r}^2 u_{t+l}^2 \right] \\
&= \frac{(1 - \alpha_{11}^2)^2}{T - k} \alpha_{11}^{2k} \sum_{r=-\infty}^{-1} \alpha_{11}^{-2(r+1)} \left[\sum_{l=0}^{k-1} \alpha_{11}^{-2(l+1)} + 2 \sum_{j=1}^{k-1} \frac{T-k-j}{T-k} \sum_{l=j}^{k-1} \alpha_{11}^{-2(l+1)} \right] \\
&= \frac{1}{T - k} \left[1 - \alpha_{11}^{2k} + 2 \sum_{j=1}^{k-1} \frac{T-k-j}{T-k} \left(1 - \alpha_{11}^{2(k-j)} \right) \right].
\end{aligned}$$

The second line uses $X'X = (T - k) \sigma_y^2$ and writes the covariance matrix $X'VX$ as the sum of the terms on the main diagonal and two times the $k - 1$ diagonal along both sides of the main diagonal that account for the MA part of the error term. The third line use the expression for σ_y^2 and expresses v_{tk} and y_t in terms of (lags of) u_t . The fourth line uses the fact that $\mathbb{E}(u_t u_{t+k}) = 0$ for $k \neq 0$, so that we can drop these terms. The fifth line takes the expectation. The final lines is straightforward algebra. Hence, we can write:

$$\begin{aligned}
\text{plim} \left[(T - k) \text{Var} \left(\widehat{\delta}_{11k} \right) \right] &= 1 - \alpha_{11}^{2k} + 2 \sum_{j=1}^{k-1} \left(1 - \alpha_{11}^{2(k-j)} \right) \\
&= 1 - \alpha_{11}^{2k} + 2(k - 1) - 2 \frac{\alpha_{11}^2 - \alpha_{11}^{2k}}{1 - \alpha_{11}^2}.
\end{aligned}$$

Proof of proposition 2:

Let e_{tk} be the estimation error in y_{tk} for country i , that is, the estimation error in the δ -coefficients times the corresponding explanatory variables. Consider the first order condition associated with the estimation of equation (5). Let I be the set of countries i that is hit by one banking crisis at time $S > k - 1$ (as before, we drop the suffix i of S); hence $d_S = 1$ and $d_t = 0$ for $t \neq S$. Let $-I$ the set of countries that is not hit by any banking crisis; hence $d_t = 0$ for all t for these countries. The condition for the fixed effect δ_{0ik} for country i reads $\sum_t (v_{tk} + e_{tk}) = 0$. By equation (7), the expectation of this condition over the error terms

u_t satisfies:

$$\begin{aligned} 0 &= \sum_{l=1}^{k-1} \gamma_{2l} + \sum_{t=1}^T \mathbb{E}[e_{tk}], \text{ for } i \in I, \\ 0 &= \sum_{t=1}^T e_{tk}, \text{ for } i \in -I. \end{aligned} \quad (12)$$

The condition for $i \in -I$ implies that the sum of the estimation errors for these countries is equal to zero. Hence, the fixed effect δ_{0ik} is estimated consistently, since the δ_{2lk} -parameters do not enter the expression for y_{tk} . The condition for $i \in I$ implies that the sum of the estimation errors must be equal to $\sum_{l=1}^{k-1} \gamma_{2l}$ for each country. Summing this equation over $i \in I$ yields $\sum_{i \in I} \sum_t e_{tk} = -I \sum_{l=1}^{k-1} \gamma_{2l}$.

The first order condition for δ_{2lk} for each $l = 1, L$ reads $\sum_i \sum_t (v_{tk} + e_{tk}) d_{t-l} = 0$. For countries without a banking crisis, $i \in -I$, $d_t = 0$ for all t . Hence, these countries do not contribute to the first order condition. For countries with a banking crisis, $i \in I$, the expected contribution to the first order condition is equal $\sum_t (v_{tk} + e_{tk}) d_{t-l} = v_{S+l,k} + e_{S+l,k}$. By equation (7), $\mathbb{E}[v_{S+l,k} d_S] = 0$, since $\mathbb{E}[u_t] = 0$ and $d_{S+k-1} d_S = 0$. Hence, the expectation of the first order condition for δ_{2lk} for $l = 1, L$ reads:

$$\begin{aligned} 0 &= \sum_{i \in I} \mathbb{E}[e_{S+l,k}] = I \left(\mathbb{E}[\widehat{\delta}_{0ik}] - \delta_{0ik} + \mathbb{E}[\widehat{\delta}_{2lk}] - \delta_{2lk} \right) \\ \mathbb{E}[\widehat{\delta}_{0ik}] - \delta_{0ik} &= - \left(\mathbb{E}[\widehat{\delta}_{2lk}] - \delta_{2lk} \right). \end{aligned} \quad (13)$$

Consider equation (12) for $i \in I$. $\sum_{i \in I} e_{tk} = 0$ for all L values of $t \in [S, S+L-1]$ by equation (13) and $\sum_{i \in I} e_{tk} = \sum_{i \in I} \left(\mathbb{E}[\widehat{\delta}_{0ik}] - \delta_{0ik} \right)$ for all other $T-L$ values of t . Hence:

$$(T-L) \left(\mathbb{E}[\widehat{\delta}_{0ik}] - \delta_{0ik} \right) = \sum_{l=1}^{k-1} \gamma_{2l},$$

satisfies the first order condition for $\mathbb{E}[\widehat{\delta}_{0ik}]$. The expression for $\mathbb{E}[\widehat{\delta}_{2lk}]$ follows from equation (13).

Figure 1: The probability limits of variances of the AR(1) process IRFs estimated analytically and through local projections

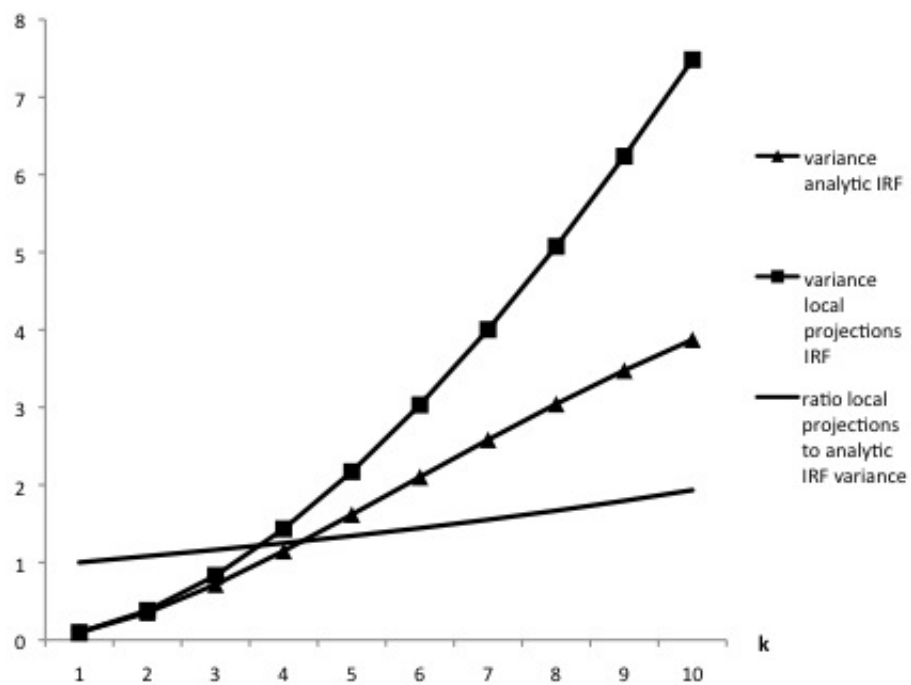


Figure 2: The IRFs from the AR(3) and AR(1) models

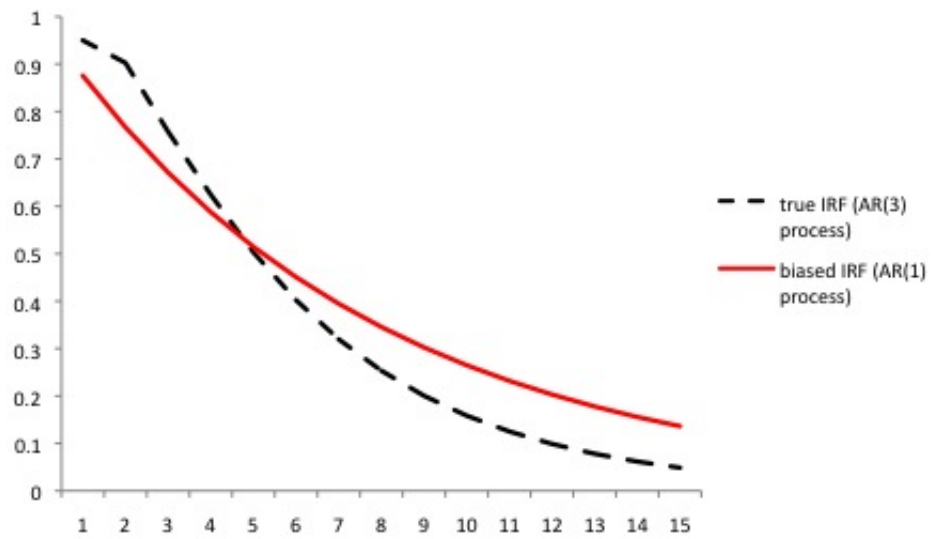


Figure 3: Monte Carlo simulation results for the analytical and local projections estimators for different specifications of the data generating process ($T = 30$).

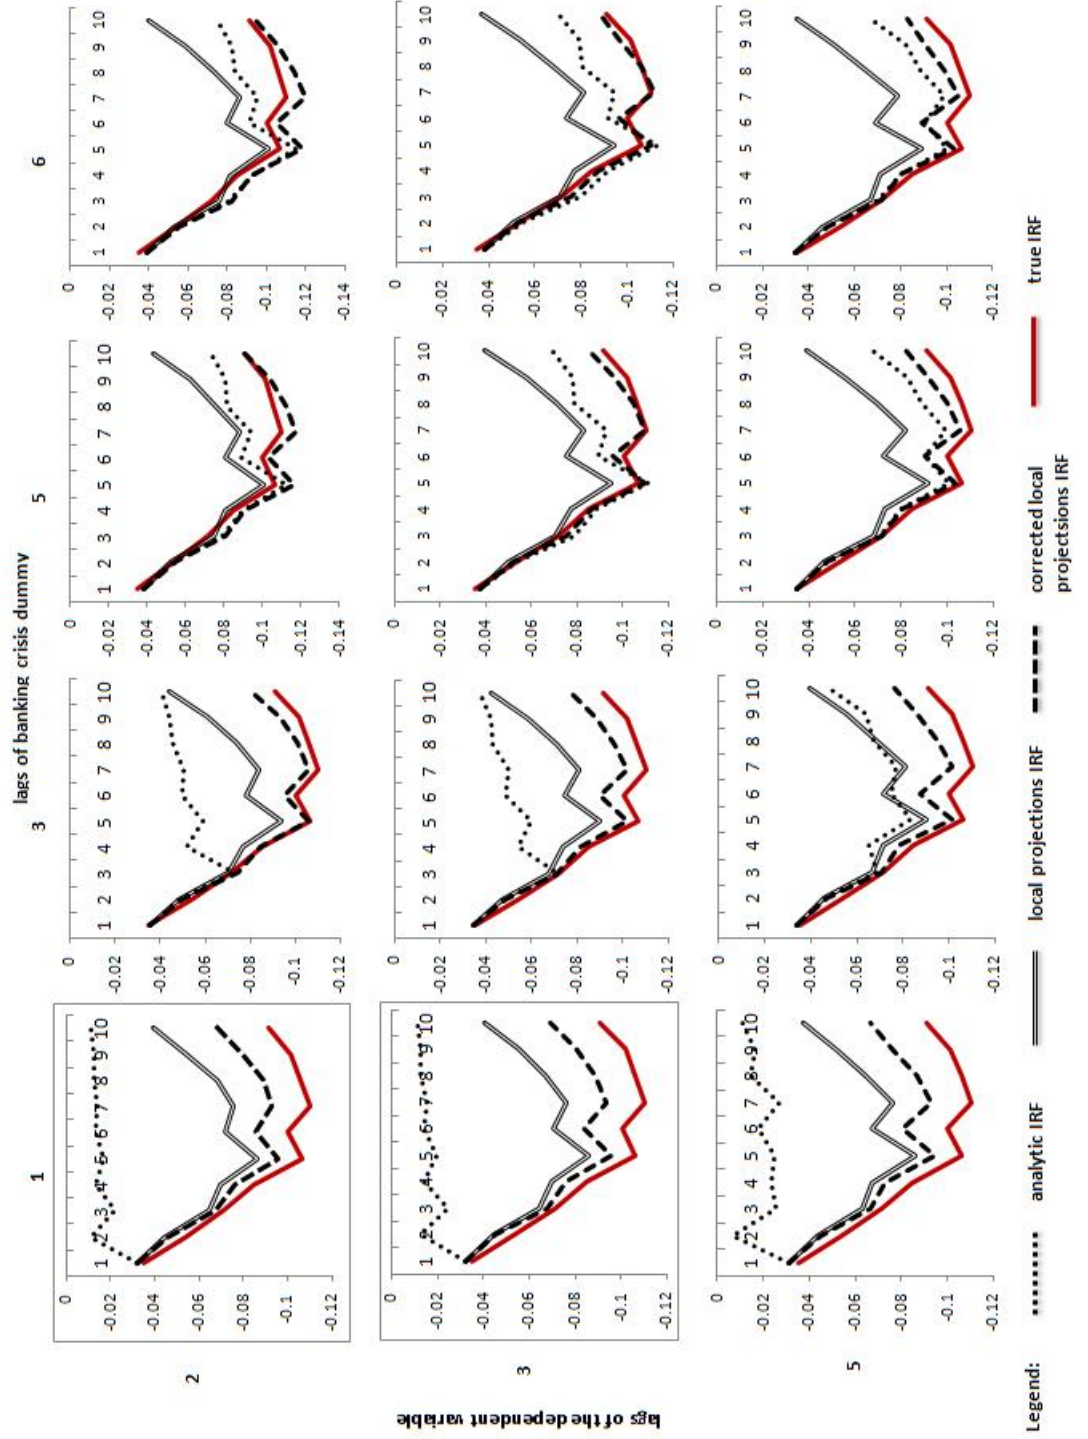
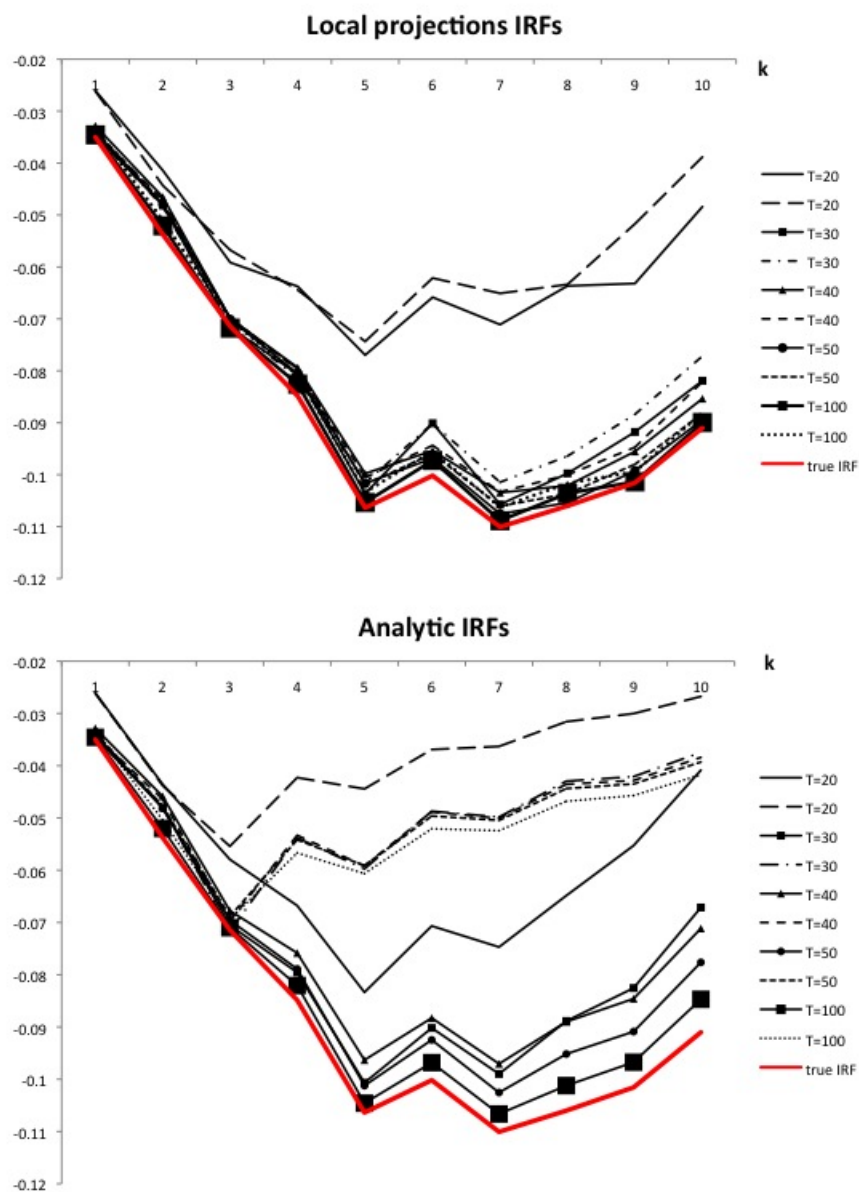
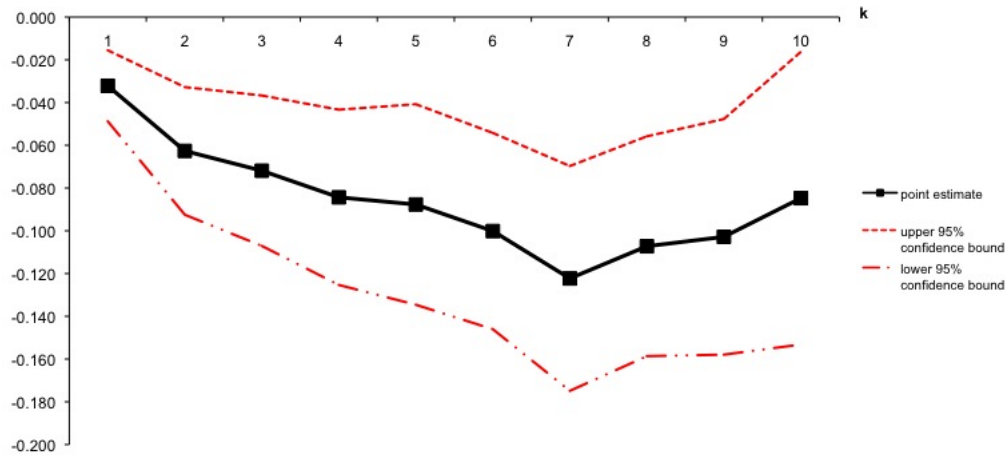


Figure 4: Simulated corrected local projections and analytical IRFs for different lengths of observations, T



Note: Solid lines denote IRFs from correctly specified model, and dashed lines denote IRFs from a misspecified model.

Figure 5: Predicted GDP loss from a banking crisis, and its 95% confidence interval.



Note: Estimates are based on corrected local projections estimator (equation (8))

Figure 6: Estimated IFRs by subsample of countries

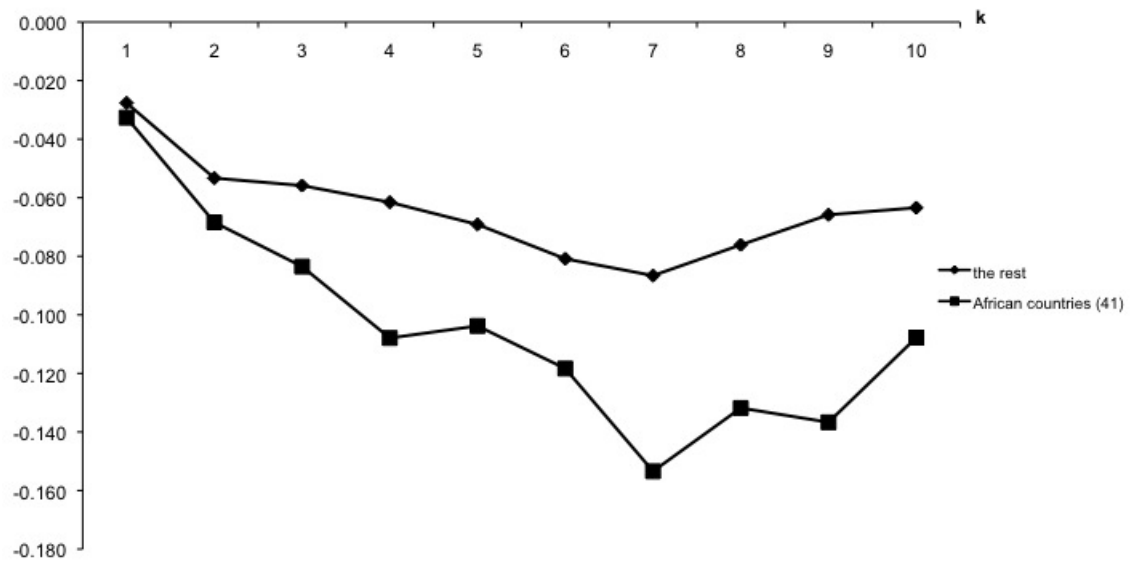


Table 1: Impulse response estimates for the effect of a banking crisis on GDP

k	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			eq. (1)	eq. (8)			eq. (9)
1	-0.035***	-0.035***	-0.032***	-0.032***	-0.033***	-0.034***	-0.050*
2	-0.066***	-0.069***	-0.058***	-0.063***	-0.062***	-0.065***	-0.063*
3	-0.071***	-0.079***	-0.060***	-0.072***	-0.067***	-0.074***	-0.091**
4	-0.077***	-0.090***	-0.064***	-0.084***	-0.074***	-0.088***	-0.106**
5	-0.069***	-0.088***	-0.056***	-0.088***	-0.066***	-0.087***	-0.102*
6	-0.069***	-0.096***	-0.057***	-0.100***	-0.069***	-0.099***	-0.109*
7	-0.084***	-0.117***	-0.069***	-0.122***	-0.084***	-0.120***	-0.138**
8	-0.058*	-0.097***	-0.041**	-0.107***	-0.051	-0.093***	-0.132*
9	-0.051	-0.089**	-0.030	-0.103***	-0.047	-0.091**	-0.122
10	-0.023	-0.063	-0.007	-0.085**	-0.021	-0.068	-0.093

Dep. var.: GDP

level/growth rate	level	level	level	level	gr.rate	gr.rate	level
Fixed effects	no	no	yes	yes	no	no	yes
Banking crisis	no	yes	no	yes	no	yes	yes
leads							

Note: from this table onwards, estimates significant at 1%, 5% and 10% levels are marked ***, ** and *, respectively.

Within-group serial correlation-robust standard errors for the local projections IRF estimates were calculated using option

`vce(cluster id)` in Stata.

Table 2: Share of residual variance in log GDP explained by banking crises

k	Single crisis	History of crises
1	0.010	0.010
2	0.014	0.019
3	0.014	0.027
4	0.019	0.038
5	0.022	0.050
6	0.027	0.063
7	0.028	0.072
8	0.023	0.081
9	0.022	0.087
10	0.018	0.087

Note: The share of residual variance explained by banking crises is calculated for each k as the ratio $\frac{r_{1,k} - r_{2,k}}{r_{1,k}}$, where $r_{1,k}$, $r_{2,k}$ are residual sums of squares (RSS). For the results in the second column, $r_{1,k}$ is the RSS of the autoregressive model without banking crises,

$$y_{t+k-1} = \delta_{0ik} + \delta_{0k}^* t + \sum_{r=1}^R \delta_{1rk} y_{t-r} + v_{tk}^1,$$

and $r_{2,k}$ is the RSS of the augmented local projections equation (8) from which our preferred IRF estimates are derived. For the results in the first column, $r_{2,k}$ is the same, and $r_{1,k}$ is the RSS of equation (8) with all banking crises happening *not* in year t are present. This ensures that we capture the share of variance explained by a single banking crisis.

Table 3: GDP growth before the start of a banking crisis

Years <i>before</i> crisis ($-k$)	Point estimate	Std. deviation
1	0.026	0.033
2	0.018	0.042
3	0.004	0.042

Note: results are based on specification (9) augmented with leads of the banking crisis dummy as follows:

$$y_{t+k-1} = \gamma_{0ik} + \gamma_{0!} t + \sum_{k=-3, k \neq 0}^{\infty} \gamma_{2k} d_{t-k} + v_t$$

The coefficients reported in the table for 1, 2 and 3 years before the crisis are those on $\gamma_{2,-1}$, $\gamma_{2,-2}$ and $\gamma_{2,-3}$, respectively.

Table 4: Impulse response estimates by subsample of countries

Specification	African countries (41)		The rest (54)	
	eq. (8)	eq. (9)	eq. (8)	eq. (9)
k				
1	-0.033**	-0.040	-0.028***	-0.024
2	-0.068**	-0.075	-0.053***	-0.021
3	-0.084***	-0.121	-0.056***	-0.041*
4	-0.108***	-0.124	-0.062***	-0.059**
5	-0.104**	-0.103	-0.069***	-0.094***
6	-0.118**	-0.138	-0.081***	-0.088***
7	-0.153***	-0.198	-0.087***	-0.088***
8	-0.132**	-0.195	-0.076***	-0.082***
9	-0.137**	-0.188	-0.066***	-0.061**
10	-0.108	-0.124	-0.063***	-0.071***