

Sander, Nicholas

Working Paper

The macroeconomic effects of portfolio equity inflows

Bank of Canada Staff Working Paper, No. 2023-31

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Sander, Nicholas (2023) : The macroeconomic effects of portfolio equity inflows, Bank of Canada Staff Working Paper, No. 2023-31, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2023-31>

This Version is available at:

<https://hdl.handle.net/10419/297416>

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.

The Macroeconomic Effects of Portfolio Equity Inflows

by Nick Sander

International Economic Analysis Department
Bank of Canada
NSander@bankofcanada.ca

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.



Acknowledgements

I would like to thank the Clausen Center, the UC Berkeley Library as well as Yuriy Gorodnichenko and Barry Eichengreen for providing funds needed to purchase the EPFR data. I am also grateful to Yuriy Gorodnichenko, Pierre-Olivier Gourinchas, David Sraer for their valuable comments and thoughts on this paper. I am also grateful to Emi Nakamura, Jon Steinsson, Amir Kermani, Julien Bengui, Geoffrey Dunbar, Ingomar Krohn, Kemal Ozhan and many others for their individual comments. I am grateful to my officemates, Marc Dordal-i-Carreras and Matthias Hoelzlein for spending far too much of their time discussing this with me and finally, to Ganesh Viswanath-Natraj for his time and effort spent helping me on this project. Thanks also to Jamie Porthiyas for excellent research assistance. The views in this paper are those of the author and do not necessarily reflect those of the Bank of Canada.

Abstract

I provide evidence that portfolio equity inflows can have expansionary effects on GDP and inflation if not offset by monetary policy. I use a shift-share instrument to estimate equity inflows based on plausibly exogenous timing of inflows into mutual funds with heterogeneous country portfolios. For countries with fixed exchange rates, GDP rises for at least two years following an exogenous inflow with a peak effect of 0.8 percent after 18 months. This is driven by rises in investment and exports, where the latter response is inconsistent with standard expenditure switching channel mechanisms. Non-fixing countries maintain GDP roughly at the same pre-shock levels but achieve this with higher interest rates.

Topics: International financial markets; International topics; Monetary policy; Business fluctuations and cycles

JEL codes: E32, F32, F44

Résumé

Nous démontrons que les flux entrants de portefeuille d'actions peuvent avoir des effets expansionnistes sur le produit intérieur brut (PIB) et l'inflation s'ils ne sont pas compensés par la politique monétaire. Nous utilisons un instrument d'analyse structurelle résiduelle pour estimer ces investissements en supposant une arrivée vraisemblablement exogène des capitaux dans des fonds communs de placement dont le portefeuille regroupe un ensemble hétérogène de pays. Dans les pays où les taux de change sont fixes, le PIB augmente pour au moins deux ans à la suite d'une entrée de capitaux exogène, laquelle contribue à cette augmentation jusqu'à hauteur de 0,8 % après 18 mois. Cela s'explique par des hausses des investissements et des exportations, phénomène qui, dans le cas des exportations, est incompatible avec les mécanismes standard de transfert des dépenses. Dans les pays où les taux de change sont flottants, le PIB se maintient sensiblement aux mêmes niveaux qu'avant le choc causé par l'entrée de capitaux exogène, et ce, grâce à un relèvement des taux d'intérêt.

Sujets : Marchés financiers internationaux, Questions internationales, Politique monétaire, Cycles et fluctuations économiques

Codes JEL : E32, F32, F44

1 Introduction

When foreign investors invest in a country, what are the effects on that country's macroeconomy? Almost all international macro models, by including a mechanism where some investors choose between domestic and international assets to determine exchange rates, implicitly take a position on the answer to this question. Typically the dominant mechanism is to emphasize expenditure switching mechanisms: capital inflows lead exchange rates (if floating) to appreciate, which in turn lowers net exports, which can lead to falls in real GDP.¹

However, empirically the correlation between gross capital inflows (foreign purchases of domestic assets) and GDP in the data tends to be positive – especially for extremely large and sudden capital inflow changes (Reinhart and Reinhart, 2008). Moreover, recent empirical work has suggested there may be a financial channel where additional capital inflows lead to a loosening of local financial conditions and possibly rises in GDP.² It is still not known which mechanism dominates at the macro level because identifying empirically the effect of capital flows at the aggregate level remains challenging.³ In this paper, I develop an instrument that can overcome this identification challenge for one type of capital inflows – portfolio equity inflows – and show that these inflows, if not offset by monetary policy, *raise* GDP. This suggests monetary policy likely should be *tightened* in response to exogenous capital inflows. These findings are in contrast to the predictions of most theoretical international macro models and suggest that these models may need additional mechanisms to expenditure-switching channels to match these findings.

Establishing the causal effect of capital inflows on macroeconomic outcomes is notoriously difficult because capital flows are themselves sensitive to local macroeconomic conditions. The literature since Fernandez-Arias (1996) tends to focus on two types of capital inflow drivers: “pull factors” representing shocks in the destination country and “push factors” representing shocks in the investor's country. Pull factors can reflect the strong fundamentals of the

¹In some cases where home and foreign goods are strong complements, GDP can rise following an (exogenous) capital inflow; however, such calibrations appear to be rare choices in practice. Examples where GDP would fall following capital inflows include Galí and Monacelli (2005), Gabaix and Maggiori (2015, p. 1396) (who both use Cobb-Douglas calibrations) and Itskhoki and Mukhin (2021).

²For example, Miranda-Agrippino and Rey (2020) and Kalemli-Özcan (2019) show that contractionary US monetary policy leads capital to flow to the US, tightening financial conditions elsewhere, Baskaya et al. (2017) show that banks lend more when capital inflows are high and Jotikasthira et al. (2012a) shows that outflows from mutual funds leads to fire sales and equity price declines.

³Two notable examples estimating macro effects of capital inflows are Blanchard et al. (2016) and Davis (2015), who both use instruments based on average global investor behavior or risk sentiment to isolate effects of capital inflows on recipient country economies. In addition, Pandolfi and Williams (2020), Broner et al. (2021) and Williams (2018) use the inclusion and removal of countries in global indexes to look at the effects of capital flows on financial and firm-level outcomes.

country “pulling” in capital from abroad. Economic reforms, strong economic growth and news about strong investment prospects are examples of factors that may lead investors to send their capital to a particular country; for the purpose of identifying the macroeconomic effects of capital flows, these are endogenous. Push factors may not directly be related to shocks in a particular country, but are still problematic from an identification perspective because they might be driven by shocks that have other worldwide macroeconomic effects – including on the destination country.⁴

The instrument I develop to address this challenge isolates small, frequent, idiosyncratic push factors into and out of a variety of countries globally. Using data from Emerging Portfolio Funds Research (EPFR), I exploit the institutional nature of mutual fund markets that offer financial services to less sophisticated investors: these investors select mutual funds and provide them their savings, while the mutual funds then decide where in the world to invest.⁵ There is considerable variation in the sets of countries each mutual fund invests in, which I exploit with a shift-share instrument. The “shares” component of the instrument is based upon the mutual funds’ portfolio held in a variety of countries around the world and is held constant in a given initial base period. Fixing each fund’s portfolio shares has the advantage of ignoring variations in capital inflows caused by mutual funds responding to news about the countries they invest in.⁶ I also show that the instrument ignores equity flows induced by mutual funds rebalancing their portfolios – selling assets that recently had above average returns to buy assets that recently had below average return in order to maintain a target level of exposure to all securities under management.⁷ Time variation in the instrument is driven by the “shifter”; I use a plausibly exogenous measure of “fund inflows” – fresh funds

⁴For instance, global risk aversion tends to rise in global downturns. Rises in global risk aversion may lead to capital flight from many countries, while the global downturn may independently cause GDP to fall in many countries (for example, through lower export demand). This would lead to a positive correlation between capital flows and GDP even if capital inflows had no direct effect on GDP. In addition Rey (2012) and Miranda-Agrippino and Rey (2020) argue investor sentiment and capital flows react to US monetary policy. However, US monetary policy shocks can affect global GDP through trade and other channels independently of their effect on global financial flows. Both examples show that global push factors may be poor sources of variation with which to construct an instrument for capital flows.

⁵Figure 4 in Miao and Pant (2012) shows that the predicted flows from the EPFR data have a considerably strong correlation with the time series of capital inflows into different regions of the world.

⁶Jotikasthira et al. (2012b) and Raddatz and Schmukler (2012) show using the same EPFR data used in this paper that clients are important actors during global “risk on” and “risk off” periods and can cause fire sales by mutual funds in equity markets around the world. Raddatz and Schmukler (2012) find that mutual funds sell more of some countries’ assets than others in “risk off” situations, which could be related to information they have about different countries. Keeping the portfolio shares fixed in a base period prevents this information being embedded in my proposed instrument.

⁷The flows induced by rebalancing go in the opposite direction to flows induced by mutual funds increasing their portfolio shares invested in high-return countries but are still endogenous.

invested into each mutual fund by their clients.

The choice of this fund inflow shifter is important as there may be two types of clients selecting mutual funds: news-insensitive clients and news-sensitive clients. News-insensitive clients are ideal for identification because their decision to pick one mutual fund over another is arguably independent of both the countries each mutual fund invests in (and of the investment prospects of each of those countries). However, news-sensitive clients may direct funds towards mutual funds investing in the countries with the strongest growth prospects and violate the exclusion condition. To ensure I focus on variation in client behavior that is likely to be exogenous, I focus on systematic but news-insensitive client investment behavior – selecting a mutual fund based on its past return history.⁸ This is a news-insensitive method of selecting a mutual fund because equity markets are well known to process news extremely quickly, making past returns a poor method of predicting future returns. However, while equity markets may reflect news quickly, macro variables may reflect news content gradually over time – making past equity returns predictive of future macroeconomic outcomes. To account for this, I construct a leave-out prediction of client inflows: for mutual funds investing in a particular country – say Argentina – I predict client inflows into that mutual fund using lagged returns these mutual funds made in all countries other than Argentina. The identification assumption then becomes that lagged returns in countries other than Argentina do not predict future macroeconomic outcomes in Argentina (and similarly for every country in the sample). I conduct many tests and robustness checks to confirm that these predicted client inflows are not systematically directed to (or away from) mutual funds exposed to countries known to be experiencing good shocks.

With this instrument, I estimate two sets of impulse responses using Jordà (2005a) local projections. In my first set of results, I trace out the dynamic effects of a 1 percent of GDP exogenous gross equity inflow for a pooled set of fixing and floating countries. Following an exogenous rise, gross equity inflows return to their pre-shock level after a year. Yet, the dynamics of overall net inflows can be complex – equity inflows from foreigners tends not to be offset by domestic residents buying equities abroad, but they tend to lead to further gross portfolio *debt* inflows (debt purchases by foreigners), perhaps by relaxing firm borrowing constraints and leading to more foreign investment in the country. However, *net* debt inflows (inflows by foreigners less outflows by domestic residents) respond negatively, suggesting that exogenous equity inflows by foreigners lead to increases in domestic savings abroad. When

⁸Many papers, starting with Zeckhauser et al. (1991), document that clients use past returns (along with other metrics) to select mutual funds (see Elton and Gruber, 2013). This finding is confirmed with my mutual fund data in Appendix A.3 and Figure 19.

looking at all inflows, a similar pattern prevails – net inflows are small and potentially negative despite being initiated by foreigners investing into a country.

Despite there being little effect on overall net inflows, I find that the overall macroeconomic effect can be large – particularly for fixed exchange rate economies.⁹ Theory suggests that *relative to non-fixing exchange rate countries*, those with fixed exchange rates should have looser monetary policy in response to an exogenous capital inflow, which would lead to higher GDP and inflation, lower imports and higher exports. This is exactly what I find.

However, while expenditure-switching mechanisms explain the *differences* in impulse responses between fixing and non-fixing countries, it cannot explain the *level* impulse responses. If expenditure switching were the dominant mechanism, then, provided the exchange rate were allowed to appreciate after an exogenous capital inflow shock, we would expect net exports to fall and this to put downward pressure on GDP and inflation. Expansionary monetary policy could offset this downward pressure on GDP and/or inflation, but in that case we should see falls in interest rates for non-fixing countries (and larger falls in interest rates for fixing countries). Instead my empirical impulse responses show that short-term interest rates *rise* for both fixing and non-fixing countries, with steeper rises for non-fixing countries. With these higher interest rates, non-fixing countries keep real GDP near to pre-shock levels, whereas fixing countries experience large rises in real GDP, peaking at around 0.8 percent after 18 months. This rise in real GDP is driven by rises in investment and exports, while for non-fixing countries exports *rise* slightly. This suggests an additional mechanism to expenditure switching accompanying equity inflows that lowers firm costs – particularly of exporting firms and is consistent with the findings of Kalemli-Ozcan et al. (2021) who show that firms in emerging economies borrow more when exchange rates appreciate – particularly firms already with high leverage.¹⁰

Overall these findings suggest that equity inflows are expansionary – even in non-crisis times. To the extent that these findings are representative of all types of inflows, this reinforces the view that global “risk-off” episodes may be contractionary – the loss of financial flows likely dominates any expenditure-switching channels from exchange rate depreciations. This has important implications for central banks in small open economies similar to those discussed in Rey (2012) and Gourinchas (2018) – changes in US monetary policy, by slowing capital flows to other countries, is likely contractionary everywhere (despite the center country’s currency

⁹I split the sample into countries with fixed exchange rate regimes and those with non-fixed regimes using the de facto classification of Ilzetzi et al. (2017).

¹⁰These findings are also consistent with Williams (2018), Pandolfi and Williams (2020) and Broner et al. (2021), who find that financial constraints are relaxed for firms when sovereign debt inflows into a country increases.

appreciating). By comparing outcomes between fixing and non-fixing countries, my results suggest that domestic monetary policy still works as conventionally understood – a *loosening* of local monetary policy can stabilize local GDP in reaction to monetary tightening in center countries.

This paper contributes to several strands of literature. First, it contributes to the literature on understanding the macroeconomic effects of capital flows. Previous work has focused on analysis of extreme events (Forbes and Warnock, 2012; Forbes and Warnock, 2014; Reinhart and Reinhart, 2008; and Mendoza, 2010), identification from entry of a country into global indexes (Williams, 2018; Broner et al., 2021; and Pandolfi and Williams, 2020) and identification using more frequent business cycle variation (Davis, 2015; Blanchard et al., 2016; and Cesa-Bianchi et al., 2018). This paper complements and extends these other works by bringing a new identification method that matches the credibility of identification based on entry of countries into global indexes but that also can be implemented on a large sample of countries to estimate global average treatment effects and to compare the effects on countries with fixed versus floating exchange rates. In particular, relative to previous work using frequent business cycle fluctuations as an instrument, my approach allows for the inclusion of region-time fixed effects. Region-time fixed effects are a crucial control as they absorb the main effects of global push factors, such as investor sentiment. By being specific to the UN subregion each country is in, they can also absorb the differences in sensitivity of some regions of the world to global investor sentiment.

Second, this paper complements the rich empirical literature studying the many microeconomic mechanisms by which capital flows can affect firms and households in different countries (a small subsample includes Kalemli-Ozcan et al., 2021; Salomao and Varela, 2022; Gopinath et al., 2017; and Gyöngyösi and Verner, 2020). By estimating the macro impacts of equity inflows, this paper can verify that some form of the financial frictions uncovered in these papers affects the macro-level dynamics of the response of countries to capital inflows.

Third, it complements the literature documenting the drivers of capital flows and arguing that global push factors are extremely important in determining when capital flows and when it dries up (see Koepke, 2015 for an excellent summary as well as Rey, 2012; Maggiori et al., 2020; and Coppola et al., 2021). This includes papers showing that global mutual funds affect the dynamics of capital flows using the same EPFR data I use (examples are Raddatz et al., 2017; Converse et al., forthcoming; Jotikasthira et al., 2012b; and Bacchetta et al., 2020). This literature has shown that mutual fund clients are often a source of change in global risk sentiment, and that client redemptions can lead to fire sales and equity price declines in countries that the affected mutual funds invest in. This paper contributes to this literature

by showing that capital flows also have macroeconomic effects and that EPFR data can be useful to construct an instrument for overall portfolio equity inflows.

Finally, this paper contributes to the literature on shift-share instruments (first proposed by Bartik, 1991) in the context of samples with a large time dimension.¹¹ Recent work by Goldsmith-Pinkham et al. (2020), Borusyak et al. (2021) and Adão et al. (2019) have illuminated two different identification assumptions needed to ensure shift-share designs work where the appropriateness of each identification assumption depends on the empirical context. I discuss that in samples covering more time periods, it is important to ensure that the results are robust to the choice of when the initial shares are chosen.

This paper is structured as follows: Section 2 describes the instrument, states the identification assumption and the second stage regressions undertaken; Section 3 describes the data; Section 4 shows the first stage results and presents tests to detect if the exclusion condition may be violated; Section 5 presents the main results; and Section 6 concludes.

2 Methodology

2.1 Clarifying Capital Inflows Concepts

In essence there are four types of activities captured by capital inflows: accumulation of foreign assets by local residents, sales of foreign assets by local residents, accumulation of liabilities held by foreigners and reductions of liabilities held by foreigners.

Traditionally, these are grouped into “gross” inflow or outflow categories, as detailed in Table 1: whenever residents of country i accumulate or reduce claims on entities or residents in j , that is referred to as a gross *outflow* from i and a gross *inflow* into j . Under this definition, gross inflows into j can be negative if foreign holders of claims on j sell them (or the claim is paid down by i). Similarly, gross outflows from j can be negative if residents in j sell their foreign assets and repatriate their savings.

Further note that inflows need to represent a flow of savings across a country border and thus capital gains or any asset revaluations are not considered inflows. However, in the traditional approaches used, re-invested dividends and interest income are considered capital inflows. Any dividends received and repatriated back into the owner’s country of residence have net 0 effect on capital inflows and outflows because the claim created by the payment

¹¹Jaeger et al. (2018) discuss some issues that can come into play when the time dimension is short but covers long time periods – such as 10 years. Without controlling for lags, shift-share instruments can estimate combinations of short-run and long-run elasticities. This paper, by contrast, uses more frequent data, which allows the inclusion of many lags and dynamic estimation to separate short-run and long-run effects.

Table 1: What transactions are recorded as inflows or outflows

<i>Activity</i>	<i>Inflow/Outflow in i</i>	<i>Inflow/Outflow in j</i>
Residents in <i>i</i> buy assets in <i>j</i>	<i>positive gross outflow from i</i>	<i>positive gross inflow into j</i>
Residents in <i>i</i> sell their assets in <i>j</i>	<i>negative gross outflow from i</i>	<i>negative gross inflow into j</i>
Residents in <i>j</i> buy assets in <i>j</i>	<i>positive gross outflow from i</i>	<i>positive gross inflow into j</i>
Residents in <i>j</i> sell their assets in <i>j</i>	<i>negative gross outflow from i</i>	<i>negative gross inflow into j</i>
Assets held by <i>i</i> in country <i>j</i> have capital gains	0	0
Residents in <i>i</i> reinvest their dividends from assets in <i>j</i>	<i>positive gross outflow from i</i>	<i>positive gross inflow into j</i>
Residents in <i>i</i> repatriate their dividends from assets in <i>j</i>	0	0

of dividends is removed once the dividend cash payment is received and repatriated.

In this paper I will focus on understanding the effects that exogenous gross *portfolio equity* inflows have on a recipient economy. While interesting, the effects of gross *outflows* on an economy are more difficult to ascertain because it is considerably more difficult to isolate exogenous variation in outflows.¹²

The definition of gross inflows into a country will be nominal – it is the amount spent by foreigners purchasing domestically listed equities. It is often noted that in a given market for every buyer there is a seller, and as such there are no inflows into markets on aggregate. Gabaix and Koijen (2022) challenge this view and propose a way to construct a theoretically consistent aggregate inflow measure based on the marginal propensity of different market participants to purchase additional securities in that market. Unfortunately, constructing this measure for the full set of countries in the sample is infeasible – flow of funds data are required. As such, the inflow measures used in the paper can be thought of as a reallocation of ownership of domestic equities towards foreigners – whether any effects estimated is due to some aggregate inflow effect á la Gabaix and Koijen (2022) or a reallocation of ownership – perhaps to investors with different risk sentiment or governance approaches – will remain unclear.

2.2 Instrument

There are two endogeneity issues to be resolved to estimate the causal effect of capital inflows on macro outcomes: good shocks or good news may attract investors and lead to capital inflows (endogenous pull factors) and global booms may lower global risk sentiment and lead to more capital flows between countries.¹³ In both cases, even if capital inflows

¹²Any exogenous movement in outflows would need to be unaffected by shocks local to the investor's country and to shocks to the investors' net worth. Inflows, by contrast, need only satisfy the former condition.

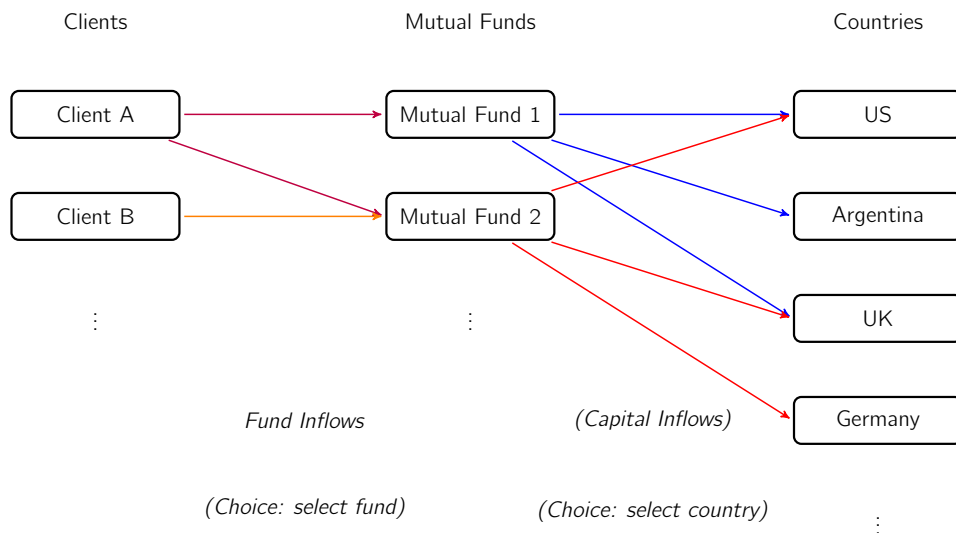
¹³To the extent that good news is capitalized into asset prices, these news shocks can also lead to *negative* inflows through rebalancing effects. This possibility is discussed later in Subsection 2.3.

had no economic effect on the recipient country's economy, we could easily see a positive correlation in the data due to these effects.

Given this, the ideal instrument would purge as much as possible the effects of country-specific news and global push factors from the predicted capital inflows to various countries. To some extent this can be done through controls, but mainly my approach is to take advantage of the role of mutual funds in allocating equity flows across countries. This market features a delegation structure where mutual funds allocate capital to countries and their clients simply select a mutual fund to invest in (as shown in Figure 1). If mutual funds are primarily the agents that respond to news about which countries are the best to invest in, then a plausibly exogenous instrument would fix the portfolios of mutual funds and leverage the behavior of clients to predict which mutual funds get money, and therefore which countries get capital inflows.

The shift-share instrument adopted in this paper ($B_{i,t}$) does exactly that: it predicts inflows into a country i at time t using the product of a fixed portfolio share and (a prediction of) inflows coming into each mutual fund. This instrument can have considerable relevance because mutual funds specialize by only investing in a very small number of countries with considerable variety in the set of countries chosen (see Figure 10). Thus a reallocation of capital invested in two random mutual funds can lead to a vastly different pattern of investment globally.

Figure 1: Delegation structure of mutual fund Investing



Specifically, the instrument is given by:

$$B_{i,t} = \sum_{k \in \mathcal{K}^{MF} \setminus \mathcal{K}^{MF}(i)} \underbrace{\frac{A_{i \leftarrow k, t_0}}{A_{k, t_0}}}_{\text{Shares}} \underbrace{\widehat{\text{Fund Inflows}}_{k,t}}_{\text{Shifter}} \quad (1)$$

where $A_{i \leftarrow k, t}$ is holdings of mutual fund k in assets listed in country i at time t ; \mathcal{K}^{MF} is the set of all mutual funds in the data; and $\mathcal{K}^{MF} \setminus \mathcal{K}^{MF}(i)$ is the set of all mutual funds in the data not domiciled in country i .

The portfolio share $A_{i \leftarrow k, t_0} \setminus A_{k, t_0}$ is fixed in some base period $t_0 < t$ so that portfolio adjustments in response to news released at period $t > t_0$ is not included in the instrument. The time variation in the instrument is generated by (a prediction of) the amount of new money clients invest in each mutual fund. These predicted “fund inflows” coming into each mutual fund is then apportioned into countries based on the t_0 portfolio shares of each mutual fund.

For reasons that will be discussed later in this section, it will be sensible to make some refinements to the exact form of the shifter (hence the notation above of $\widehat{\text{Fund Inflows}}_{k,t}$ to emphasize that the shifter I will use will be a refinement of overall inflows into mutual fund k : $\text{inflows}_{k,t}$).

2.3 Unpacking the Instrument

Shift-share instruments have received renewed attention recently. This literature shows that shift-share instruments can achieve identification in two ways – random assignment of the shares á la Goldsmith-Pinkham et al. (2020) or random assignment of the shifters conditional on the shares (Borusyak et al., 2021; Adão et al., 2019). Mathematically, both approaches are equally valid to achieve identification; however, the plausibility of each identification assumption depends on the economic context. For understanding the effects of portfolio equity inflows, identification is better thought of as coming from the random assignment of the shifters (client inflows into a mutual fund) conditional on the mutual funds’ investment behavior (the shares).¹⁴ To make this clear, in this subsection I will use some simple definitions and accounting relationships to unpack how the instrument operates in this context and be precise about the identification assumption being applied. A more mathematically rigorous approach is discussed in Appendix B.

¹⁴The Goldsmith-Pinkham et al. (2020) exclusion condition in this context would be that mutual fund portfolio shares have nothing to do with the (future) economic performance of countries. This is highly implausible.

First consider the following (slightly atypical) definition of gross capital inflows:

$$\text{inflows}_{i,t} \equiv d\text{GFL}_{i,t} - \underbrace{r_{i,t}\text{GFL}_{i,t}}_{\text{Capital Gains}} \quad (2)$$

Capital inflows represent the accumulation of claims by foreigners on domestic residents or institutions. This equals the difference between the growth in gross foreign liabilities (GFL) – the value of all claims owed to foreigners – in excess of the amounts earned by foreigners from holding these claims. To be consistent with traditional measures of capital inflows, $r_{i,t}$ should exclude dividends and contain *only* revaluations; however, to focus on how the instrument might achieve identification I will refer to $r_{i,t}$ as the return on equity.¹⁵

Letting \mathcal{K} denote all asset holders and $\mathcal{K}\setminus\mathcal{K}(i)$ denote all asset holders domiciled outside country i , then GFL can be disaggregated according to the following equation:

$$\text{GFL}_{i,t} = \sum_{k \in \mathcal{K}\setminus\mathcal{K}(i)} \omega_{i \leftarrow k,t} A_{k,t} \quad (3)$$

where the notation $\omega_{i \leftarrow k,t}$ refers to the share of agent k 's assets ($A_{k,t}$) invested in country i at time t . Taking the total derivative of this equation gives:¹⁶

$$d\text{GFL}_{i,t} = \sum_{k \in \mathcal{K}\setminus\mathcal{K}(i)} d\omega_{i \leftarrow k,t} A_{k,t} + \omega_{i \leftarrow k,t} dA_{k,t} \quad (4)$$

This equation states that GFL can increase for two reasons: investors with a fixed set of assets tilt their portfolios toward i and away from some other country ($d\omega_{i \leftarrow k,t} > 0$), or investors already exposed to i have more assets to invest ($dA_{k,t} > 0$). Note that returns earned in country i are reflected in this equation through $dA_{k,t}$.

For the remainder of this discussion, all asset holders will be assumed to be mutual funds. Given this, I use the concept of “fund inflows” (notation: $\text{inflows}_{k,t}$) as the mutual fund equivalent of the country inflows:

¹⁵The inflow data used in this paper is constructed according to the standard approaches and therefore treats $r_{i,t}$ in Equation 2 as capital gains/revaluations.

¹⁶It is without loss of generality to ignore changes in the numbers of investors. One could instead define that the set $\mathcal{K}\setminus\mathcal{K}(i)$ includes all *potential* investors of country i and set $\omega_{i \leftarrow k,t} = 0$ for those not currently investing in i .

$$\text{inflows}_{k,t} \equiv dA_{k,t} - r_{k,t}A_{k,t} \quad (5)$$

Putting Equations 2 to 5 together produces the following exact decomposition of country-level inflows ($\text{inflows}_{i,t}$):¹⁷

$$\text{inflows}_{i,t} = \sum_{k \in \mathcal{K} \setminus \mathcal{K}(i)} \left(\underbrace{d\omega_{i \leftarrow k,t} A_{k,t}}_{\text{Portfolio Adjustment}} - \underbrace{(r_{i,t} - r_{k,t})\omega_{i \leftarrow k,t} A_{k,t}}_{\text{Portfolio Rebalancing}} + \underbrace{\omega_{i \leftarrow k,t} \text{inflows}_{k,t}}_{\text{Allocation of Fund Inflows}} \right) \quad (6)$$

There are three terms driving inflows: first there is portfolio adjustment coming from a mutual fund deciding to increase its holdings of country i 's assets holding fixed the fund's assets under management $A_{k,t}$. The second term reflects portfolio rebalancing required when returns in country i are different from the average return each fund earns. Without this adjustment, if country i 's return is above the average return of a fund k , then the share of assets of that fund held in country i (in dollar terms) would increase. To keep the portfolio share invested in country i fixed, the mutual fund would need to sell some of its holdings in country i and reallocate them to other countries. The final term reflects what happens when the mutual fund grows larger by attracting more money from clients (either existing or new). In this case, the mutual fund allocates this money to countries in proportion to its current portfolio shares.

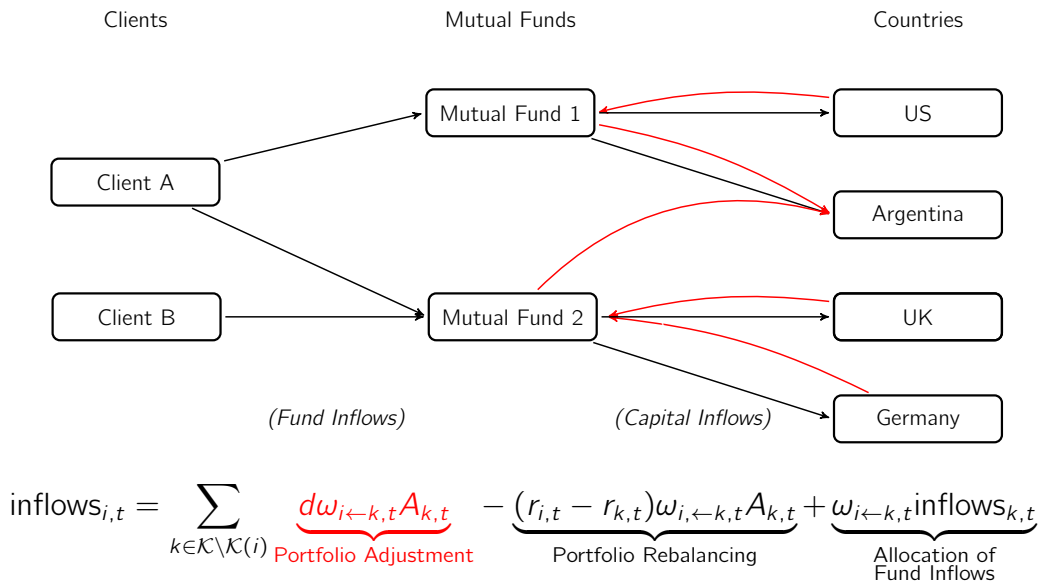
By comparing Equation 1 to Equation 6, it is clear that the instrument is solely related to the third term.¹⁸ This is deliberate because the first two terms are likely to be contaminated by news about country i 's growth prospects. To see why, consider as a motivating example an announcement by the Argentinean government of a series of policy reforms that are expected to (and will) raise Argentina's economic growth in the future. One might expect that actively managed mutual funds would respond to this news by tilting their portfolios toward Argentina. Figure 2 shows an example of this with two mutual funds. Prior to the news, one fund invests in Argentina and the US and the other fund invests in the UK and Germany. These initial portfolios are shown in black with straight lines. Upon receiving the news, the first mutual

¹⁷While this decomposition is exact, the reader should note that this is with the caveat that characterization of direct investors in country i as receiving inflows from some other agent is not a sensible way to characterize all investors in country i . However, this is a coherent description of the subset of investors included in the construction of the instrument.

¹⁸The only differences between the two expressions are the timing of the portfolio shares: $\omega_{i \leftarrow k,t}$ in Equation 1 versus $\omega_{i \leftarrow k,t_0}$, $t_0 < t$ in the instrument construction and the fact that fund inflows are predicted in the instrument.

fund might like to increase its portfolio holdings in Argentina and does so by selling holdings in its US portfolio. The second fund, while not having any holdings in Argentina, might decide to start investing in Argentina. These portfolio adjustments are shown in Figure 2 by the red arrows (with curved lines). Note that these adjustments occur holding fixed the inflows from clients. These are all endogenous – these portfolio adjustments are in response to news about a policy reform in Argentina, and leaving this variation out of the instrument seems sensible.

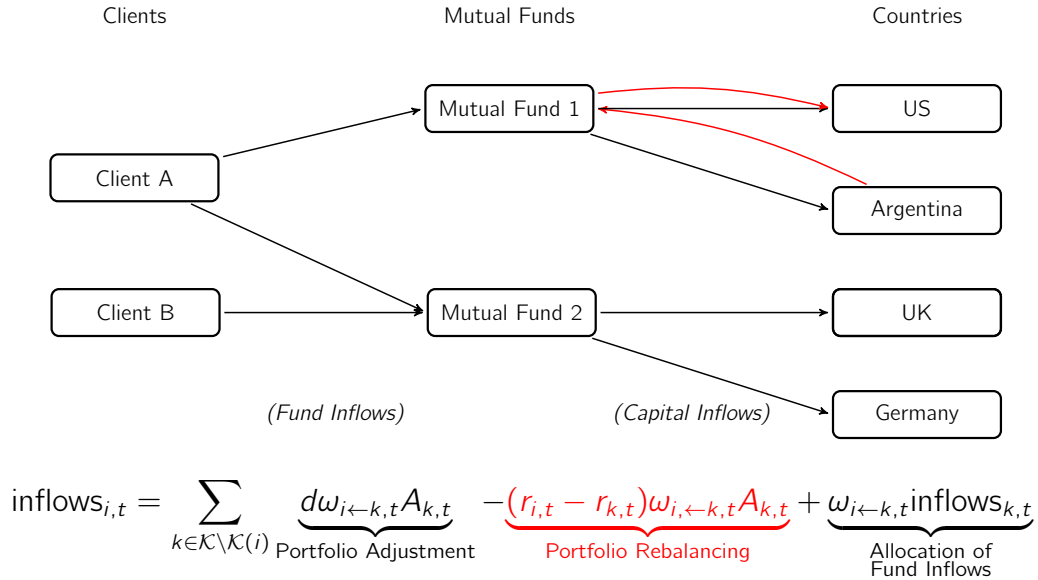
Figure 2: Portfolio adjustment towards Argentina due to expectations of high future growth



It might appear that for passive (index) funds, focusing on client inflows is less important because the portfolio adjustment effects shown in Figure 2 are absent. Unfortunately, the portfolio rebalancing term in the equation applies to index funds and is equally problematic. Figure 3 shows the portfolio-rebalancing effects (with red curved lines) of the same policy announcement in Argentina. Upon announcement, it would be reasonable to expect Argentina's equity and bond markets to rally; raising returns for mutual funds with existing investments in Argentina. Unless matched with higher returns in other countries, any index fund with an investment in Argentina (shown as "Mutual Fund 1" in Figure 3) will have a higher share of their portfolio invested in Argentina. To correct for this, Mutual Fund 1, for example, will need to sell some of its Argentinean assets and purchase additional US assets. The flow of capital this creates goes in the opposite direction to the portfolio adjustment term (money flows out of Argentina rather than in) but, as the cause of these flows was news about Argentina, these

negative inflows still constitute variation that a suitable instrument should ignore.

Figure 3: Portfolio rebalancing due to high returns earned from Argentina



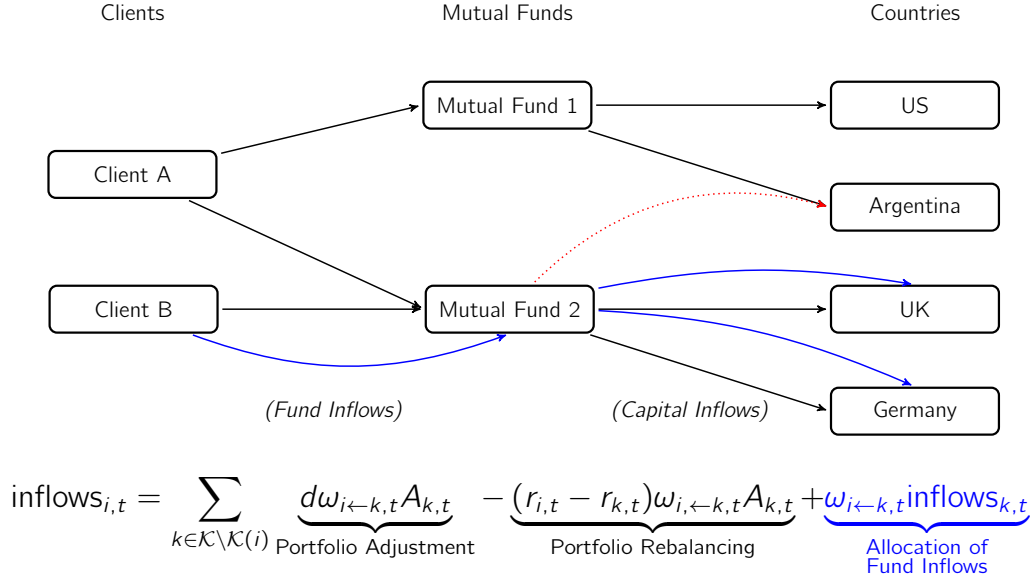
The final term is the allocation by each fund of inflows from clients towards the countries the mutual fund currently invests in. The instrument by contrast allocates these inflows towards countries based on *previous* countries the fund invested in. This is shown in Figure 4 with blue curved lines, where Client B receives a wealth shock that she invests in Mutual Fund 2 and the instrument predicts that this wealth will be allocated into the UK and Germany. Furthermore, if this inflow occurred after the hypothetical policy announcement discussed above, and Mutual Fund 2 responded by starting to invest in Argentina, the instrument would ignore this development and keep the portfolio shares fixed at their initial level.

2.4 Capturing Exogenous Variation in Client Inflows

Intuitively, if clients are allocating money to mutual funds for reasons unrelated to news about the country each fund invests in, then mutual fund inflows can be considered exogenous and the causal effects of capital inflows identified.¹⁹ However, one could imagine that clients are

¹⁹There has been recent work on shift-share designs, such as Goldsmith-Pinkham et al. (2020), Borusyak et al. (2021), Jaeger et al. (2018) and Adão et al. (2019), who bring to light the different identification assumptions one can rely on in shift-share designs. Moreover, this literature shows that the appropriateness of these various assumptions is context dependent. In this paper, the identification assumption closely resembles that discussed

Figure 4: Relevance of the instrument is driven by client-level variation



themselves behaving in similar ways to mutual funds and are therefore (directly or indirectly through rebalancing) responding to country news.

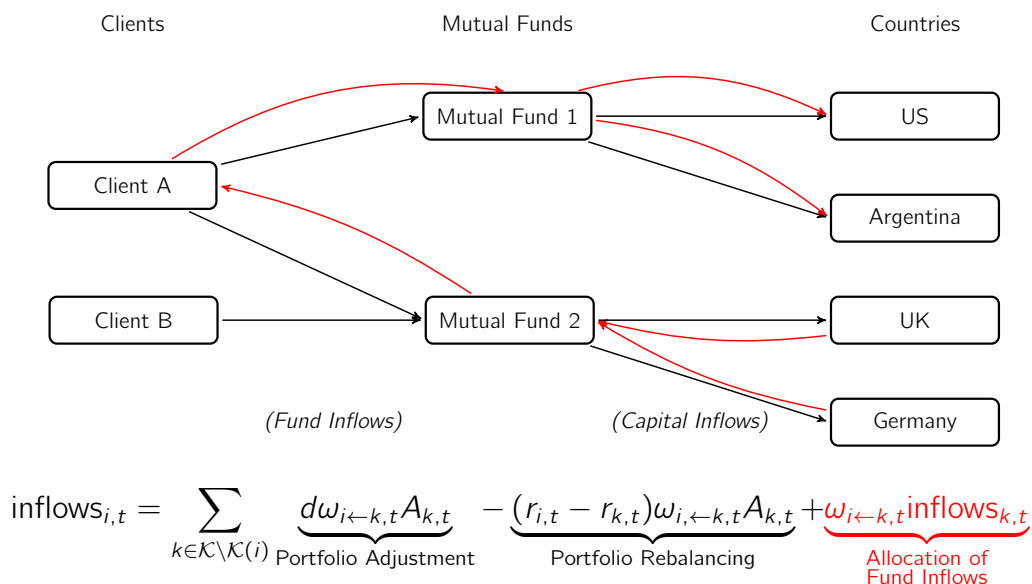
As an example of client portfolio adjustment, consider the reaction of Client A in the above example when she learns about the Argentinean policy announcement.²⁰ As shown in Figure 5, it is reasonable to expect that Client A might want to increase her exposure to Argentinean assets and so would move some of her savings from Mutual Fund 2 towards Mutual Fund 1. This would lead to outflows from the UK and Germany and inflows into the US and Argentina. As a result, the instrument's predicted inflows into Argentina would include responses to news about this country and would violate the exclusion condition.

In particular, mutual fund choice need not be based directly on news about country performance but could be based indirectly on variables that predict future country shocks like past equity returns. If equity markets quickly price in news about rises in future GDP, and client inflows follow past returns earned by the mutual fund, then a particular fund's inflows could easily be predicting future economic growth in the countries the fund invests in.

in Borusyak et al. (2021) (and is one case discussed in Goldsmith-Pinkham et al. (2020)). I provide further discussion of this in Appendix B.1 and refer interested readers there.

²⁰Note that the data used in the instrument construction has no information on client identities or their portfolios across mutual funds. All that is observed is the total inflows each fund received from all their clients. The client structure shown here is entirely for the purposes of discussing threats to identification.

Figure 5: Endogenous variation captured by the client inflows – good news about Argentina



This issue is significant: a large literature on mutual fund inflows started by Gruber (1996) shows that clients *do* use past fund returns to pick the mutual funds they invest in (and this applies to the mutual fund data used here). Given the speed at which financial markets price in news, good news about Argentina could lead to large returns posted by mutual funds exposed to Argentina well in advance of when the good news is visible in Argentina's GDP. If client inflows react to these returns sufficiently quickly, then they might precede the GDP change and therefore not be exogenous.

Fortunately, the high sensitivity of clients to the past returns of mutual funds offers another avenue to construct an exogenous instrument: using past returns *from other countries* to predict which mutual funds get inflows. To do this I first run the following regression:²¹

$$\frac{\text{inflows}_{k,t}}{A_{k,t}} = \gamma_t + \sum_{p=1}^{10} \delta_p r_{k,t-p} + e_{k,t} \quad (7)$$

²¹The left-hand variable is the inflow *rate* so that large mutual funds in terms of assets do not dominate the regression results and because the identification assumption is best thought of in terms of inflow rates rather than levels (see Appendix C.1). Inflow rates are calculated assuming that the timing of inflows within a quarter is constant. See Appendix A.1 for details.

where $r_{k,t-p}$ is the return of mutual fund k made in period $t - p$.²² Then for each mutual fund k I construct a set of “leave-out returns” for each country i that this fund k invests in:²³

$$\hat{r}_{k \setminus r_i, t-p} = \sum_{j \neq i} \frac{\omega_{j \leftarrow k, t} r_{j, t-p}}{1 - \omega_{j \leftarrow k, t}} \quad (8)$$

and produce a set of predicted inflows for each k according to:

$$\text{inflows}_{k,t} \approx \text{inflows}_{k \setminus r_i, t} \equiv A_{k,t} \left(\gamma_t + \sum_{p=1}^{10} \delta_p \hat{r}_{k \setminus r_i, t-p} \right) \quad (9)$$

where each predicted shifter $\text{inflows}_{k \setminus r_i, t}$ ignores the contribution the lagged return in country i had on overall inflows into k .²⁴

Finally, when predicting overall inflows into country i , I replace the shifter in the overall instrument ($\text{inflows}_{k,t}$) with $\text{inflows}_{k \setminus r_i, t}$:

$$\tilde{B}_{i,t} = \sum_{k \in \mathcal{K} \setminus \mathcal{K}(i)} \omega_{j \leftarrow k, t_0} \text{inflows}_{k \setminus r_i, t}$$

where now the shifter used varies by the country i whose inflows are being predicted.

Figure 6 shows in a diagram how to understand this approach. Consider a simple example where there are two mutual funds: one invests in the US and Argentina, and the second invests in Argentina and Brazil.

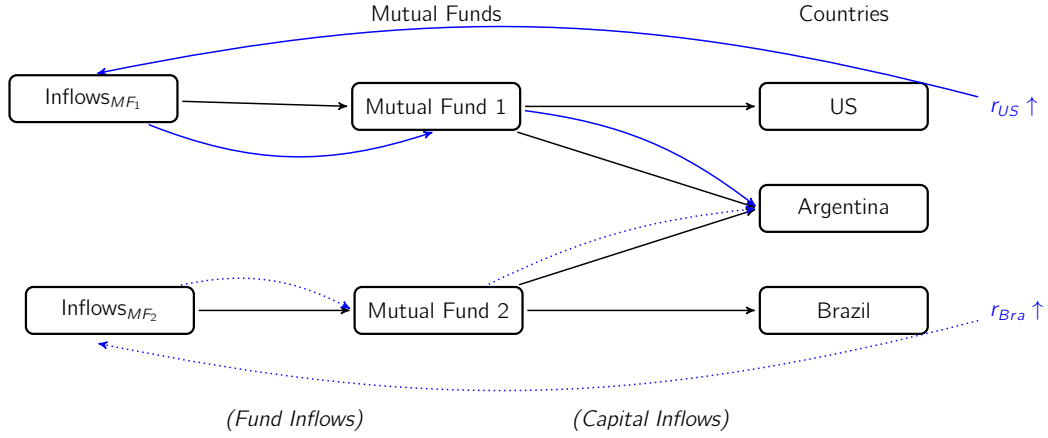
The new instrument uses an Argentina-specific approximation for each shifter when predicting inflows into Argentina (and a Brazil-specific approximation for Brazil and so on). Inflows into Mutual Fund 1 are predicted using lagged returns from the US, and inflows into Mutual Fund 2 are predicted using lagged returns from Brazil – in both cases lagged returns in Argentina are ignored. With an estimated coefficient on lagged returns of $\hat{\delta}_p$ (ignoring for now the time fixed effects), predicted returns into Argentina is given by:

²²The 10 lags are chosen by BIC.

²³The country-level returns $r_{i,t-p}$ are not available directly in the EPFR data. See Appendix A.2 for details on how these country-level returns are constructed.

²⁴Under this construction, “dedicated” funds or funds that invest only in a single country are dropped from this instrument.

Figure 6: Instrument refinement: Variation used to predict inflows into Argentina



$$\tilde{B}_{ARG,t} = \omega_{ARG \leftarrow MF_1, t_0} \underbrace{A_{MF_1, t} \sum_{p=1}^{10} \hat{\delta}_p r_{US, t-p}}_{\equiv \text{inflows}_{MF_1 \setminus ARG, t}} + \omega_{ARG \leftarrow MF_2, t_0} \underbrace{A_{MF_2, t} \sum_{p=1}^{10} \hat{\delta}_p r_{BRA, t-p}}_{\equiv \text{inflows}_{MF_2 \setminus ARG, t}}$$

$$\tilde{B}_{ARG,t} = \omega_{ARG \leftarrow MF_1, t_0} \underbrace{A_{MF_1, t} \sum_{p=1}^{10} \hat{\delta}_p r_{US, t-p}}_{\equiv \text{inflows}_{MF_1 \setminus ARG, t}} + \omega_{ARG \leftarrow MF_2, t_0} \underbrace{A_{MF_2, t} \sum_{p=1}^{10} \hat{\delta}_p r_{BRA, t-p}}_{\equiv \text{inflows}_{MF_2 \setminus ARG, t}}$$

And in the same example to predict inflows into Brazil, the instrument would be constructed from only Mutual Fund 2 (as Mutual Fund 1 does not invest in Brazil) and would now use the lagged return in *Argentina* as the shifter for inflows into Brazil:

$$\tilde{B}_{BRA,t} = \omega_{BRA \leftarrow MF_2, t_0} \underbrace{A_{MF_2, t} \sum_{p=1}^{10} \hat{\delta}_p r_{ARG, t-p}}_{\equiv \text{inflows}_{MF_2 \setminus BRA, t}}$$

The exclusion condition for this instrument is that lagged returns that mutual funds (with positive holdings in country *i*) make in countries *other than country i* have no predictive power for future returns or economic outcomes in *i*. For sufficiently diversified mutual funds this is likely to hold, however for specialized mutual funds one still might be concerned about

spatial correlation between country shocks: in the above example, lagged returns in Brazil could be predicting future shocks in Argentina given that the two economies are likely to have correlated business cycles. I address this concern in two ways: in a robustness exercise I use a “regional-leave-out” shifter – only lagged returns from countries in a different *region* to i (defined by the UN’s M49 area codes) are used, and again with a “continent-leave-out” shifter. These exercises deliver qualitatively similar patterns and is shown in Figure 28 in Appendix D.1. Moreover, given that the concern is that mutual funds investing in countries with strong growth prospects will get more inflows from clients than mutual funds investing in countries with weak prospects, in Section 4.1 I conduct several tests based on those suggested in Borusyak et al. (2021) to show that these predicted inflows are not affected by some measure of country news.

While the EPFR contains data on bond inflows, unfortunately the same instrument refinement does not achieve a strong enough first stage. Therefore, in this paper I focus solely on portfolio equity inflows.

2.5 Second Stage – Dynamic Responses to Capital Inflows

I use Jordá (2005b) local projections to estimate the dynamic response of a country-level macroeconomic outcome Y (for example, GDP, consumption, investment) to a 1 percent of GDP increase in portfolio equity inflows. Local projections have several advantages over traditional VAR-based approaches – namely, other than being linear, estimates are non-parametric and the regressions can easily be integrated into more traditional two-stage-least-square setups, which in this context makes it straightforward to apply the correct standard error calculations required in shift-share contexts (Adão et al., 2019). Furthermore, in this paper I will investigate the effects of many macroeconomic variables and split my sample into fixing and floating countries; in a VAR context, this would require the estimation of a large number of parameters, however in a local projections framework considerably fewer parameters are needed. Finally, while not strictly infeasible it is not typical in many panel VAR applications to include time fixed effects, whereas in this context where capital inflows react strongly to global macro shocks (see for instance Rey, 2012), controlling for time fixed effects is crucial.

In a panel of quarterly data on many countries (each denoted by i) this implies the following regressions:

$$\frac{Y_{i,t+h} - Y_{i,t-1}}{Z_{i,t-1}} = \alpha_i^h + \gamma_{c(i),t}^h + \sum_{p=0}^P \beta_p^h \frac{\text{Equity Inflows}_{i,t-p}}{\text{GDP}_{i,t-p}} + \sum_{l=1}^P \psi_l \frac{\Delta Y_{i,t-l}}{Z_{i,t-p-1}} + \theta^h \text{Sum of shares}_{i,t} + \epsilon_{i,t}^e \quad (10)$$

where Y is the outcome variable of interest, α_i^h are country fixed effects in their respective regressions, $\gamma_{j(i),t}^h$ are continent-time fixed effects (where $c(i)$ denotes the continent of i) and the controls include four lags of the outcome variable and of the relevant inflow variable. The continent-time fixed effects are an important control because they remove any common variation in push factors at the continent level. As mentioned above, world economic booms lead to falls in global risk sentiment and rises in capital inflows (and outflows) around the world. Time fixed effects control for this confounding mechanism, and continent-time fixed effects further control for this mechanism operating at the continent level (where a boom in Asia leads to higher investor confidence in investments in Asia and the redirection of inflows towards Asia).²⁵

One additional control is included as suggested by Borusyak et al. (2021) in contexts where the shares over mutual funds in the shift-share instrument used do not sum to one for each country.²⁶ The reason for this control is that it captures implicitly the heterogeneous exposure of countries to investors not in the sample. As these investors' behavior is not observable, it is prudent to identify $\{\beta^h\}_{h=1}^H$ with variation in predicted inflows to investors *in the sample* and to hold constant each country's overall exposure to all other investors.

As is typical for this type of regression, I include lags of the dependent and endogenous variable to reduce autocorrelation in the residuals. I set $p = 4$ so that four lags are included in the baseline specification.²⁷

The dynamic response is estimated by running H regressions, each with a different lead of the outcome variable Y_{t+h} $h \in \{0, \dots, H\}$. The collection of coefficients $\{\beta_0^h\}_{h=0}^H$ represent the cumulative impulse response of outcome variable Y (as a percent of Z in period $t - 1$) to a 1 percent of GDP increase in equity.

²⁵The results are almost identical to including subregion-time fixed effects using the UN's M49 subregion classification.

²⁶In Borusyak et al. (2021), this control is needed if the shares don't sum to one *for each cross-sectional unit* i – that is, $\sum_{k \in \mathcal{K} \setminus \mathcal{K}(i)} \omega_{i \leftarrow k, t} \neq 1$. In this context, however, the shares $\omega_{i \leftarrow k, t}$ are portfolio shares and so would sum to one *for each mutual fund* k . A further complication is that the shares often don't sum to one for each mutual fund because some funds are held in cash and not all countries each mutual fund invests in are in my sample. For both reasons, controlling for the sum of shares is important for the identification assumption to be valid.

²⁷The results are unchanged from including anything from two to eight lags. See Figure 27 in Appendix D.

In most cases, Z is set equal to the outcome variable in the period prior to the inflow $Z_{i,t-1} = Y_{i,t-1}$ to estimate a cumulative impulse response of Y in percent. However, for components of GDP (consumption, investment, and so on) I set $Z_{i,t-1} = \text{GDP}_{i,t-1}$ to reflect (cumulative) changes in the shares of GDP of these variables. This normalization does not affect the qualitative findings but provides a more convenient interpretation of the responses of these variables.

These regressions are estimated with two-stage-least-squares with the instruments described in the previous subsection.

3 Data

Macroeconomic data come from a variety of sources detailed in Table 3 in Appendix A.1. Inflow data come from the IMF's Balance of Payments (BoP) database – specifically from “net incurrence of liabilities” in the financial account. Macroeconomic data come largely from the data stream and is supplemented with data from the World Bank, the OECD and the IMF's International Financial Statistics (IFS) database. For each country and variable, the data come from only one source (that is, there is no splicing), however the source for each variable might differ by country depending on the quality of data of each source. I collect further data on equity prices, returns and bond returns from the Global Financial Database (GFD). I will show results split by exchange rate regime, using data from Ilzetzki et al. (2017), where I use their coarse classification to determine which countries have fixed exchange rate regimes in a given period. I then collapse all other countries into a group I denote “non-fixers,” which include countries with “dirty” floats as well as pure floats.²⁸

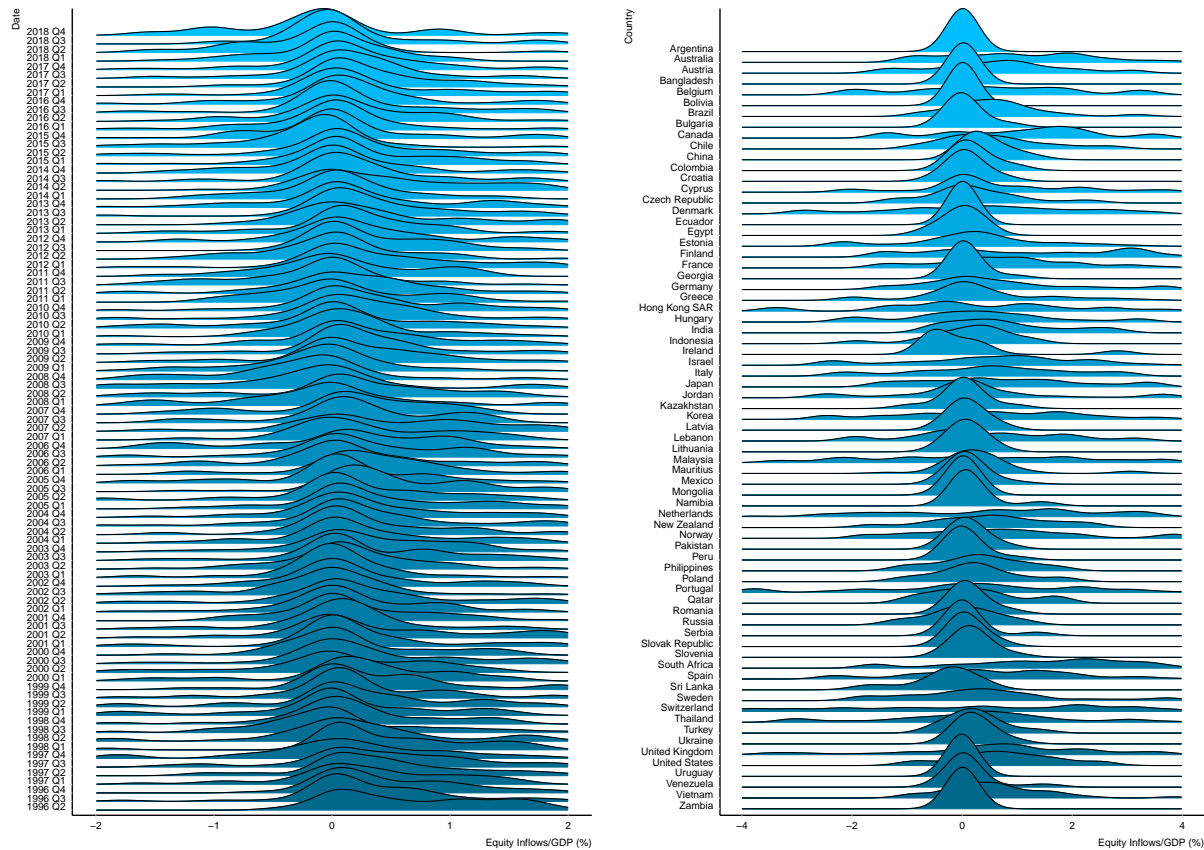
The data collected covers 67 countries in an unbalanced panel and allows for regressions to be run from 1996 Q2 to 2019 Q1, although for reasons discussed in the next section it will be appropriate to begin the regressions at later dates.²⁹ The Balance of Payment Flows are normalized by GDP and include categories of portfolio equity, portfolio debt, banking, foreign direct investment and other. Running the empirical specifications with market capitalization in the denominator produces similar results.

Figure 7 shows the the distributions of portfolio equity inflows both over time (left panel) and by country (right panel). While the mean of inflows changes slightly over time, there can be even larger changes in skewness – for instance, during 2008 negative skewness of the inflow

²⁸The Ilzetzki et al. (2017) ends in 2016, and so I assume that exchange rate regimes remain unchanged from their 2016 values over the remainder of the sample (to 2018 Q4). Figure 17 in Appendix A.1 shows the classifications over time for all countries in the sample.

²⁹For more details on which countries are in the sample, see Figure 16 and Table 2 in Appendix A.1.

Figure 7: Equity inflows by date and country



(a) By Date

(b) By Country

Notes: The panels show kernel densities of equity inflows as a percent of GDP over time (left) and country (right) for all countries in my sample.

distribution widens considerably. The right panel shows considerable variation by countries where advanced economies tend to have higher inflows on average as well as more variable inflows.

The mutual fund-level data come from Emerging Portfolio Fund Research (EPFR). EPFR contains monthly data on equity and bond fund flows and country allocations for approximately 24,000 equity funds (as of 2018).³⁰ The data starts in 1996 Q2 for equity mutual funds.³¹ The mutual fund data consist of two databases. The first, “FFShareClass,” provides a detailed

³⁰For reasons discussed in Appendix A.2, the sample of funds used for instrument construction contains approximately 1600 equity funds. The reason for this is that the “CountryAllocations” EPFR database covers a considerably smaller set of funds than the “FFShareClass” database. EPFR tends to estimate portfolios for the missing funds using some imputation methods, which unfortunately risk invalidating the identification assumptions required in this paper.

³¹EPFR includes mutual fund information for bond mutual funds as well as mixed and other types of mutual funds. Unfortunately, there are too few mutual funds to construct sufficiently strong instruments for portfolio debt inflows.

account of fund flows and assets, including the specific type of equity or bond class that the fund is investing in. For example, this can include attributes such as credit quality, currency and sector of the fund's investment. These monthly data include assets under management at the beginning and end of the period as well as the change in asset valuation due to both market movements and exchange rate movements. Inflows into each fund represents new money entering the fund and is computed by the following formula:

$$\text{inflow}_{k,t} = A_{k,t+1} - A_{k,t} - \underbrace{(\text{Forex Change}_{k,t} + \text{Portfolio Change}_{k,t})}_{\text{USD Return}}$$

where $A_{k,t}$ and $A_{k,t+1}$ are the beginning and end period of assets, and $r_{k,t}$ measures the market return of the fund (in USD).³²

As discussed, this includes only new money invested in the fund and does not capture changes in asset valuation due to market returns. This corresponds to the concept of $\text{inflows}_{k,t}$ as defined in Equation 5. Fund flows are aggregated to the level of the fund, which can then be matched to the second database, "CountryAllocations," which contains fund-level country portfolios. This dataset provides the fund's self-reported percentage of total assets invested in each country. For each fund, the base period portfolio share in country i is taken from this database.

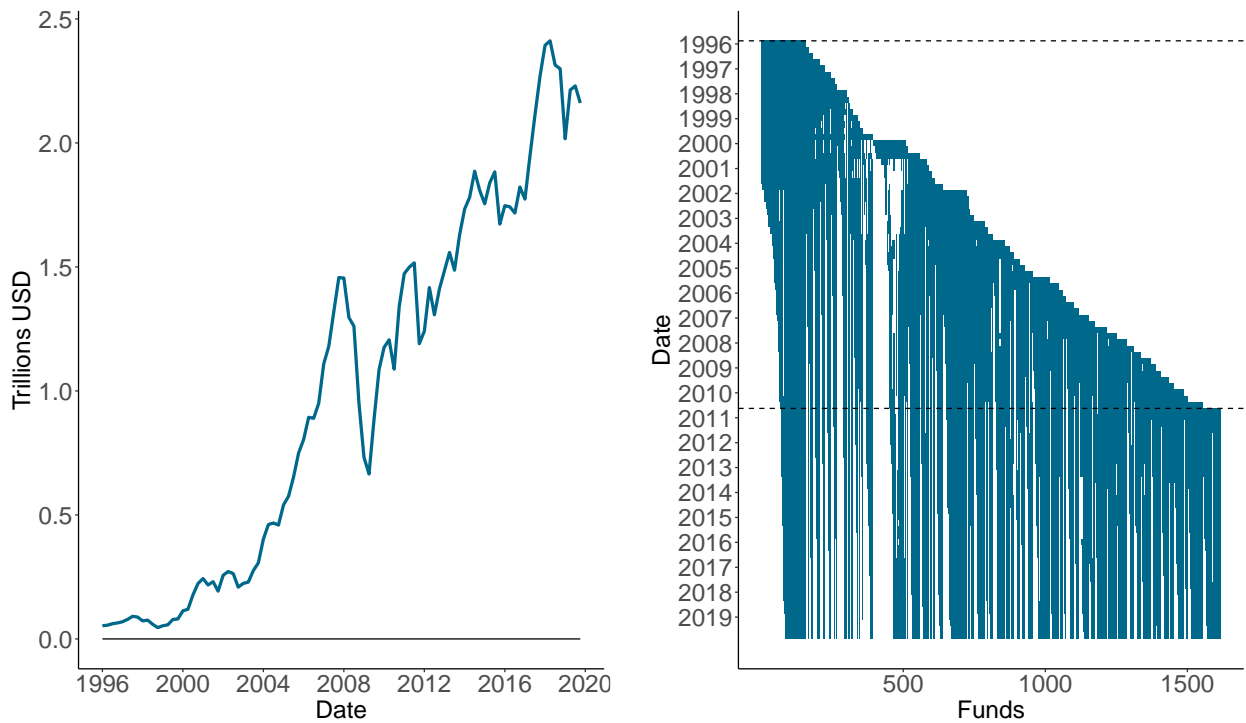
To construct all versions of the instrument discussed above I remove dedicated funds (those that only invest in one country). I also filter out the country of domicile in computing portfolio shares. For example, for a fund domiciled in the US, the US portfolio share is set to 0 even if the fund invests in the US. This is because capital inflows into the US should not include investments made by US residents in the US. Furthermore, recent work by Coppola et al. (2021) has shown that because of tax havens there is considerable mismeasurement in international investment positions and suggested that considerable adjustments are needed to capture correct international investment positions. Fortunately, because the capital *flow* data are not bilateral and are netted across financial transactions, a capital flow from the US to China through the Cayman Islands will correctly be recorded as a gross outflow from the US, a gross inflow to China and a gross inflow of 0 on net in the Cayman Islands.³³ However, in

³²Both the asset positions and overall mutual fund return are data provided by EPFR.

³³One potential issue may arise, however, if a Chinese firm lists a security in the Cayman Islands and then makes transfers from the Cayman affiliate firm to its head office. In this case, the flow from the US to the Cayman Islands will be (correctly) recorded as an equity inflow, but the flow from the Cayman Islands to China may be (incorrectly) recorded as a different type of inflow. The netting off would occur across all types of flows but not in terms of equity inflows. While Coppola et al. (2021) have kindly provided the data to make adjustments to bilateral international investment positions, there unfortunately isn't data available on how to make adjustments to capital flows nor is there data on bilateral flows that may be needed to correctly perform

the context of mutual funds, there may be some issues where funds domiciled in the Cayman Islands take funds from US clients, or funds in Ireland or Luxembourg take funds from all across the Euro Area. Therefore, for funds in the Cayman Islands I set their US shares to 0, and for funds in Ireland and Luxembourg I set their portfolio shares in all of the Euro Area to 0 as well. Further details on processing the instrument from the EPFR data can be found in Appendix A.2.

Figure 8: EPFR coverage – Assets under management (left) and fund prevalence in sample (right)



(a) Assets under management (AUM)

(b) Fund length in the sample

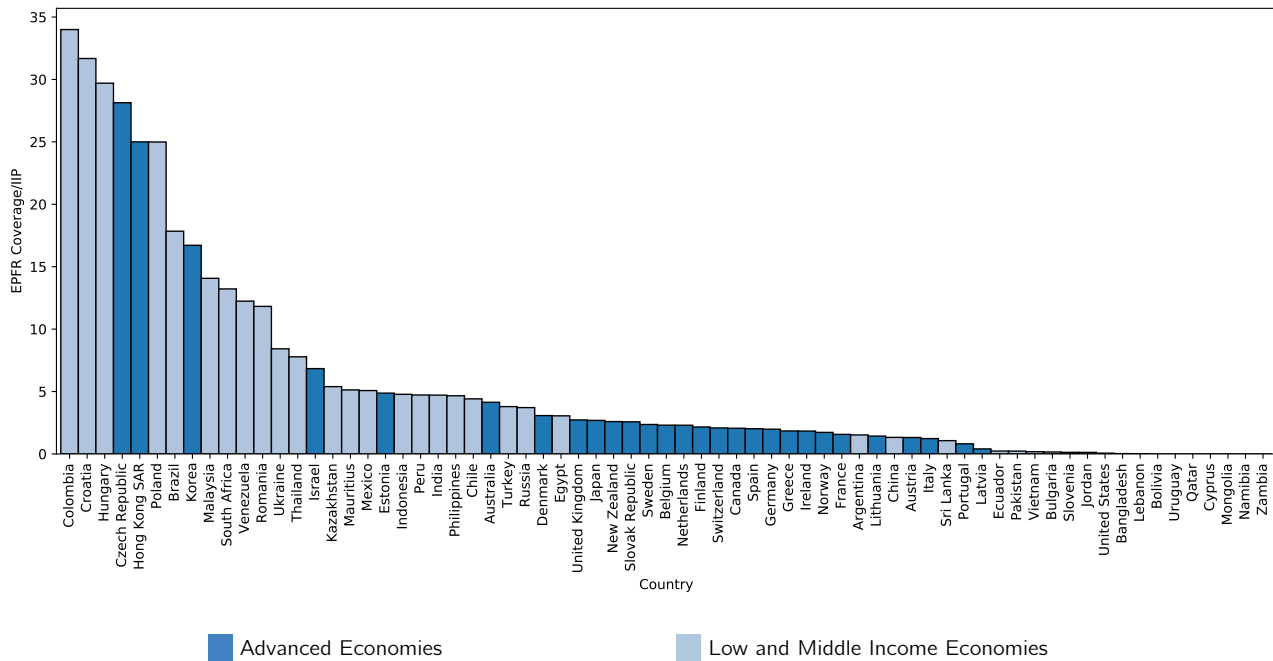
Source: EPFR. Author's own calculations

After merging the fund flows and country allocations database, the data are aggregated to a quarterly frequency and merged with the Balance of Payments data. EPFR data have been improving over time, and as such the number of funds covered tends to increase over time. The impact of this is shown in Figure 8. The left panel shows the total assets under management for the funds in the sample. In the second half of the sample, assets under management is impressive – over \$2 trillion USD – however, because the instrument construction requires that funds report in the base period, there will be a trade-off between instrument strength

the adjustment.

and sample length. This issue can be further seen in the right panel, which shows the sample period on the vertical axis that each fund is in the sample for. Mutual funds are ordered so that funds to the left enter the data early and funds to the right enter the data later.³⁴ Then within each entering cohort, funds are ordered by when they first leave. Of the first 200 funds that enter the data at the beginning of the sample – 1996 Q2 – only about half report until the end of the sample. This remains a problem until about 2004/2005, when several large cohorts that report until the end of the sample period have entered the data.

Figure 9: Coverage of foreign equity claims by EPFR in 2005 Q1



Notes: This graph shows the ratio of the total amount of assets in the EPFR allocated to each country in the sample as a share of that country's total equities owned by foreigners (in percent). The total international equity liability position was obtained from the IMF's Balance of Payments database.

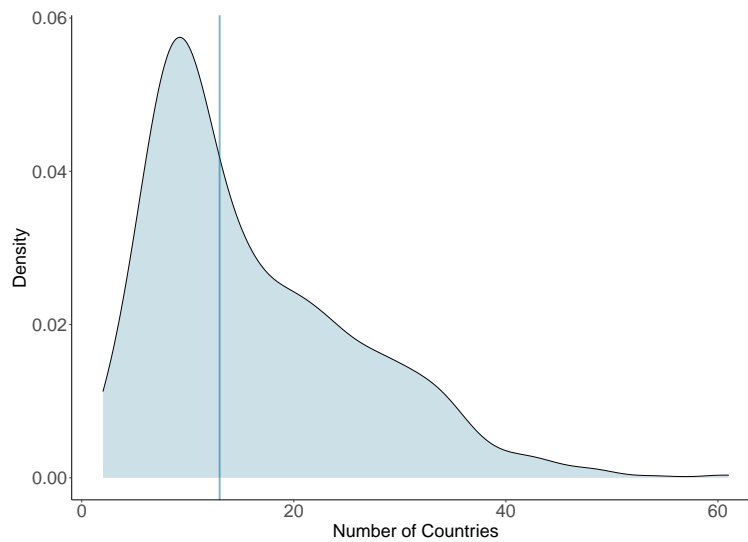
Furthermore, by 2005 the EPFR coverage of some countries' international equity liabilities (Equity IIP) – equity claims held by foreigners – is impressive. Figure 9 shows the share of Equity IIP that is covered by the matched mutual funds in the sample. For the median country, EPFR covers 3-4 percent of Equity IIP, but this rises to above 30 percent for Colombia, Croatia and Hungary. Furthermore, for the Czech Republic, Hong Kong, Poland, Brazil and Korea the coverage is above 10 percent of total Equity IIP.

Finally, for the instrument to have relevance, it should be the case that mutual fund portfolios are considerably heterogeneous by country. Figure 10 confirms this by showing the

³⁴The right panel does not show funds that enter the sample after 2010 Q4. These comprise another 1000 funds in both EPFR databases.

distribution of number of countries invested in by the equity funds in the EPFR data. The median mutual fund invests in 16 countries, and the modal fund invests in 10 countries. Most funds at the low end specialize in a particular region of the world (such as Latin America), and the funds that invest in more countries tend to be global funds investing all around the world. This supports the theory that as the pattern of inflows between funds shifts over time, that distribution of capital flows between countries will shift considerably in response.

Figure 10: Number of countries invested in



Notes: This figure shows the distribution of the number of countries funds in EPFR invest in. The vertical line show the median of the distribution.

4 First Stage Results

A nice feature of this dataset relative to typical shift-share settings is that it is a long panel. This allows an assessment of the instrument's validity (and second stage estimates) over many choices of base years. If the instrument is really capturing the variation in capital inflows, then the F-statistic should be high regardless of the base period chosen. In particular, changing the base period by a quarter should have little impact on the F-statistic. By contrast, if the first stage F-statistic were high only for a small number of very different base period choices, this would imply that overall the instrument is weak and the few F-statistics above 10 are a statistical artifact. I therefore show all F-statistics across many years of base period choices.

The first stage in this setting is given by the following equation:

$$\text{Equity Inflows}_{i,t} = \delta_i + \kappa_t + \psi B_{i,t} + \theta' \text{controls}_{i,t} + \xi_{i,t}$$

where as before i indicates country, t is time and $B_{i,t}$ is the shift-share instrument introduced in Section 2.2.

The second stage regressions include lags of the dependent variables as controls, and as such imply different first stage regressions for each dependent variable.³⁵ Therefore, for robustness, I run separate first stage regressions for each specification used (one per left-hand-side variable) for a variety of base periods for the shift-share instrument covering 1996 Q2 to 2010 Q4. The results for pooled regressions are shown in the left panel of Figure 11. The x-axis shows the base period for the instrument construction and is also the beginning of the sample for each corresponding first stage regression. Each grey line then represents a single regression specification (one for computing the equity inflow impulse response, another for real GDP impulse response and so on) and the thicker blue line shows the *minimum* F-statistic across all these specifications. Thus, when the blue line is above 10, it means that over all first stage specifications I run, the F-statistics are above 10. This is an extremely conservative approach to ensuring that the instrument is strong.

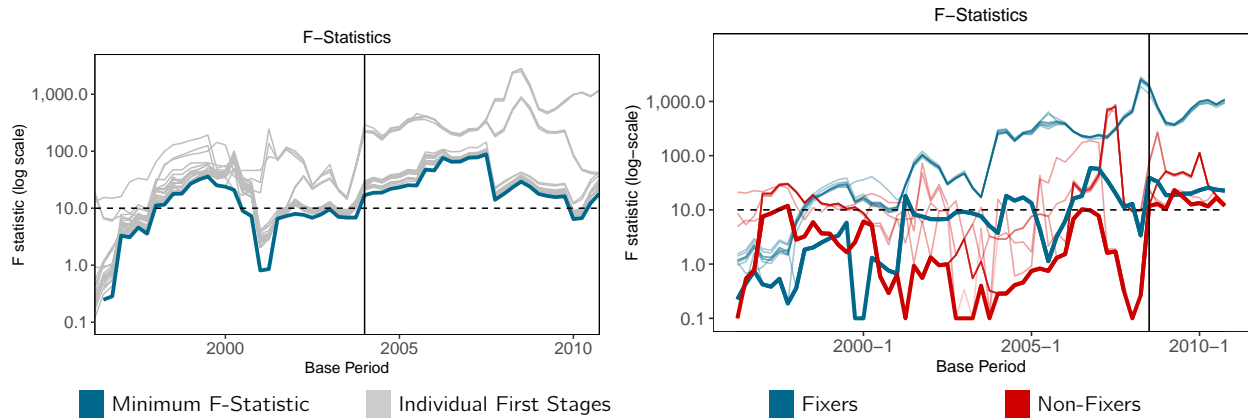
Around 2004 Q1 all F-statistics remain robustly above 10 for the remainder of the period (bar 6 months in early 2010) and as the scale is logarithmic, for some variables the F-statistic is extremely strong. To maximize the sample used, I therefore opt to run the pooled regressions starting in 2004 Q1. This corresponds to the time when there is a large influx of funds into the EPFR data that remain in the data until the end of the sample (see Figure 8 panel (b)). Broadly speaking, selecting alternative dates with F-statistics above 10 has little effect on the impulse responses reported in Section 5. These are presented in Figures 26, 27 and 28 in Appendix D.

Panel (b) of Figure 11 shows the F-statistics from all first stages when the sample is split for each time period into countries with fixed exchange rate regimes and those with other regimes (either “dirty” regimes or floating regimes). The light blue lines show F-statistics for the sample of fixers and the red lines for the non-fixers. The thicker lines show the minimum F-statistics for each group. For fixed exchange rate countries the F-statistics are above 10 for almost all variables (the one exception is short-term interest rates), and for floating exchange rate regimes all but two first stages are above 10 robustly after 2005.³⁶ However, it is not until

³⁵But note that changing the horizon of the impulse response involves taking only a lead of the dependent variable without changing the right-hand-side set of variables. Thus there is a single first stage for each variable.

³⁶The two variables for which the F-statistic is below 10 over this period are the short-term interest rate and

Figure 11: F-statistics for pooled (left) and by exchange rate regime (right)



Notes: Each panel shows the F-statistics over a variety of first stage regressions. Each Jorda local projection regression has a different right-hand-side value for each variable, necessitating a different first stage regression. Furthermore, the shift-share instrument can have a variety of different base periods. In the left panel, F-statistics for regressions covering all countries in the sample are computed. For each left-hand-side variable I compute an impulse response for (real GDP, real consumption, and so on) there is a grey line representing the F-statistics from the corresponding first stage regression for that variable as the instrument's base period is changed. Then the solid blue line in this panel shows the *minimum* F-statistic across all first stage regressions. When this line is above 10, it means that the F-statistics from *all* specifications are above 10. In the right panel, the same information is shown for two subsamples – in blue the subsample of countries that have fixed exchange rate regimes, and in red the subsample of countries with non-fixed exchange rates. As in the left panel, the lighter thin lines each show a first stage regression of a particular variable and the thicker line in each color shows the minimum F-statistic for that subsample respectively. Standard errors in the first stage are calculated using the approach from Adão et al. (2019).

2008 Q2 that all variables are robustly above 10 and in the interests of being conservative, this is the period I choose to present the baseline results.

4.1 Testing for Threats to Identification

The previous section has shown that the instrument is relevant for a variety of base years, but it is also important that the exclusion condition holds. Recall that the exclusion condition is that the rate of predicted inflows into mutual funds based on returns made in countries other than i are not correlated with news about country i . This can be tested by constructing measures of news about country i and seeing if there are differential inflows into funds with high exposure to i relative to funds with low exposure.

Formally, the test is:³⁷

CPI.

³⁷See Borusyak et al. (2021) for details.

$$\underbrace{\sum_{i=1}^N \overbrace{\omega_{i \leftarrow k, t_0} A_{k, t}}^{\equiv \tilde{\omega}_{i \leftarrow k, t}} \text{news}_{i, t}}_{\text{Exposure of } k \text{ to news from all countries}} \parallel \frac{\widehat{\text{inflows}}_{k, t}}{A_{k, t}} \Big| \text{controls}_{k, t} \quad (11)$$

where $\omega_{i \leftarrow k, t_0}$ are the portfolio shares of mutual fund k in i in the base period t_0 and $\text{controls}_{k, t} \equiv \sum_{i=1}^N \tilde{\omega}_{i \leftarrow k, t} \text{controls}_{i, t}$ – controls to be used in these tests are the controls in the original regressions converted to vary by mutual fund k instead of country i . The exclusion condition is expressed in terms of inflow *rates* as it is more plausible that the *rate* of inflows are as good as randomly assigned between large and small mutual funds than it is to argue that the *level* of inflows are as good as random.³⁸ Either is sufficient for identification to be achieved, but when testing for threats to identification it is important to express the exposure weights appropriately – in this case they are $\omega_{i \leftarrow k, t_0} A_{k, t}$.³⁹

I present two sets of tests. The first set considers several crisis events and compares the distribution of inflow rates between funds that invest in the affected countries to those that do not. Since crisis events are large and suggest large losses for affected mutual funds and their clients, if a violation to exclusion were to occur it should be visible in these events. The second test involves using two measures of future country performance – IMF forecasts and future country returns – that vary more continuously. Other than the Asia Crisis, where the number of EPFR funds is low, I find little evidence of systematic violations to the exclusion condition.

4.1.1 Tests Using Crisis Events

For the first test I use three large crises to see if clients of mutual funds withdrew money from funds exposed to the affected countries. I consider the set of countries affected during the 1997 Asia Crisis,⁴⁰ Argentina's 2001 Peso Crisis and Greece during its debt crisis at the end of 2009. In all cases, these events had sudden stops where investors pulled money out of the respective countries, which would suggest that these events are plausible candidates for

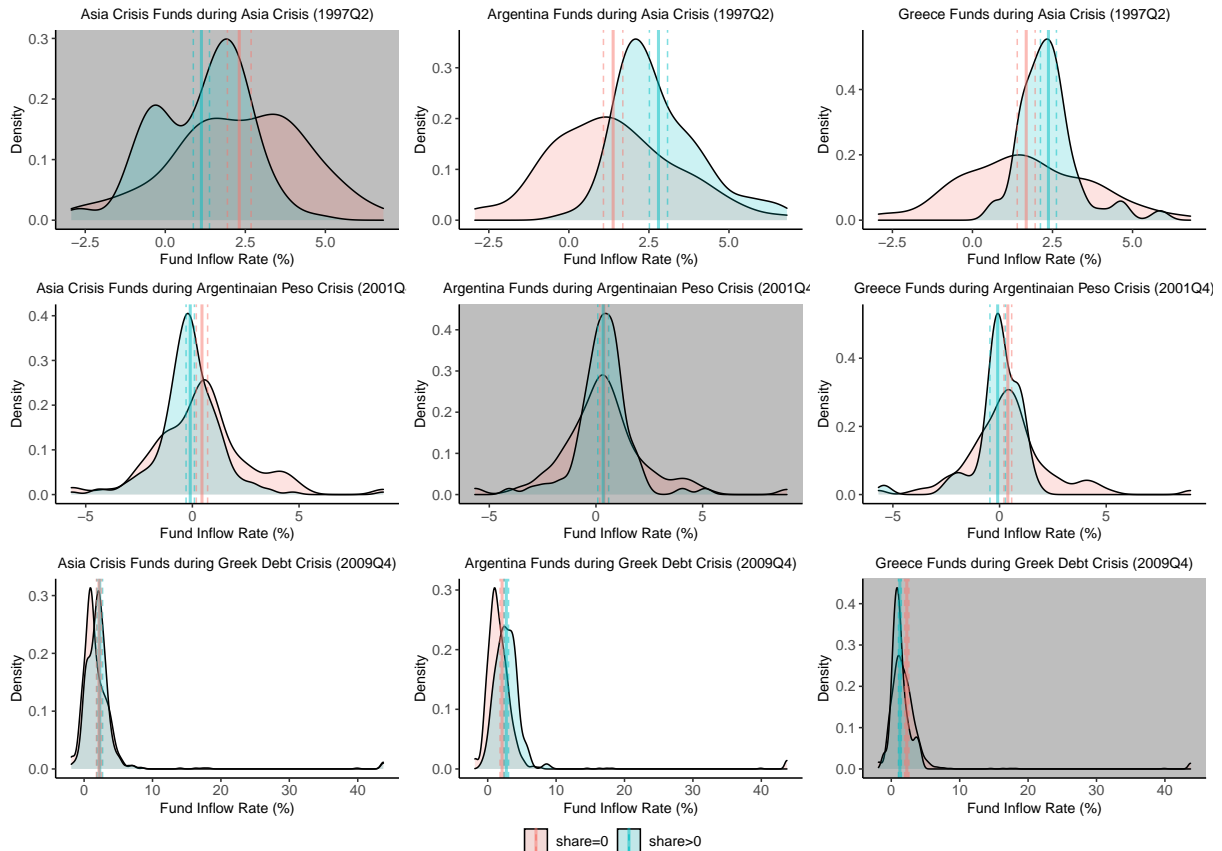
³⁸To deal with outlier inflow rates, inflow rates are calculated assuming that inflow rates occur at a constant rate throughout the quarter. See technical Appendix A.2 for details.

³⁹Note that this identification assumption is similar in spirit to the granular instrumental variables approach in Gabaix and Koijen (2020) with a different method to calculate the idiosyncratic shocks. Instead of estimating inflows idiosyncratic to each mutual fund (which requires that all common shocks be measured and controlled for), I predict inflows from each fund's lagged returns from other countries.

⁴⁰These countries are Thailand, Malaysia, Philippines, South Korea, Hong Kong and Indonesia. Figures 23 and 24 in Appendix C.3 show the results for each country individually.

assessing if clients are reacting to the exposure of their mutual funds to country shocks.

Figure 12: Inflow rate distributions for equity funds during three crises



Notes: Each graph shows the distribution of funds invested in the labelled country during either the Asia Crisis (1997 Q1), the Argentinian Peso crisis (2001 Q4) or the Greek Debt crisis. The panels with grey backgrounds show the inflow rate distributions of mutual fund investments and not investing in the country primarily affected by the crisis. The panels with white backgrounds are placebos: they show the same two mutual fund inflow distributions for funds invested in a country that was not (directly) affected by the particular crisis in question. The vertical lines show the respective means of each distribution with 95 percent confidence bands.

The results are shown in Figure 12. Each panel shows the inflow rate distributions in percent for mutual funds that invest positively in the country being considered (in blue) and for funds that are unexposed (in pink).⁴¹ Consider the top left panel, which focuses on the Asia Crisis. The inflow rate distributions for both exposed and unexposed funds are very wide: inflows vary a lot across individual mutual funds. They can be as low as -50 percent in this period and as high as 100 percent (a doubling of size). Next, based on this panel, it appears

⁴¹Formally, to test Equation 11, the exposure measure should be continuous rather than zero versus non-zero. The results are presented this way because the majority of funds have 0 exposure to any of these six countries during this period, and portfolio shares tend to be low (the same is true for Argentina and Greece during their crises). As such, the starkest comparison is between no exposure and some exposure. Bin-scatter plots showing tests more in line with Equation 11 are available upon request.

that the funds invested in the six affected Asian countries are more likely to receive negative inflow rates during this period but are also more likely to receive very positive inflow rates. The net effect is that the mean of this distribution – while lower – is not significantly different from the mean inflow rate of funds with no direct exposure to countries affected by the Asia Crisis. Given the extreme nature of this event, this striking result appears to occur for two reasons: first, the variance of the inflow distribution is wide, and second, the share of each fund invested in each country (or the set of six countries affected by the Asia Crisis) tends to be low. Conditional on investing a positive amount in a country, the share invested is typically less than 10 percent. As such, clients might not have sufficient incentive to go through the trouble of liquidating their investments in the affected mutual funds.

The top-center and top-right panels show a placebo test for this crisis. These panels repeat the exercise *during the Asia crisis* but compares mutual funds that do and do not invest in Argentina (top-center panel) and Greece (top-right panel). Both countries are geographically far away from the affected countries and are not typically considered to be connected (at least at the genesis of the crisis). In both cases, the inflow distributions also vary considerably for funds invested in these countries from those that are not. This suggests that the visual differences in inflows for funds exposed to and not exposed to countries affected by the Asia Crisis are simply statistical artifacts of the very volatile inflow rate distribution.

The next row repeats the exercise for the Argentinian Peso Crisis for the same three sets of countries (now the Asia Crisis countries – middle left panel – and Greece – middle right panel – are placebos), and the bottom row repeats the same exercise for the Greek debt crisis. In all cases the means of the two distributions shown are not statistically different and the inflow rate distributions are visually distinct in similar ways for both the affected and placebo countries. In this figure the distributions are only shown for the exact period the crisis occurs. However, the exclusion condition applies both to inflow rates before and after the crisis event.⁴² Figure 22 in Appendix C.3 shows these distributions for the affected countries up to a year before and after the crisis date. The results are similar.

4.1.2 Tests Using Regular Measures of Future Country Performance

This second test is based on that proposed in Borusyak et al. (2021). The logic of the test is to test the linear regression equivalent of Equation 11 with some measure of news $\phi_{i,t,t+h}$.⁴³

First the news is converted to vary *across mutual funds* and time by scaling by the base

⁴²Crises tend to have prolonged implications for future GDP growth in periods well after they begin. As such, similar inflow distributions even after the crisis event might still be required for identification.

⁴³Appendix C.1 contains more details on the approach.

period share term:

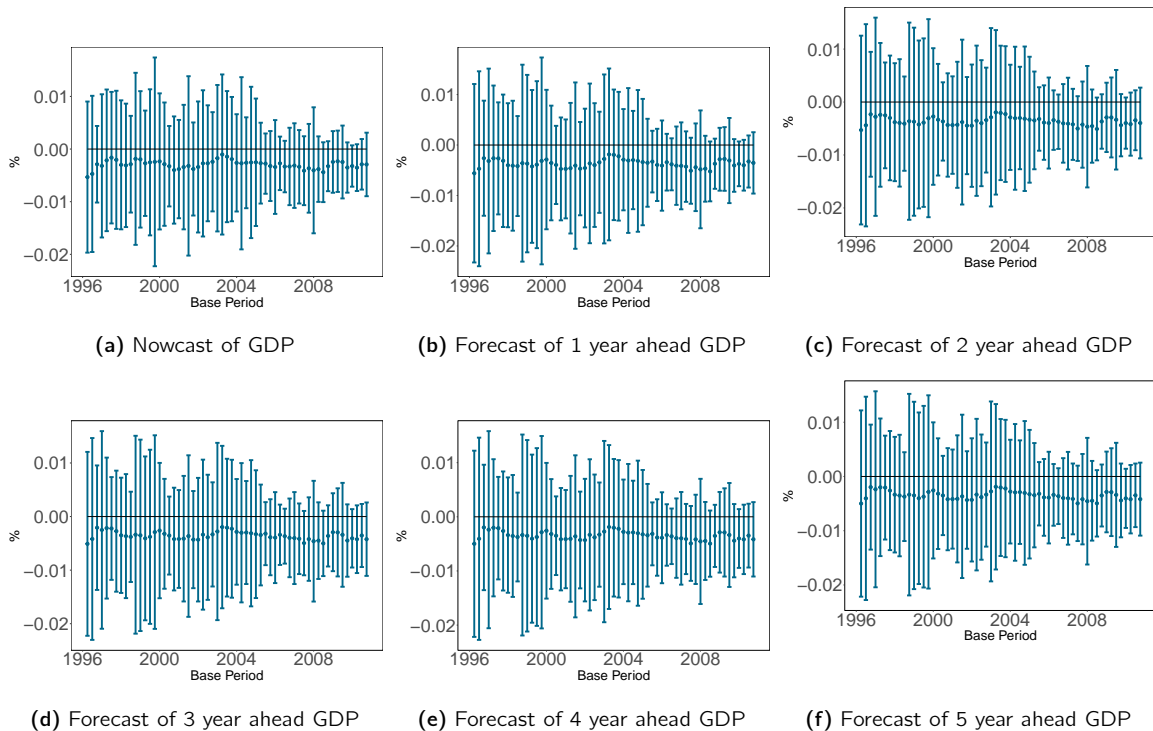
$$\tilde{\phi}_{k,t,t+h} = \sum_{i \in \mathcal{I}} \omega_{\leftarrow k,t_0} \phi_{i,t,t+h}$$

The new term $\tilde{\phi}_{k,t,t+h}$ is a measure of each mutual fund's exposure to countries with a high $\phi_{i,t,t+h}$. The test is captured in the following regression:

$$\frac{\text{inflows}_{k,t}}{A_{k,t}} = \alpha + \delta \tilde{\phi}_{k,t,t+h} + \text{controls}_{k,t} + \nu_{k,t+h} \quad (12)$$

where the controls are converted to vary the same way across mutual funds as $\phi_{i,t,t+h}$.

Figure 13: Test of inflow rate response of equity funds to exposure to IMF GDP forecasts



Notes: Each panel shows regressions from Equation 12 for a variety of base period choices when constructing the instrument. Each dot is a point estimate. Confidence intervals are 95 percent intervals. Standard errors are clustered by mutual fund and time. A coefficient of x should be interpreted as a fund exposed to countries with 1 percent higher GDP forecasts than another should have a higher inflow rate of x percent.

Two measures of ϕ are used: the first is historical IMF forecasts from the World Economic Outlook (WEO) report. These forecasts are made twice a year and cover every country in the sample, which allows for the tests to be run across many mutual funds investing all around

the world.

The results are shown in Figure 13. Each panel in Figure 13 presents a different forecast horizon. The top-left panel shows the relationship between contemporaneous GDP forecasts from the IMF and fund inflow rates, the top-center panel uses one-year-ahead forecasts and so on. As with the first stage graphs, every panel presents output from many regressions, each using a different base period, to convince the reader that the choice base period is irrelevant for the results of this test. Each dot is a point estimate corresponding to a particular base period, and the 95 percent confidence interval is given by the vertical lines. In all cases, the results are insignificant. Furthermore, the point estimates are very small and negative – if the countries one fund invests in are expected to grow by 1 pp more than the countries of another mutual fund, the first fund should expect *outflows* of 0.05 percent. Relative to the overall variation in the sample, a coefficient of 0.05 percent is minuscule. Put together, this suggests that inflows into bond funds are not strongly affected by news, as captured by the IMF forecasts.

Put together, neither of these types of tests are unable to detect movements in inflows that imply violations to exclusion. This suggests that the instrument as formulated is appropriate for estimating impulse responses to capital inflows.

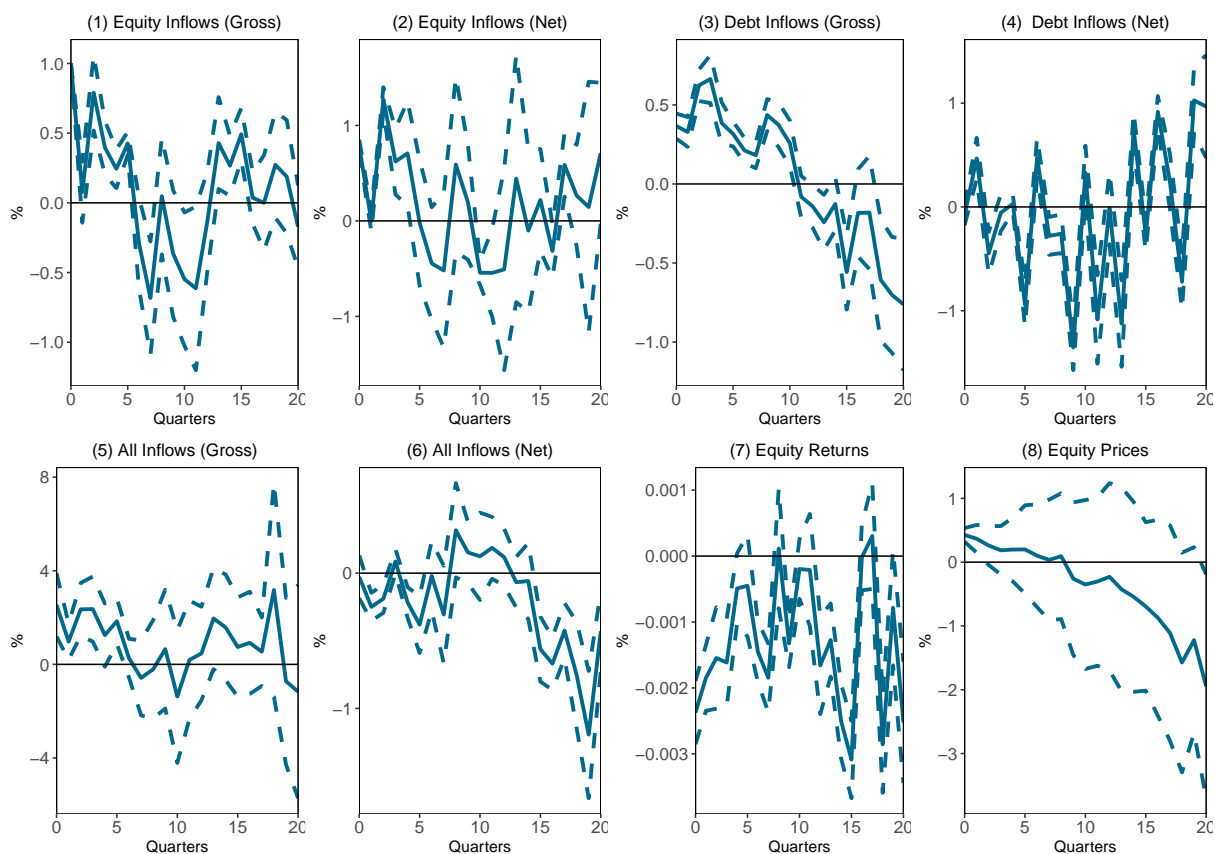
5 Second Stage Results

Figure 14 shows the cumulative impulse response results estimated from the full country sample using Equation 10 up to $H = 20$ quarters ahead. All regressions are instrumented with the leave-out instrument described in Section 2.4. Standard errors are calculated as in Adão et al. (2019), who suggest how to construct conservative standard errors for shift-share designs. These standard errors are robust to arbitrary country/time error structures and even generalize Anderson and Rubin (1949) standard errors for weak instruments.

The top-left panel (1) shows the cumulative dynamic response of gross equity inflows to an exogenous 1 percent of GDP increase in gross equity inflows occurring at quarter 0 (i.e., it is plotting the set $\{\hat{\beta}_h\}_{h=0}^H$ from Equation 10). While the coefficient must be 1 on impact, the dynamics of this impulse response come from the data and estimation procedure. These show that gross equity inflows are persistent – they remain elevated for around 18 months after the initial shock.

Panel (2) shows the dynamic response of *net* equity inflows (foreign equity inflows - domestic equity outflows), and the resulting dynamics are very similar to gross equity inflows. In very liquid financial markets one might expect that a sudden increase in demand for one

Figure 14: Cumulative impulse responses to a 1 percent of GDP increase in equity inflows



Notes: This figure shows the second stage results to a 1 percent of GDP increase in equity inflows. The base period of the instrument is 2004 Q1. The dashed lines represent 95 percent confidence bands calculated using the method of Adão et al. (2019).

country's equities would lead to domestic investors capitalizing on the increased demand for domestic equities by reallocating their portfolios towards foreign equities, yet this is not what appears to be happening.

Next, Panels (3) and (4) show the dynamic response of gross and net portfolio debt inflows respectively. They show that exogenous gross portfolio inflows appear to also lead to gross portfolio debt inflows, yet little or even negative responses of net debt inflows. This suggests that perhaps domestic investors are using the funds from selling equities to foreigners to purchase bonds and other assets abroad rather than equities. This is confirmed in panels (5) and (6), which show the responses of gross and net inflows of all types respectively. Here it is clear that when foreigners purchase equities, domestic investors appear to re-invest the funds abroad in some form or another, leading to 0 or negative net inflows overall.

Finally, panels (7) and (8) show the responses of equity returns and equity prices respectively. Equity prices rise modestly and remain insignificantly different from 0 after two

quarters and cumulative equity returns fall modestly. This serves as an additional sanity check on whether the equity inflows are likely to be exogenous – endogenous inflows that chase returns should raise equity prices *and predict positive future returns*. The fact that returns are negative is consistent with these inflows occurring for reasons other than a change in country fundamentals.

Next, Figure 15 shows the responses of real activity variables, nominal exchange rates, CPI and short-term interest rates separately for countries with fixed exchange rate regimes (blue) and non-fixed regimes (red).⁴⁴ Panel (1) repeats the impulse response of gross equity inflows separately for both fixers and non-fixers. While there are differences in the precision of the estimates, the dynamic responses are somewhat similar for equity inflows. However, the responses of all other variables are strikingly different.

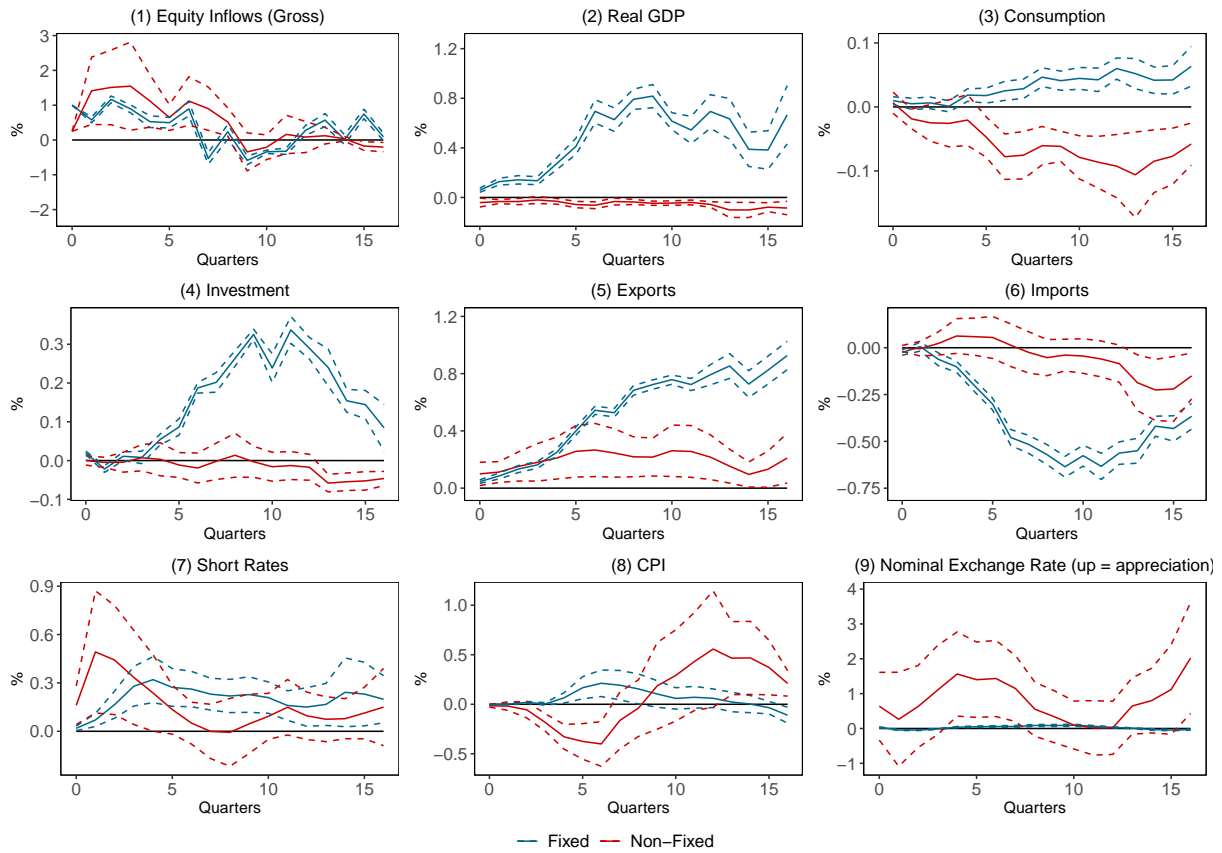
First, theory is the sharpest in focusing on the *difference* in reactions between fixing and floating exchange rate regimes. Capital inflows should create pressure for the exchange rate to appreciate, and theory suggests that whatever the overall macroeconomic impact of capital inflows, fixed exchange rate countries need to run looser monetary policy than floating exchange rate countries to keep the exchange rate stable. All else equal, this should lead fixed exchange rate countries to have *higher* GDP, consumption, investment and inflation than non-fixing countries. Moreover, without an exchange rate appreciation, exports for fixing countries should be higher than non-fixers and imports should be lower.

All these predictions are reflected in Figure 15 – fixing economies end up with *rises* in real GDP, consumption and investment, whereas floating economies face modest falls in all three variables. Exports are higher for fixing economies and imports are lower and inflation is higher too. Most importantly, the response of short-term interest rates suggest that the average fixing country's monetary policy stance is looser than the average non-fixing country. This is consistent with macroeconomic theory that inflation targeting countries prefer to allow their exchange rates to appreciate in response to exogenous rises in demand for their currency (whether induced by equity inflows or other shocks). This leads floating countries to make smaller monetary policy adjustments in response to these shocks than fixing countries that must maintain the exchange rate at the pre-shock level .

However, canonical macro models solely focus on the exchange rate-to-net exports channel when linking capital inflows to macro outcomes, and this channel alone suggests that fixing

⁴⁴Data from Ilzetzki et al. (2017) is used to classify countries as fixers or non-fixers using their coarse classification where a coarse classification of 1 denotes fixers and classifications of 2-6 denote non-fixers. Equation 10 is run for $H = 20$ on two separate samples of fixing countries and non-fixing countries (where countries can switch samples upon changing their regime – See Figure 17 in Appendix A.1). The instrument is the same as described in Section 2.4.

Figure 15: Fixed versus floating exchange rate regimes



Notes: This figure shows the second stage results to a 1 percent of GDP increase in debt inflows. The base period chosen is 2007 Q4. The dashed lines represent 95 percent confidence bands calculated using the method of Adão et al. (2019).

countries need to cut monetary policy to maintain their exchange rate peg. By contrast, panel (7) shows that short-term interest rates rise even for countries with fixed exchange rates, suggesting that capital inflows are raising demand in some manner. Further corroborating evidence is that despite higher short-term rates in reaction to this shock, CPI *rises* for countries with fixed exchange rates. Finally, if the exchange rate were to remain fixed one would expect very small reactions in exports to a capital inflows shock, yet for fixed exchange rate countries export volumes increase considerably, suggesting that equity inflows might somehow be facilitating firms to export more. Given the large rise in investment as well, it is possible that equity inflows might relax borrowing constraints for some firms, allowing them to expand production.⁴⁵

This suggests that the exchange rate-to-net exports channel, while useful for comparing

⁴⁵Kalemli-Ozcan et al. (2021), for instance, finds that firm borrowing increases when exchange rates appreciate.

outcomes across different monetary policy regimes, nonetheless is not the only mechanism at play when capital inflows enter a country. Indeed, despite monetary policy in non-fixing economies tightening sharply in response to exogenous equity inflows, real GDP is almost completely stabilized slightly below pre-shock levels. There is some evidence of an overreaction of monetary policy in terms of generating initial temporary falls in CPI, however macro outcomes for non-fixers appear considerably more stable than for fixing countries, suggesting that monetary policy flexibility is a useful tool to use in response to inflow shocks.

5.1 Robustness

In the interests of brevity I will give a summary of the robustness checks detailed in Appendix D. First, Figure 26 shows that the baseline results are robust to picking any other base period other than 2008 Q3 for the main results shown in Figure 15. Second, Figure 27 shows that the results are robust to the choice of lags in the local projections.

Further, the responses of variables to equity inflows are robust to further instrument refinements. Extending the leave-out definition to leave out all fund returns in the subregion of a country and even the whole continent of a country gives similar point estimates for activity variables. For these refinements the identification assumptions are very plausible: lagged returns in countries far away from country i are not correlated with future economic events in i conditional on country and time fixed effects and lags of outcome variables. However, the reader should note that in both cases the first stage is not robustly above 10 across all base periods, and therefore the second stage standard errors are large. Nonetheless, the point estimates remain similar to the baseline estimates discussed in Figure 15. Figure 28 shows these results.

6 Conclusion

This paper estimates the causal effects of capital inflows into equity on business cycle variables. The results suggest that capital inflows can have large and expansionary effects on GDP if not offset by monetary policy. For countries with fixed exchange rates the rise in GDP could be as high as 0.7 percent after 18 months, driven largely by investment and exports. For non-fixing countries, the expansionary pressure appears to be neutralized by monetary policy yielding rates approximately 50 basis points higher than if no inflow occurred.

To the extent that equity inflows being expansionary implies other types of inflows are as well, this reinforces the view that “risk-off” episodes are contractionary globally – the loss

of financial flows likely dominates any expenditure-switching channels from exchange rate depreciations. This raises questions as to the optimality of these inflows – do we want macroeconomic booms following equity inflow events, or are these best avoided? Furthermore, it suggests an asymmetry in monetary policy impacts similar to those discussed in Rey (2012) and Gourinchas (2018) – changes in the center country’s monetary policy, by slowing capital flows to other countries, is likely contractionary everywhere (despite the center country’s currency appreciating). By comparing outcomes between fixing and non-fixing countries, my results suggest that domestic monetary policy still works as conventionally understood. Putting these two together then implies non-center countries may respond by *loosening* monetary policy in reaction to a monetary tightening of the center country.

Finally there is the issue of capital flow measures. Given investors may be fickle, does the fact that equity inflows are expansionary warrant the introduction of capital flow measures to stem the macro effects of risk-on/risk-off dynamics? Or are these measures more likely to cause harm than good?

References

- Adão, Rodrigo, Michal Kolesár, and Eduardo Morales**, “Shift-Share Designs: Theory and Inference,” *The Quarterly Journal of Economics*, 08 2019, 134 (4), 1949–2010.
- Ananchotikul, Nasha and Ms Longmei Zhang**, *Portfolio Flows, Global Risk Aversion and Asset Prices in Emerging Markets* number 14-156. In ‘1.’, International Monetary Fund, 2014.
- Anderson, Theodore W. and Herman Rubin**, “Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations,” *The Annals of Mathematical Statistics*, 1949, 20 (1), 46–63.
- Bacchetta, Philippe, Simon Tišče, and Eric van Wincoop**, “International Portfolio Choice with Frictions: Evidence from Mutual Funds,” Swiss Finance Institute Research Paper 46 June 2020.
- Bartik, Timothy J.**, *Who Benefits from State and Local Economic Development Policies?*, W.E. Upjohn Institute for Employment Research, September 1991.
- Baskaya, Yusuf Soner, Julian di Giovanni, Kalemli-Özcan, José-Luis Peydro, and Mehmet Fatih Ulu**, “Capital Flows and the International Credit Channel,” *Journal of International Economics*, 2017, 108, S15–S22. 39th Annual NBER International Seminar on Macroeconomics.
- Blanchard, Olivier, Jonathan D Ostry, Atish R Ghosh, and Marcos Chamon**, “Capital Flows: Expansionary or Contractionary?,” *American Economic Review*, 2016, 106 (5), 565–69.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-Experimental Shift-Share Research Designs,” *The Review of Economic Studies*, 06 2021, 89 (1), 181–213.
- Broner, Fernando, Alberto Martin, Lorenzo Pandolfi, and Tomás Williams**, “Winners and Losers from Sovereign Debt Inflows,” *Journal of International Economics*, 2021, 130, 103446. NBER International Seminar on Macroeconomics 2020.
- Cesa-Bianchi, Ambrogio, Andrea Ferrero, and Alessandro Rebucci**, “International Credit Supply Shocks,” *Journal of International Economics*, 2018, 112, 219–237.
- Converse, Nathan L., Eduardo Levy Yeyati, and Tomás Williams**, “How ETFs Amplify the Global Financial Cycle in Emerging Markets,” *Review of Financial Studies*, forthcoming.

- Coppola, Antonio, Matteo Maggiori, Brent Neiman, and Jesse Schreger**, “Redrawing the Map of Global Capital Flows: The Role of Cross-Border Financing and Tax Havens*,” *The Quarterly Journal of Economics*, 05 2021, 136 (3), 1499–1556.
- Davis, J. Scott**, “The Macroeconomic Effects of Debt- and Equity-Based Capital Inflows,” *Journal of Macroeconomics*, 2015, 46, 81–95.
- Elton, Edwin J. and Martin J. Gruber**, “Mutual Funds,” in G.M. Constantinides, M. Harris, and R. M. Stulz, eds., *Handbook of the Economics of Finance*, Vol. 2 of *Handbook of the Economics of Finance*, Elsevier, 2013, chapter 0, pp. 1011–1061.
- Fernandez-Arias, Eduardo**, “The New Wave of Private Capital Inflows: Push or Pull?,” *Journal of Development Economics*, 1996, 48 (2), 389–418.
- Forbes, Kristin J and Francis E Warnock**, “Capital Flow Waves: Surges, Stops, Flight, and Retrenchment,” *Journal of International Economics*, 2012, 88 (2), 235–251.
- Forbes, Kristin J. and Francis E. Warnock**, “Debt-and Equity-Led Capital Flow Episodes,” in Miguel Fuentes D., Claudio E. Raddatz, and Carmen M. Reinhart, eds., *Capital Mobility and Monetary Policy*, Vol. 18 of *Central Banking, Analysis, and Economic Policies Book Series*, Central Bank of Chile, April 2014, chapter 9, pp. 291–322.
- Fratzscher, Marcel, Marco Lo Duca, and Roland Straub**, “On the International Spillovers of US Quantitative Easing,” *The Economic Journal*, 2017, 128 (608), 330–377.
- Gabaix, Xavier and Matteo Maggiori**, “International Liquidity and Exchange Rate Dynamics,” *The Quarterly Journal of Economics*, 2015, 130 (3), 1369–1420.
- **and Ralph S. J Koijen**, “Granular Instrumental Variables,” Working Paper 28204, National Bureau of Economic Research December 2020.
- **and Ralph S. J. Koijen**, “In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis,” Swiss Finance Institute Research Paper 91 May 2022.
- Galí, Jordi and Tommaso Monacelli**, “Monetary Policy and Exchange Rate Volatility in a Small Open Economy,” *The Review of Economic Studies*, 07 2005, 72 (3), 707–734.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik Instruments: What, When, Why, and How,” Technical Report 8 August 2020.

Gopinath, Gita, Sebnem Kalemli-Ozcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez, “Capital Allocation and Productivity in South Europe*,” *The Quarterly Journal of Economics*, 06 2017, 132 (4), 1915–1967.

Gourinchas, Pierre-Olivier, “Monetary Policy Transmission in Emerging Markets: An Application to Chile,” 2018.

Gruber, Martin J., “Another Puzzle: The Growth in Actively Managed Mutual Funds,” *The Journal of Finance*, 1996, 51 (3), 783–810.

Gyöngyösi, Győző and Emil Verner, “Household Debt Revaluation and the Real Economy: Evidence from a Foreign Currency Debt Crisis,” *American Economic Review*, September 2020, 110 (9), 2667–2702.

Ilzetzki, Ethan, Carmen M. Reinhart, and Kenneth S. Rogoff, “Exchange Arrangements Entering the 21st Century: Which Anchor Will Hold?,” NBER Working Papers 23134, National Bureau of Economic Research, Inc February 2017.

Itskhoki, Oleg and Dmitry Mukhin, “Mussa Puzzle Redux,” Working Paper 28950, National Bureau of Economic Research June 2021.

Jaeger, David A., Joakim Ruist, and Jan Stuhler, “Shift-Share Instruments and the Impact of Immigration,” NBER Working Papers 24285, National Bureau of Economic Research, Inc February 2018.

Jordà, Òscar, “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 2005, 95 (1), 161–182.

Jordà, Òscar, “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, March 2005, 95 (1), 161–182.

Jotikasthira, Chotibhak, Christian Lundblad, and Tarun Ramadorai, “Asset Fire Sales and Purchases and the International Transmission of Funding Shocks,” *The Journal of Finance*, 2012, 67 (6), 2015–2050.

—, —, —, and —, “Asset Fire Sales and Purchases and the International Transmission of Funding Shocks,” *The Journal of Finance*, 2012, 67 (6), 2015–2050.

Kalemli-Özcan, Şebnem, “U.S. Monetary Policy and International Risk Spillovers,” in “Jackson Hole Economic Symposium” 2019.

Kalemli-Ozcan, Sebnem, Xiaoxi Liu, and Ilhyock Shim, “Exchange Rate Fluctuations and Firm Leverage,” *IMF Economic Review*, March 2021, 69 (1), 90–121.

Koepke, Robin, “What Drives Capital Flows to Emerging Markets? A Survey of the Empirical Literature,” 2015.

Kolesar, Michal, Raj Chetty, John N Friedman, Edward L Glaeser, and Guido W Imbens, “Identification and Inference with Many Invalid Instruments,” Working Paper 17519, National Bureau of Economic Research October 2011.

Maggiore, Matteo, Brent Neiman, and Jesse Schreger, “International Currencies and Capital Allocation,” *Journal of Political Economy*, 2020, 128 (6), 2019–2066.

Mendoza, Enrique G., “Sudden Stops, Financial Crises, and Leverage,” *American Economic Review*, December 2010, 100 (5), 1941–66.

Miao, Yanliang and Malika Pant, “Coincident Indicators of Capital Flows,” IMF Working Papers 2012/055, International Monetary Fund 2012.

Miranda-Agrippino, Silvia and Hélène Rey, “U.S. Monetary Policy and the Global Financial Cycle,” *The Review of Economic Studies*, 05 2020, 87 (6), 2754–2776.

Pandolfi, Lorenzo and Tomás Williams, “Real Effects of Sovereign Debt Inflow Shocks,” *AEA Papers and Proceedings*, May 2020, 110, 511–15.

Raddatz, Claudio and Sergio Schmukler, “On the International Transmission of Shocks: Micro-evidence from Mutual Fund Portfolios,” *Journal of International Economics*, 2012, 88 (2), 357–374.

—, **Sergio L. Schmukler, and Tomás Williams**, “International Asset Allocations and Capital Flows: The Benchmark Effect,” *Journal of International Economics*, 2017, 108, 413–430.

Reinhart, Carmen M and Vincent R Reinhart, “Capital Flow Bonanzas: An Encompassing View of the Past and Present,” Working Paper 14321, National Bureau of Economic Research September 2008.

Rey, Hélène, “Dilemma Not Trilemma: The Global Financial Cycle and Monetary Policy Independence,” Technical Report, Jackson Hole Economic Symposium 2012.

Salomao, Juliana and Liliana Varela, “Exchange Rate Exposure and Firm Dynamics [Credit Constraints and the Cyclicity of R&D Investment: Evidence from France],” *Review of Economic Studies*, 2022, 89 (1), 481–514.

Williams, Tomás, “Capital Inflows, Sovereign Debt and Bank Lending: Micro-Evidence from an Emerging Market,” *The Review of Financial Studies*, 03 2018, 31 (12), 4958–4994.

Zeckhauser, R., J. Patel, and D. Hendricks, “Nonrational Actors and Financial Market Behavior,” *Theory and Decision*, 1991, 31, 257–287.

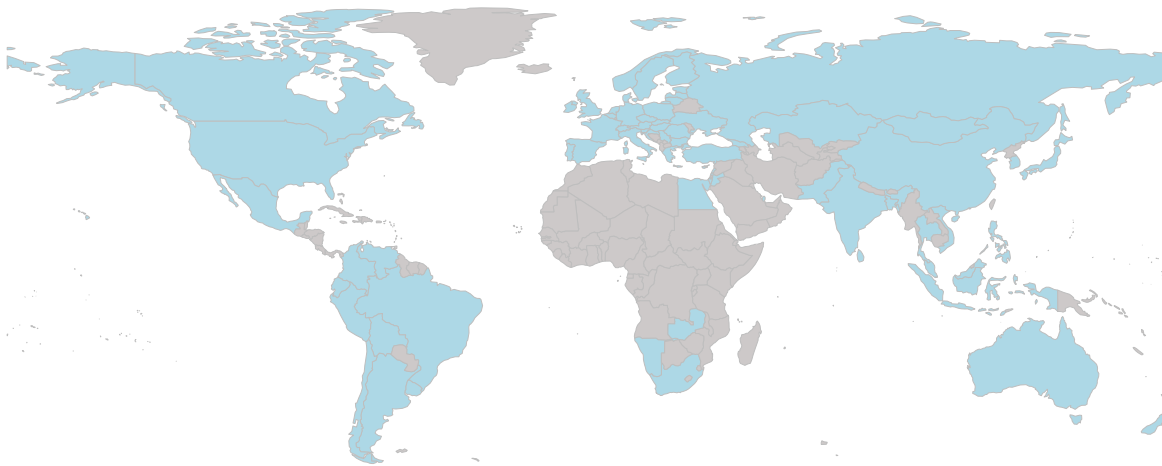
Appendices

A Data

A.1 Details on Macro Data

Figure 16 and Table 2 show the countries included in the sample. There are 71 countries made up of 31 advanced and 40 emerging economies. Most of the world is covered fairly well except for Africa.

Figure 16: Countries in the sample



This graph shows in blue the set of countries in the sample. This is based on whether a country is included in the regressions of real GDP on equity inflows for at least one base period of the instrument.

In the baseline estimates the United States is included, and this may appear problematic given the US is atypical because it is often the source of capital inflows to other countries and the destination when inflows into other countries dry up. The inclusion or otherwise of the US has little effect on my estimates, and so I default to include them. The reason the US has little effect on my estimates is because each country is given equal weight in the regression analysis (and all variables analyzed are scaled to be size independent). As such, the US's inclusion or exclusion effectively removes 1 of the 70 countries in the sample and should not be expected to have a large impact on the estimates.

Table 3 shows the sources for the data used in the paper. Where two sources exist for a series, the longer series is chosen and there is no splicing. For the exchange rate regime

Table 2: Countries in sample

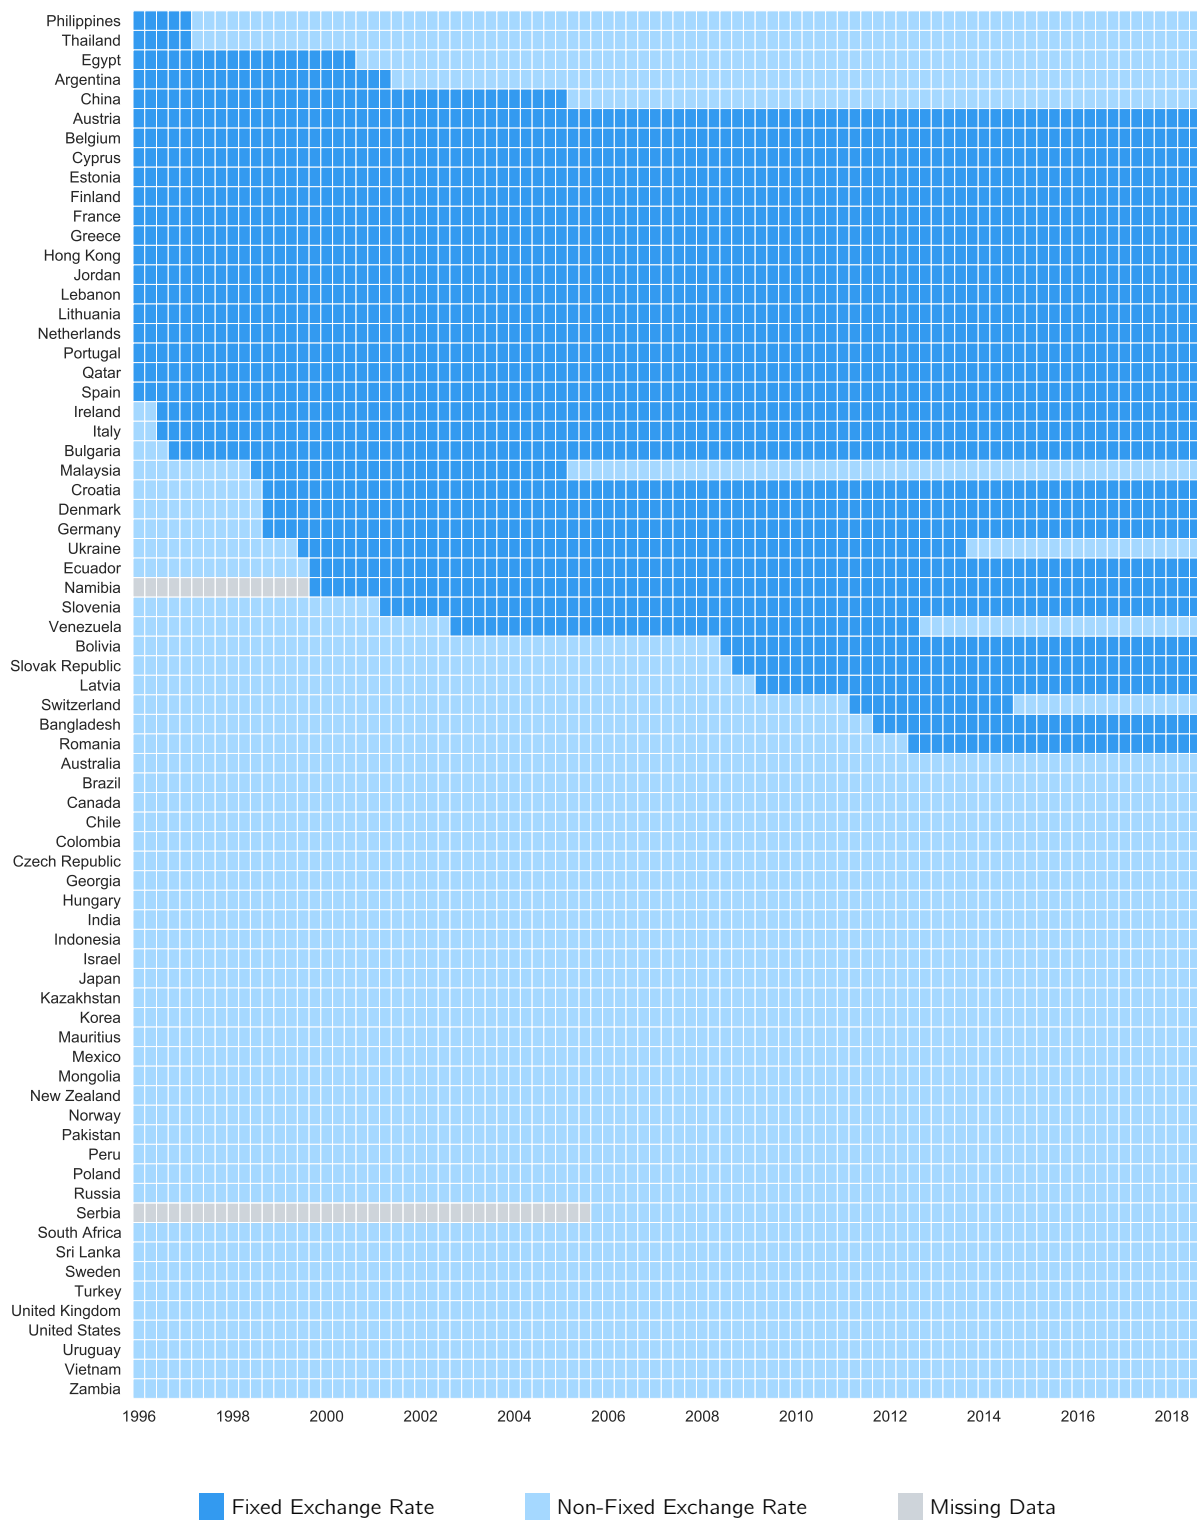
<i>Advanced Economies</i>				
Australia	Estonia	Israel	New Zealand	Switzerland
Austria	Finland	Italy	Norway	United Kingdom
Belgium	France	Japan	Portugal	United States
Canada	Germany	Korea	Slovak Republic	
Cyprus	Greece	Latvia	Slovenia	
Czech Republic	Hong Kong	Lithuania	Spain	
Denmark	Ireland	Netherlands	Sweden	
<i>Middle-Income Countries</i>				
Argentina	Ecuador	Lebanon	Philippines	Thailand
Bolivia	Egypt	Malaysia	Poland	Turkey
Brazil	Georgia	Mauritius	Qatar	Ukraine
Bulgaria	Hungary	Mexico	Romania	Uruguay
Chile	India	Mongolia	Russia	Venezuela
China	Indonesia	Namibia	Serbia	
Colombia	Jordan	Pakistan	South Africa	
Croatia	Kazakhstan	Peru	Sri Lanka	
<i>Low-Income Developing Countries</i>				
Bangladesh	Vietnam	Zambia		

classification the data ends in 2016, whereas the sample used here continues on to 2018 Q4. To fill in the missing data, I assume that exchange rate regimes are unchanged from the last value in the data. I collapse the coarse classification of Ilzetki et al. (2017) into two categories: fixers (coarse classification of 1) and non-fixers (coarse classification larger than 1). Figure 17 shows the classifications by country and time.

A.2 Details on EPFR Data

In this section I describe some more details of the EPFR data used to construct the mutual fund-based instrument. Essentially, there are three potential issues in the data. The first issue speaks to whether the EPFR data matches the Balance of Payments flows, given it deals exclusively with mutual fund flows. Studies using EPFR data (Miao and Pant (2012); Fratzscher et al. (2017); Ananchotikul and Zhang (2014)) show that there is a fairly strong correlation between aggregate EPFR and BOP flows (in particular see Figure 4 of Miao and

Figure 17: Exchange rate regimes by country and time



Notes: Data obtained from Ilzetzki et al. (2017). These data end in 2016, so for all subsequent dates, I assume that the exchange rate regime for each country remains as it was at the end of the recorded sample.

Table 3: Data sources

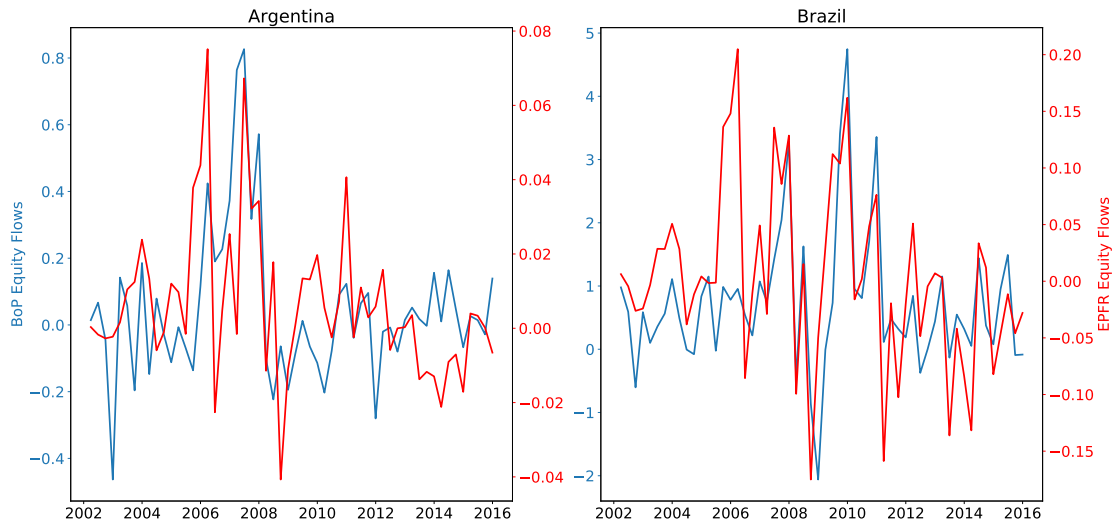
<i>Variable</i>	<i>Source:</i>
Equity Inflows/Outflows	IMF BoP Database
Debt Inflows/Outflows	
All Inflows/Outflows	
Broad Nominal Exchange Rates	BIS
Equity Returns	Global Financial Database
Equity Prices	
3-Month Interest Rates	
Real GDP	Datastream/IMF's IFS Statistics
Real Consumption	
Real Investment	
Real Exports	
Real Imports	
CPI	IMF IFS Statistics
Exchange Rate Regimes	Ilzetzki et al. (2017)

Pant, 2012).

Another concern in the data is the size of the country allocations database. Only approximately 10 percent of funds in the FF-ShareClass database are represented in the country allocations database. Some of the missing funds are dedicated funds – they invest in only one country – but a large number of global and global emerging market (GEM) funds (approximately 80 percent) do not report their portfolios. For the approach used in this paper having accurate portfolio data is essential, and as such the sample of funds I can use is considerably smaller than some other papers using EPFR data.⁴⁶ For what follows I will ignore the funds that do not have data on country allocations.

⁴⁶When EPFR compute aggregate country flows, they match funds with no entry in the country allocations database to a portfolio based on funds in the country allocations database with the same geographic focus (geo-focus). For example, suppose there is an unmatched fund with Latin America as the geographic focus. They take the average country allocations of all funds with that geo-focus to assign the allocation for the fund with missing country allocations data.

Figure 18: Scaling capital flows with Argentina and Brazil



Source: IMF BOP database and EPFR. Both flows are relative to GDP.

The final issue with the data is an issue of asymmetric coverage: flows from the EPFR data cover more of some country's flows than others (Figure 9 in the main paper shows this in terms of asset positions relative to total foreign equity investment in a country, and the same issue holds for inflows). For example, let us take two countries, Argentina and Brazil. Figure 18 shows the time series variation in the EPFR fund flows relative to GDP (in red, right scale) and the Balance of Payments inflows (in blue, left scale). While there is very similar cyclical behavior of the EPFR and total equity inflows into both countries, there is approximately a 10:1 ratio of total to EPFR flows for Argentina, but a 25:1 ratio for Brazil.

Estimating a first stage without accounting for this variation in coverage by country will lead to misleading standard errors. In the cases of Argentina and Brazil, running a first stage for Argentina only would suggest a coefficient near 10, but Brazil would require a coefficient closer to 25. Pooling both countries into one sample (which I do with 70 countries) would produce a coefficient somewhere between 10 and 25; this would underfit the variation for Brazil and overfit it for Argentina. These deviations would be reflected in the standard errors of the coefficients and would produce, in my view, misleading F-statistics.

To account for the scale of magnitude difference in the flows, I re-scale the instrument so that the standard deviation of overall inflows and the instrument are equal country by country.

A.3 Computing Predicted Inflows

In this subsection I provide some more details about how the predicted inflows term $\text{inflows}_{k \setminus r, t}$ is calculated. First, while in the EPFR data we observe the overall return (both portfolio and currency) in dollars, there is not data available on how much each mutual fund earned in one country versus another. However, it is possible to estimate the average return earned by each fund in each country by regressing the overall return on the portfolio shares of each fund in each country:

$$r_{k,t} = \alpha_k + \sum_{\tau=1}^T \sum_{i=1}^N \omega_{i \leftarrow k, \tau} \times \mathbf{1}_{t=\tau} \underbrace{r_{i,\tau}}_{\text{coefficients}} + e_{k,t}$$

The coefficients $\hat{r}_{i,\tau}$ then represent the average interest rate earned in country i at time τ by mutual funds in the sample.

Once we have estimates of the country-level returns, I then run the regression in Equation 7 (repeated here for convenience):

$$\frac{\text{inflows}_{k,t}}{A_{k,t}} = \gamma_t + \sum_{p=1}^{10} \delta_p r_{k,t-p} + e_{k,t}$$

where the number of lags (10) was chosen by the Bayesian Information Criterion.

One further complication is that the mutual funds in the EPFR data can often receive very large inflows that, depending on the timing of when they are received during the month, can make dividing by the beginning-of-month asset position an inaccurate method to measure returns. To demonstrate with a concrete example, let's consider a mutual fund with \$100 of assets under management at the beginning of the month. The very same day, imagine that it receives another \$100 of inflows so that for most of the month its assets under management is \$200. Then imagine it makes a 5 percent return on investment at the end of the month, earning \$10. Using the beginning-of-month assets under management of \$100 would lead this return to be overstated as 10 percent instead of 5 percent, making the return distribution across mutual funds unrealistically wide.

To correct for this, I assume that inflows and returns accrue at a constant rate during the month. While this may not be accurate for many individual instances, it serves to smooth out the more volatile values of flow rates or returns that would be obtained if one divided by the beginning-of-period assets.

The approach taken assumes that:

$$dA_{k,t} = (1 + \bar{r}_k + \bar{f}_k)A_{k,t}$$

where $A_{k,t}$ are assets under management at time t for mutual fund k , \bar{r}_k is the instantaneous return that that mutual fund earns (assumed constant within a quarter) and \bar{f}_k is the instantaneous inflow rate of mutual fund k (also assumed constant within a quarter). Solving this first order ODE yields:

$$\frac{A_{k,t+1} - A_{k,t}}{A_{k,t}} = e^{1+\bar{r}+\bar{f}} \quad (13)$$

We also observe in the data

$$Flow_{k,t} = \bar{f}_k \int_0^1 A_{k,t+s} ds = \bar{f}_k A_{k,t} e^{1+\bar{r}_k+\bar{f}_k} \quad (14)$$

$$Return_{k,t} = \bar{r}_k \int_0^1 A_{k,t+s} ds = \bar{r}_k A_{k,t} e^{1+\bar{r}_k+\bar{f}_k} \quad (15)$$

Putting Equations 13, 14 and 15 together gives the following expressions for the geometric average flow and return rates:

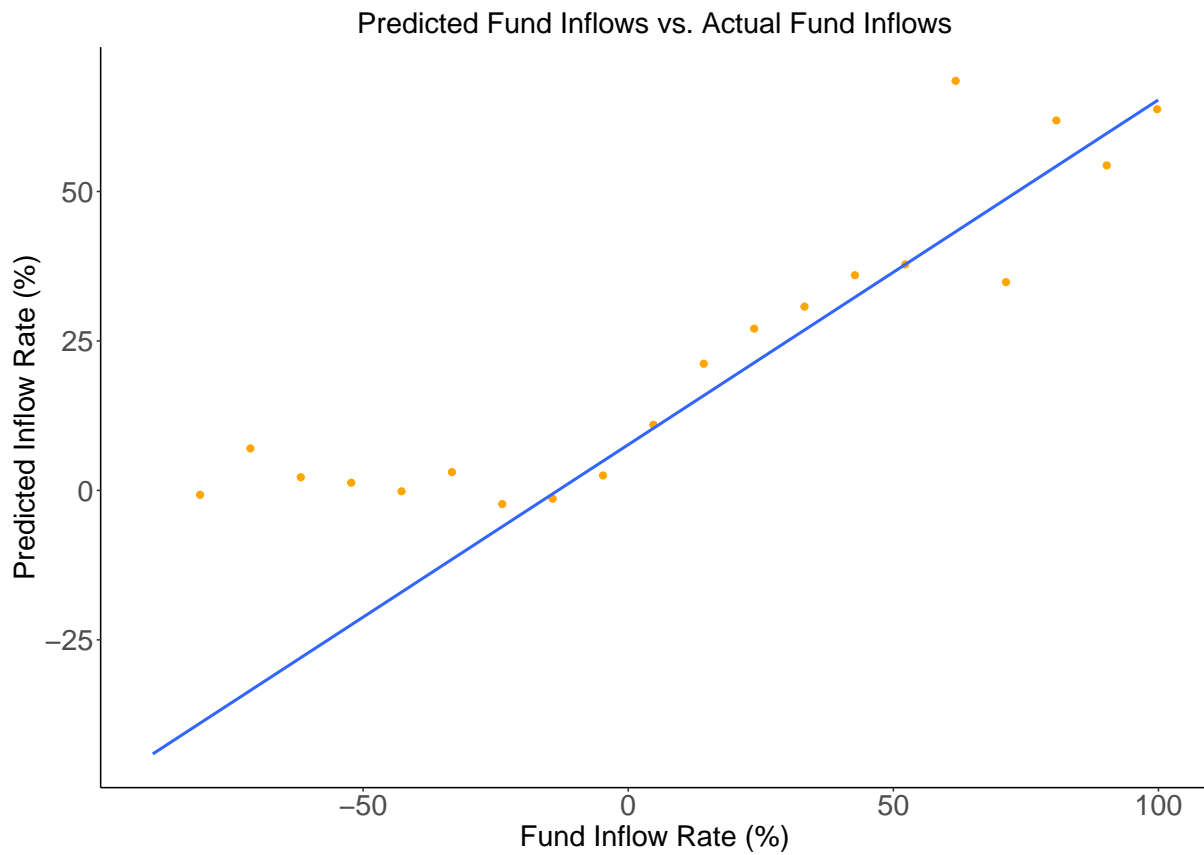
$$\bar{f}_k = \frac{Flow_{k,t}}{A_{k,t+1} - A_{k,t}}$$

$$\bar{r}_k = \frac{Return_{k,t}}{A_{k,t+1} - A_{k,t}}$$

Calculating the flow and return rates this way reduces to a large extent the outliers in both. With these, I then run the regressions in Equation 7 and then use Equations 8 and 9 to construct predicted inflows.

The fit of these predicted inflows is shown in Figure 19, which plots a bin-scatter with the actual inflow rate on the horizontal against the predicted inflow rate on the vertical. For positive inflow rates the fit is excellent, and there is some evidence that the fit is less good for negative inflow rates. However, the bulk of actual inflows are positive, which leads a line of best fit to put almost all weight on the data where inflow rates are positive. The slope of the line of best fit is about 0.7, which connotes an excellent fit of the predicted inflows to actual inflows.

Figure 19: Bin-scatter of the fit between actual and predicted inflows



Notes: This figure shows a bin-scatter with 30 evenly spaced bins of the fund inflow rate on the horizontal axis against the predicted inflow rate on the vertical axis. A linear line of best fit is shown in blue.

B More Details on Shift-Share Identification

B.1 Mathematical Treatment of Identification

The purpose of this section is to contextualize the identification requirements discussed in this paper with recent work in the literature on shift-share instruments. To do so I will use the framework introduced in Kolesar et al. (2011).

Consider the following system of second stage and first stage equations (omitting controls and fixed effects for simplicity):

$$\begin{aligned} \frac{Y_{i,t+h} - Y_{i,t-1}}{Y_{i,t-1}} &= \beta^h \frac{\text{inflows}_{i,t}}{Y_{i,t}} + \epsilon_{i,t+h} && \text{Second Stage} \\ \frac{\text{inflows}_{i,t}}{Y_{i,t}} &= a + b \underbrace{\sum_{k \in \mathcal{K} \setminus \mathcal{K}(i)} \frac{\omega_{i \leftarrow k, t_0} \text{inflows}_{k,t}}{Y_{i,t}}}_{\text{Instrument}} + \underbrace{\sum_{s=0}^S \gamma_s E_t(Y_{i,t+s})}_{\text{Endogeneity}} + e_{i,t} && \text{First Stage} \end{aligned}$$

For convenience, let me start by re-writing the first stage to have the following split between the shares and the shifters:

$$\frac{\text{inflows}_{i,t}}{Y_{i,t}} = a + b \boldsymbol{\omega}'_{i,t_0,t} \mathbf{FundFlows}_t + \sum_{s=0}^S \gamma_s E_t Y_{i,t+s} + e_{i,t}$$

where $\boldsymbol{\omega}_{i,t_0,t} = \left\{ \frac{\omega_{i \leftarrow k, t_0} A_{k,t}}{Y_{i,t}} \right\}_{k \in \mathcal{K} \setminus \mathcal{K}(i)}$, $\mathbf{FundFlows}_t = \left\{ \frac{\text{inflows}_{k,t}}{A_{k,t}} \right\}_{k \in \mathcal{K} \setminus \mathcal{K}(i)}$ and $\{\gamma_s\}_{s=0}^S$ determines how inflows respond to news about future Y .

Note that the shifter has now been written as the inflow *rate* from clients (or the amount of inflows relative to the assets managed by the fund). As discussed, it is more plausible that the inflow *rate* is randomly assigned to mutual funds conditional on shares rather than the level of inflows.

Without loss of generality, the second stage error term can be decomposed into the following terms:

$$\epsilon_{i,t+h} = \boldsymbol{\omega}'_{i,t_0,t} \boldsymbol{\Gamma}_t^h + \vartheta_{i,t+h} \quad (16)$$

where $\boldsymbol{\Gamma}_t^h$ is an arbitrary $\mathcal{K} \setminus \mathcal{K}(i) \times 1$ vector.

In this formulation, the error term contains a structure that correlates with the portfo-

lio shares of mutual funds. In general, this decomposition of the second stage error term is without loss of generality,⁴⁷ and in this context can be assigned an ex-ante or ex-post interpretation.

Under the ex-ante interpretation, a high $\Gamma_{k,t}^h$ would reflect information advantages of a particular mutual fund. If one mutual fund k is associated with a large $\Gamma_{k,t}$, then it might have a great deal of information about which countries will have the largest shocks. If it wants to be exposed to these shocks (which would correspond to $\Gamma_{k,t}^h \gg 0$), then it would likely choose a portfolio $\{\omega_{i \leftarrow k, t_0}\}_{i \in \mathcal{I}}$ that will correlate positively with the countries with the best shocks. Another way to conceptualize a non-zero $\Gamma_{k,t}^h$ might be a measure of ex-post connection to the error term. In this conceptualization, mutual funds pick different portfolios, and some get lucky and have a positive loading on future shocks. $\Gamma_{k,t}^h$ then measures how the portfolio of mutual fund k correlated with later error terms. Whether an information advantage exists or not, one can ask what conditions must be imposed on $\mathbf{\Gamma}_t^h$ to achieve identification in this system.

In Goldsmith-Pinkham et al. (2020), they essentially argue that the most sensible identification assumption in many contexts is that the shares in the shift-share instruments be themselves exogenous (i.e., $\mathbf{\Gamma}_t^h = \mathbf{0} \forall t, h$). In the context here, this identification assumption would be too strict to be plausible as actively managed mutual funds have considerable incentive to obtain useful information about which countries to invest in.

One might hope that the results on efficient markets hypothesis (EMH) would suggest that this condition is plausible, however the EMH best put states that even if $\Gamma_{k,t} > 0$, then $E_t[\Gamma_{k,t+s}] = 0 \forall s$. In other words, the same mutual fund cannot systematically beat the market, but in one particular quarter it may receive some information that allows it to beat the market. Unfortunately, this particular identification condition requires $\Gamma_{k,t+s} = 0$.

Borusyak et al. (2021) and Adão et al. (2019) argue that a weaker identification assumption is possible to still achieve identification: that $\mathbf{\Gamma}_t^h \perp \mathbf{FundFlows}_t \forall h$. In this context, the requirement is that the distribution of information across mutual funds (or ex-post exposure of mutual funds to shocks) is uncorrelated with the distribution of inflows into mutual funds. This is the identification condition I rely on in this paper. The insight for this identification assumption to work is that while $\text{Cov}(\omega_{i \leftarrow k, t_0} \Gamma_{k,t}^h, \omega_{i \leftarrow k, t_0} \frac{\text{inflow}_{k,t}}{A_{k,t}} | k) \neq 0$ for each individual mutual fund (since both terms contain the same vector of country shares), the exclusion condition is imposed at the aggregated level: $\text{Cov}(\sum_{k \in \mathcal{K} \setminus \mathcal{K}(i)} \omega_{i \leftarrow k, t_0, t} \Gamma_{k,t}^h, \sum_{k \in \mathcal{K} \setminus \mathcal{K}(i)} \omega_{i \leftarrow k, t_0, t} \frac{\text{inflow}_{k,t}}{A_{k,t}}) = 0$. The (weighted) sums can have 0 correlation even while each component is mechanically correlated provided there is no correlation between the two weights used – hence the independence

⁴⁷ $\mathbf{\Gamma}$ can be set to any value, including the zero vector.

assumption between fund inflows and the exposure term ($\Gamma_{k,t}$) – and that \mathcal{K} is large.⁴⁸

In the notation of Borusyak et al. (2021), the exclusion condition is that:

$$E \left[\frac{\text{inflows}_{k,t}}{A_{k,t}} \mid \frac{E_i(\omega_{i \leftarrow k, t_0, t} \epsilon_{i, t+h})}{E_i(\omega_{i \leftarrow k, t_0, t})} \right] = \mu_t^h \quad \forall k, t, h \in \{0, \dots, H\} \quad (17)$$

where μ_t^h is a time-specific (and horizon h -specific) constant.

This condition states that within each time period the shifters (in my context, fund inflows) need to be randomly assigned across funds conditional on each fund's exposure to future country shocks $\epsilon_{i, t+h}$. This is captured by the term $\frac{E_i(\omega_{i \leftarrow k, t_0} \epsilon_{i, t+h})}{E_i(\omega_i)}$, which maps each country shock to each mutual fund k based on k 's portfolio share invested in i in the base period.

Recall that the shifter as defined here is the inflow rate into each mutual fund. The reason this is a rate is that the shifter needs to be randomly assigned *across funds* within each time period t . Given that some funds are much larger than others for the entirety of the sample, using the dollar amount of inflows as the shifter would render the identification assumption implausible. However, by defining the shifter as being relative to the size of the fund, it is much more plausible that the shifter in this situation is randomly assigned.

Furthermore, refining the instrument to use lagged leave-out returns as predictions of the flow rate involves the simple replacement of the inflow rate with the fitted values from the regression in Equation 7. In this case, the random assignment condition can be stated as: Each mutual fund's exposure to lagged returns outside of i is randomly assigned with respect to future country i shocks.

B.1.1 Why Control for the Sum of Shares?

In Section 2.5 I mentioned that the Borusyak et al. (2021) identification requires controlling for the sum of shares in the IV regression using the shift-share instrument. The reason for this is that in this scenario not all investors in each country (indeed not even all mutual funds) are captured by my data. However, having missing funds is equivalent to including them with a shifter of 0. Given that the random assignment allows that the mean of the shifters be $\mu_t^h \neq 0$, with missing data it *requires* that the shifters (conditional on shock exposure) be mean 0. Given that the unconditional mean of the flow rate is not 0 this is clearly violated. Either de-meaning the shifter or controlling for the sum of the shares corrects this issue.

Fortunately, whether the sum of shares is included or not matters little for my findings,

⁴⁸There are additional regularity conditions required, mainly to ensure that each mutual fund remains small as $\mathcal{K} \rightarrow \infty$.

and given the technical arguments for it the results shown in Section 5 and the subsequent robustness checks include it.

B.1.2 Standard Errors under This Shift-Share Identification

While the identification assumption **FundFlows** $\perp\!\!\!\perp$ Γ is appropriate to recover the correct slope coefficient, the notion that the shares in the instrument could have non-zero weight in the second stage error term has implications for the appropriate standard errors. How to address this is discussed by Adão et al. (2019).

The intuition is that if two regions share common shocks, then one would want to cluster the standard errors with these two regions in one cluster. This is the logic behind choosing geographic clusters or country clusters in panel data – correlations between data within a cluster are unrestricted, and the resulting standard errors are robust to any such correlations. In the context discussed in Adão et al. (2019), Boston and California both have large information technology (IT) sectors, which means that any shock that affects IT is likely to affect both regions at the same time. This is a problem for inference even if it is not an issue for identification.

In this context, the key dimension connecting countries is mutual funds. Mutual funds are small and are unlikely to affect the distribution of country shocks in ways such as IT would for California or Boston. However, there still might be an issue if some mutual funds are able to predict country shocks in some way and tilt their portfolio shares towards countries with “good” shocks. In such a case there could be a connection between the vector of country shocks and some mutual fund’s portfolio shares.⁴⁹ While I do not expect this to have material impacts on the standard errors, the correction suggested in Adão et al. (2019) is robust to country-clustering when country fixed effects are included (which I do) and generalizes even Anderson-Rubin standard errors if the instrument is weak. As such, it is more conservative to produce confidence intervals with their standard errors.

I implement their (time-clustered) standard error with one slight difference: the correct standard error is calculated as follows:

$$\hat{se}(\beta) = \frac{\sqrt{\sum_{t=1}^T \sum_{(k,k') \in \mathcal{K} \setminus \mathcal{K}(i)} \hat{X}_{k,t} \hat{R}_{k,t} \hat{R}_{k',t} \hat{X}_{k',t}}}{\sum_{i=1}^N \dot{X}^2}$$

⁴⁹Note that if this occurs, it is not a threat to identification unless clients systematically invest in the mutual funds best able to do this.

where \hat{X} are residualized equity inflows, $\hat{R}_{k,t} = \sum_{i=1}^N \omega_{i \leftarrow k, t_0} \hat{\epsilon}_{i,t}$ where $\hat{\epsilon}_{i,t}$ are the regression residuals and $\hat{\chi}_{k,t}$ should be obtained by regressing \hat{X} on the shares $\omega_{i \leftarrow k, t_0}$. However, in my context I have fewer countries ($N = 67$) than mutual funds ($N \in [300, 1617]$), which makes such a regression infeasible. Fortunately, because $\hat{\chi}_{k,t} \in [0, FundFlows_{k,t}]$ and positively affects the standard error, I can simply use the raw shifters $\hat{\chi}_{k,t} = FundFlows_{k,t}$ to create conservative versions of Adão et al. (2019) standard errors.

B.2 Testing Predictions Consistent with Exclusion Holding

Borusyak et al. (2021) propose a tractable test to see if the distribution of fund exposures to country shocks (the vector Γ_t^h defined in Equation 16) – to the extent measurable – is correlated with the inflow rates across mutual funds. Note that the exclusion restriction can be expressed as follows:⁵⁰

$$\begin{aligned}
E[B_{i,t} \epsilon_{i,t+h}] &= E \left[\sum_k \frac{\omega_{i \leftarrow k, t_0} A_{k,t}}{Y_{i,t}} \frac{inflows_{k,t}}{A_{k,t}} \epsilon_{i,t+h} \middle| controls_{i,t} \right] \quad \forall h > 0 \\
&= \sum_k E \left[\frac{\omega_{i \leftarrow k, t_0} A_{k,t}}{Y_{i,t}} \frac{inflows_{k,t}}{A_{k,t}} \epsilon_{i,t+h} \middle| controls_{i,t} \right] \quad \forall h > 0 \\
&= \sum_k E \left[\underbrace{E_i \left(\frac{\omega_{i \leftarrow k, t_0} A_{k,t}}{Y_{i,t}} \epsilon_{i,t+h} \middle| k, t \right)}_{\equiv \phi_{k,t+s}} \frac{inflows_{k,t}}{A_{k,t}} \middle| controls_{i,t} \right] \quad \forall h > 0
\end{aligned}$$

The variable $\phi_{k,t+s}$ converts an error term that varies at the country level and time into a variable that varies over mutual funds and time. This is extremely convenient because mutual fund inflow rates ($\frac{inflows_{k,t}}{A_{k,t}}$) also vary over the same dimensions as this $\phi_{k,t}$ variable.

Defining $\phi_{k,t+s}$ this way conveniently summarizes the distribution of exposures of mutual funds to country shocks regardless of the connecting mechanism.

The Borusyak et al. (2021) test involves constructing an observable approximation of $\phi_{k,t+h} = E_i[\omega'_{i \leftarrow k, t_0} \epsilon_{i,t+h}]$ using a proxy for $\epsilon_{i,t+h}$: $\tilde{\epsilon}_{i,t+h}$. This proxy represents an observable variable that is strongly believed to be correlated with the error term in the second stage regression.⁵¹ With such a proxy variable I can run a simple regression of mutual fund exposure

⁵⁰Note that because in this context news violates exclusion and a dynamic response is being estimated, the exclusion restriction must be satisfied in terms of both current and future error terms.

⁵¹It is also important to choose a proxy that we believe is not affected causally by capital inflows. Otherwise, failing this test could simply reflect the inflow rates directing capital towards some countries and causing the proxy variable to move.

to this proxy $\tilde{\phi}_{t+s} = \sum_i \omega_{i \leftarrow k, t_0} \tilde{\epsilon}_{i, t+h}$ on the shifters (inflows $_{k,t}$):

$$\text{inflows}_{k,t} = \alpha + \delta \tilde{\phi}_{k,t+h} + \text{controls}_{k,t} + \nu_{k,t+h}$$

where $\text{controls}_{k,t} = \sum_i \omega_{i \leftarrow k, t_0} \text{controls}_{i,t}$ and come from the controls included in the first and second stage and are constructed as follows:

$$\text{controls}_{k,t+s} = \sum_i \omega_{i \leftarrow k, t_0} \text{controls}_{i,t}$$

A simple hypothesis test of $\delta = 0$ gives us a sense of whether the distribution of shifters ($\frac{\text{inflows}_{k,t}}{A_{k,t}}$) is correlated with the proxy of the error term ($\tilde{\phi}_{k,t+s}$). This is a particularly nice test because it takes a multi-dimensional identification condition (shares \times shifters \times errors) and simplifies it down to a single test statistic. If the null is rejected, then mutual funds with high information (sensible initial portfolios) tend to be the ones that receive future inflows. This would suggest that the instrument includes responses to news about each country and would not be exogenous. If the null cannot be rejected and the confidence interval of the estimates of δ are small, then we can be confident there is unlikely to be a large connection between mutual fund exposures to country shocks/news and their inflows.

The mechanisms that are concerning in the capital inflows context are:

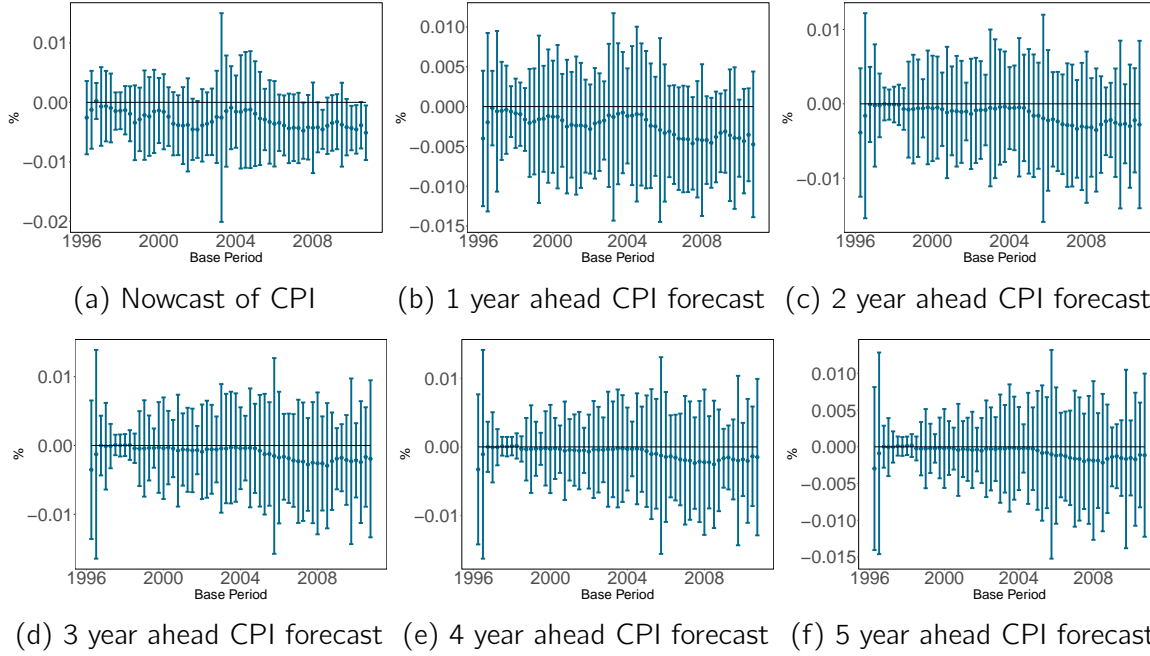
- Clients selecting mutual funds based on news about the countries they invest in. Here I find a measure of news that varies by country and create corresponding $\tilde{\phi}_{k,t}^{news}$ exposure measures.
- Clients can detect the “best” funds that are investing in the countries about to get the highest returns. In this case, I construct a measure summarizing exposure to future high country returns by estimating these returns from the overall returns in EPFR and portfolio shares. This is presented in Appendix C.1.

C Tests for Possible Violations of the Exclusion Condition

C.1 IMF forecasts and country level returns

Figure 20 shows the results when looking at mutual fund exposure to high IMF forecasts of CPI.

Figure 20: Correlation of fund inflows and IMF CPI forecasts



Notes: Each panel shows regressions from Equation 12 for a variety of base period choices when constructing the instrument. Each dot is a point estimate. Confidence intervals are 95 percent intervals. Standard errors are clustered by mutual fund and time.

C.2 Returns Test Using Mutual Fund Returns Data

As discussed in Section C.1, when testing whether inflows can predict future counterfactual fixed portfolio returns it is possible to calculate these fixed portfolio returns using the mutual fund returns data rather than the country index returns data used above. In this section, I show the results from using the direct mutual fund returns instead.

The country index measure assumes that within countries mutual funds all make the same return as either each country's main equity index or 10 year bond index. Using this measure has the advantage of relating to returns most likely to be correlated with a country-level error term; however, it assumes that mutual funds within a country essentially operate as index funds. Since this is unrealistic, the second measure is constructed from EPFR data on the overall returns made from each mutual fund.

Constructing this measure requires two steps. First I find the average return made by each mutual fund in the sample from country i in time period t by regressing the overall mutual fund returns on their portfolio shares in each country and in each time period:

$$r_{k,t} = \gamma_k + \sum_{\tau} \sum_i \beta_{i,\tau} \omega_{i \leftarrow k, \tau} \times \mathbb{1}_{\tau=t} + v_{i,t}$$

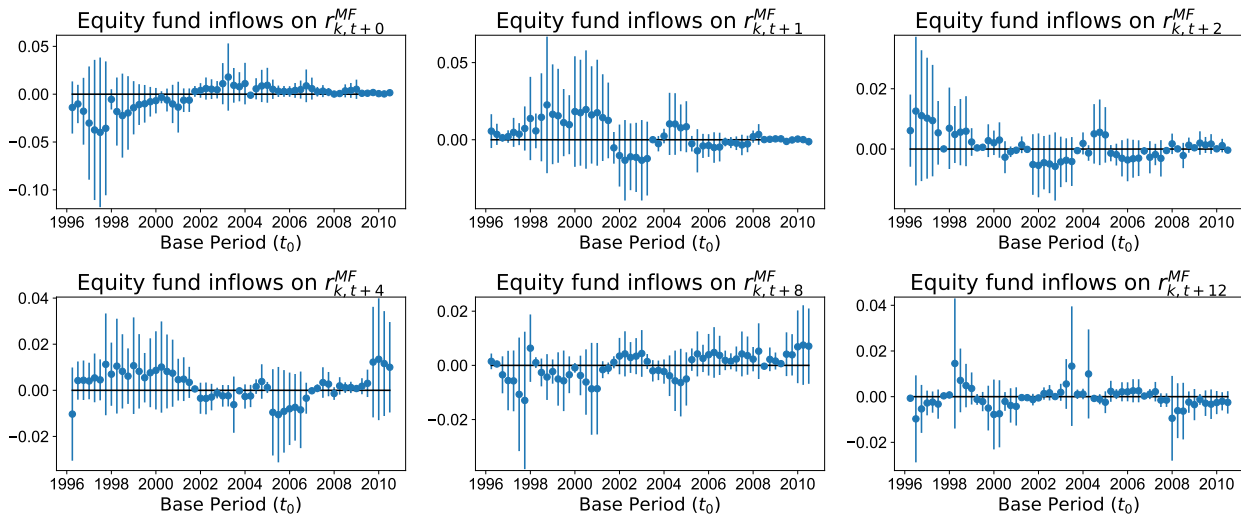
Note that there is a regression coefficient for each country at each time period in the sample ($\beta_{i,t}$). This will be used as the estimate of the counterfactual return made for investments in country i in time t . Essentially, this uses cross-sectional variation in returns and portfolios over mutual funds to identify the marginal effect an additional percentage point of investment in each country has on the return on the mutual fund. A fund fixed effect is included in the regression to control for any potential composition bias between the countries funds invest in and each fund's "quality" or ability to earn high returns. If, say, the best performing funds tend to invest more in Germany, then the counterfactual return for a low performing fund that shifted into Germany would be significantly less than the (share weighted) average returns of funds that invest in China. The fund fixed effect is my attempt to remove this confounding variation.

The estimated regression coefficients are then used calculate the counterfactual return of mutual funds based on base period portfolio weights:

$$\tilde{r}_{k,t+s}^{MP} = \sum_i \omega_{i \leftarrow k, t_0} \beta_{i,t+s}$$

The results are shown in Figure 21 and show no significant correlation between this measure of future fixed portfolio counterfactual returns and current fund inflows for either equity funds or bond funds across a variety of base period choices.

Figure 21: Correlation of fund inflows and imputed mutual fund returns



Notes: Dots are point estimates. Confidence intervals are 95 percent intervals. Standard errors are clustered by mutual fund.

C.3 Crisis Inflow Distributions

This subsection presents additional figures comparing the distribution of funds exposed to a crisis event to those not exposed. Figure 22 shows the same distributions for the affected funds both before and after the event. The panels in grey represent the distributions at the time of the crisis and are repeated from the diagonals in Figure 12. Since the crisis likely had implications for economic prospects for a considerable period after the genesis of the crisis, an appropriate view would be that exclusion would require all panels shown in this figure to not have systematically different mean fund inflow rates between exposed and unexposed funds. This is the case for all 32 panels with the exception of the Asia Crisis funds one period before the crisis where the difference in means is marginally significant. Strangely, this appears to be because unexposed funds experience extremely large inflows rather than due to exposed funds experiencing outflows. The number of funds in the sample not investing in the affected countries during this period is a mere 36, suggesting that this is an anomaly rather than a cause for concern.

Finally, Figures 23 and 24 show the same time distributions for each individual country affected by the Asia Crisis with similar inflow patterns comparing exposure to individual countries rather than the full set.

Figure 22: Inflow rate distributions for all funds during three crises

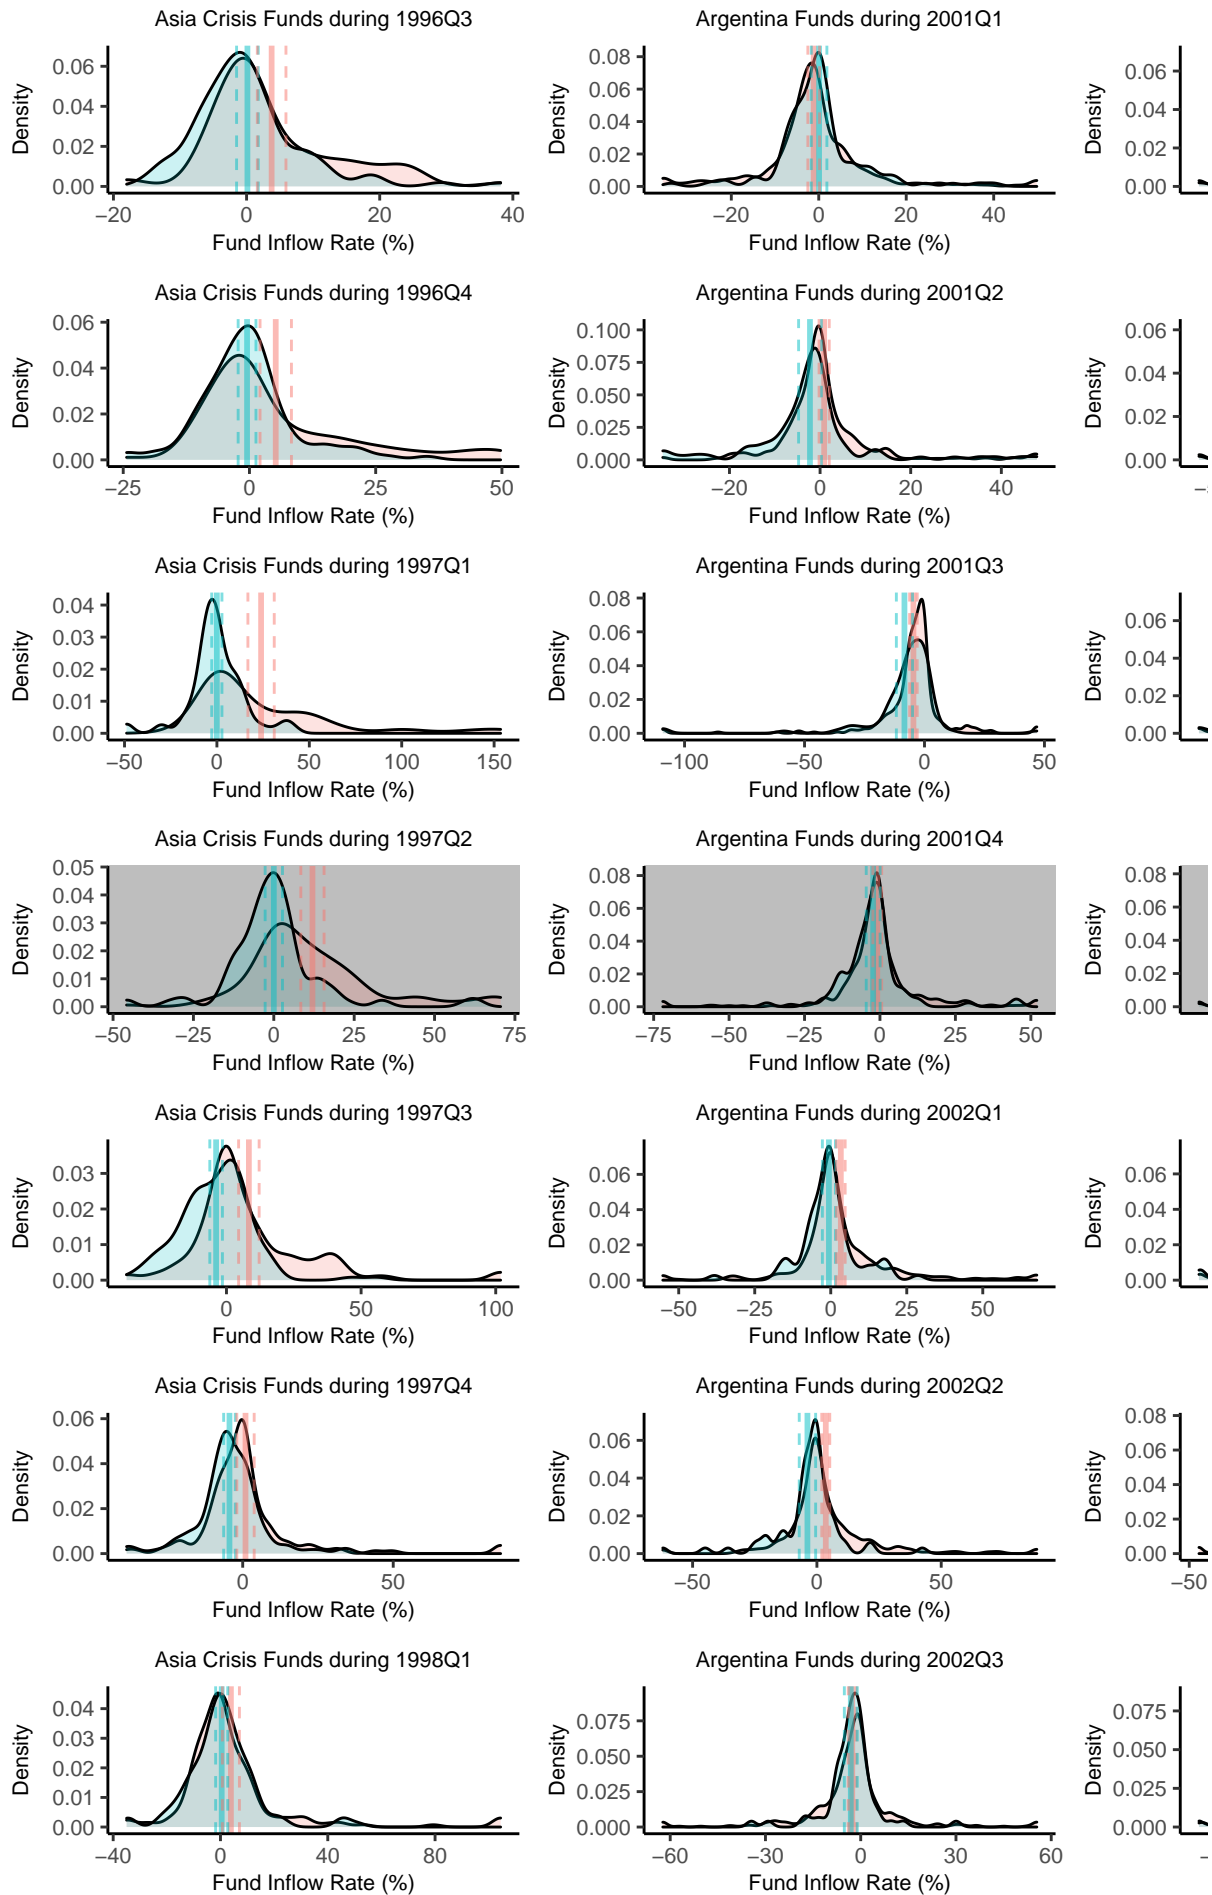
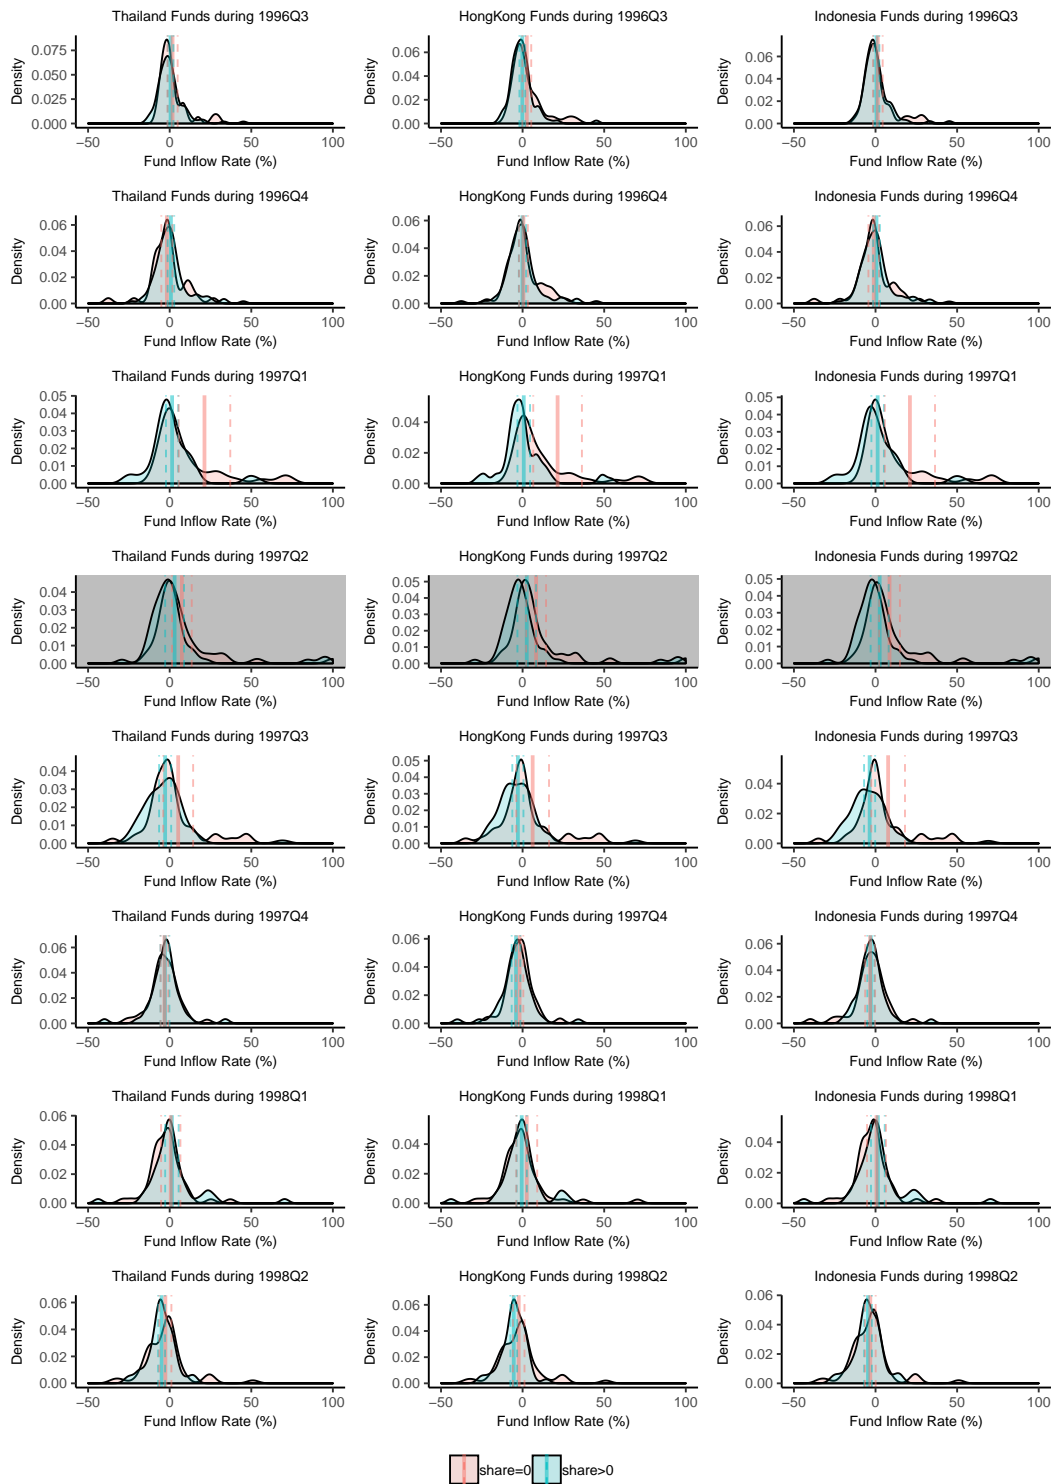
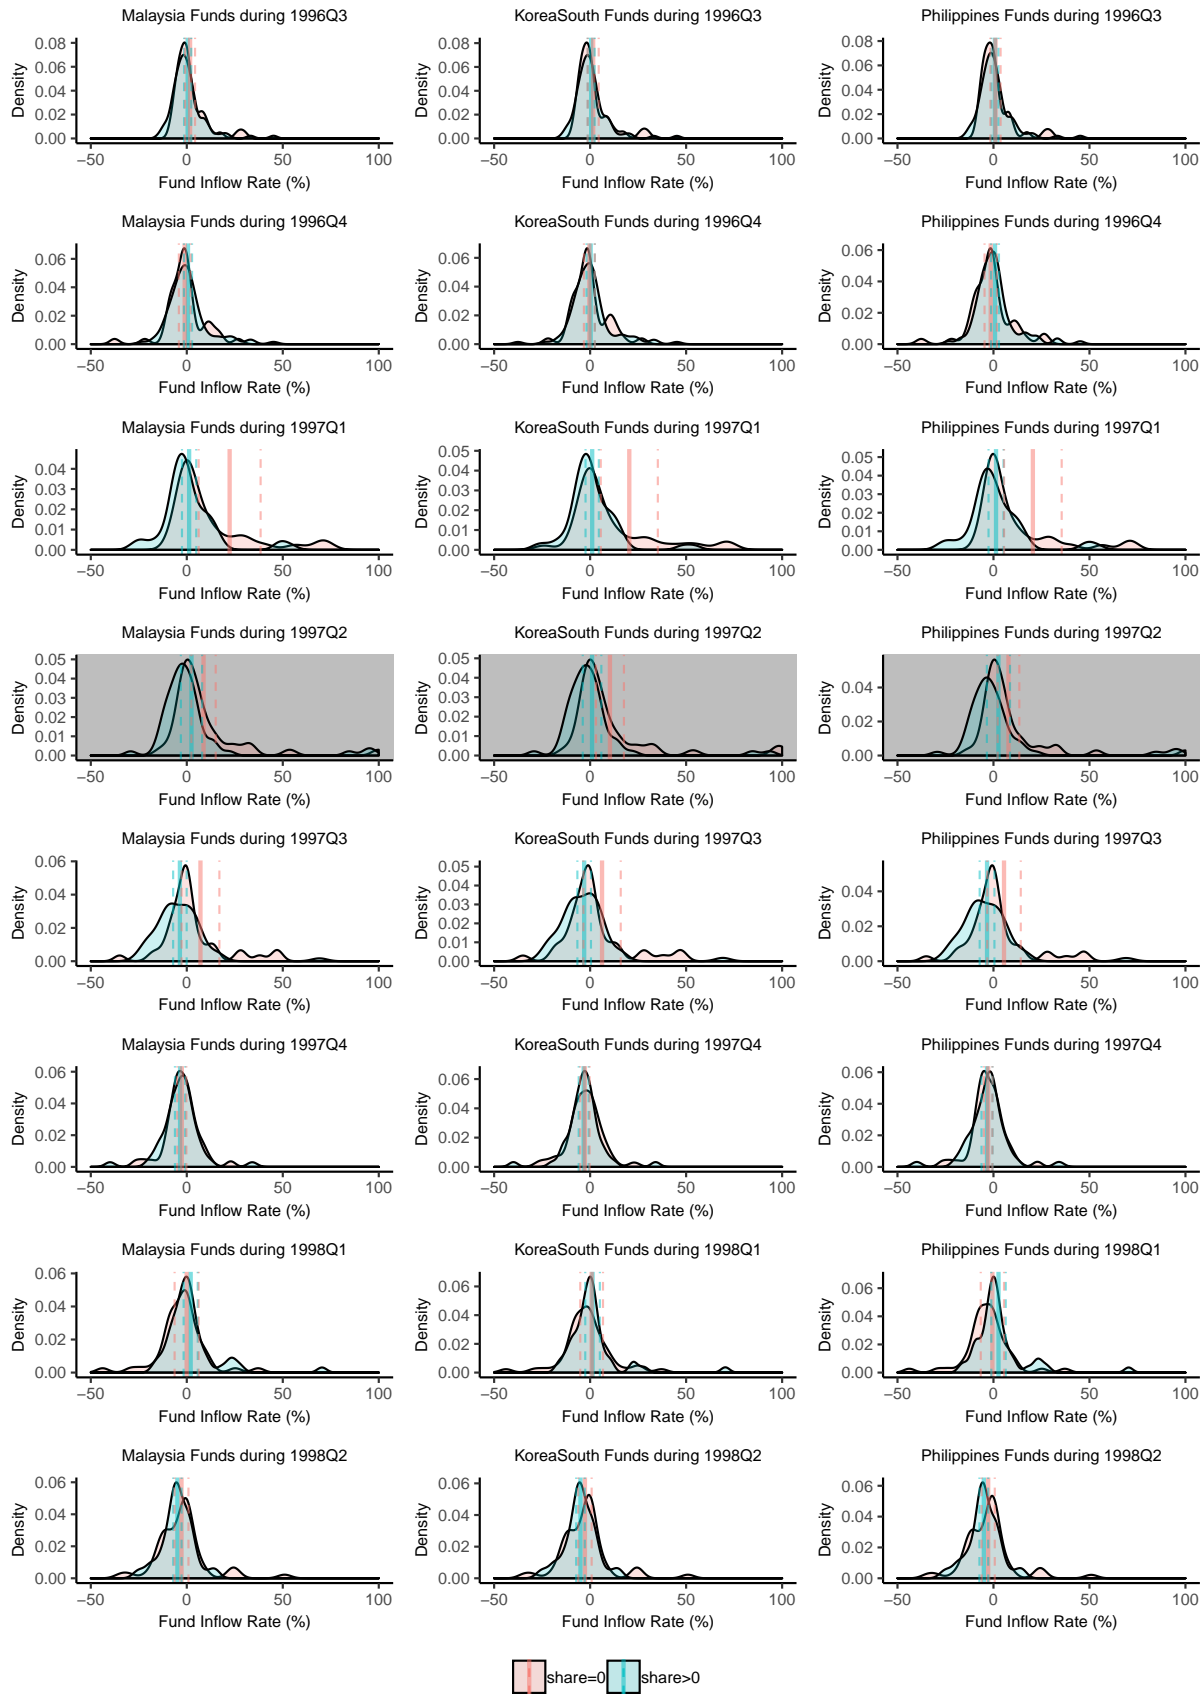


Figure 23: Inflow rate distributions for countries affected by the Asia Crisis (1 of 2)



Notes: Each graph shows the distribution of funds investing in the labelled country during the Asia Crisis (1997 Q2). The panels with grey backgrounds show the inflow rate distributions of mutual funds investing and not investing in the country primarily affected by the crisis. The panels with white backgrounds show the distributions prior to and after the crisis date. The vertical lines show the respective means of each distribution with 95 percent confidence bands. Each row is the same number of quarters before the crisis and each column of columns 1 to 3 is a different country.

Figure 24: Inflow rate distributions for countries affected by the Asia Crisis (2 of 2)

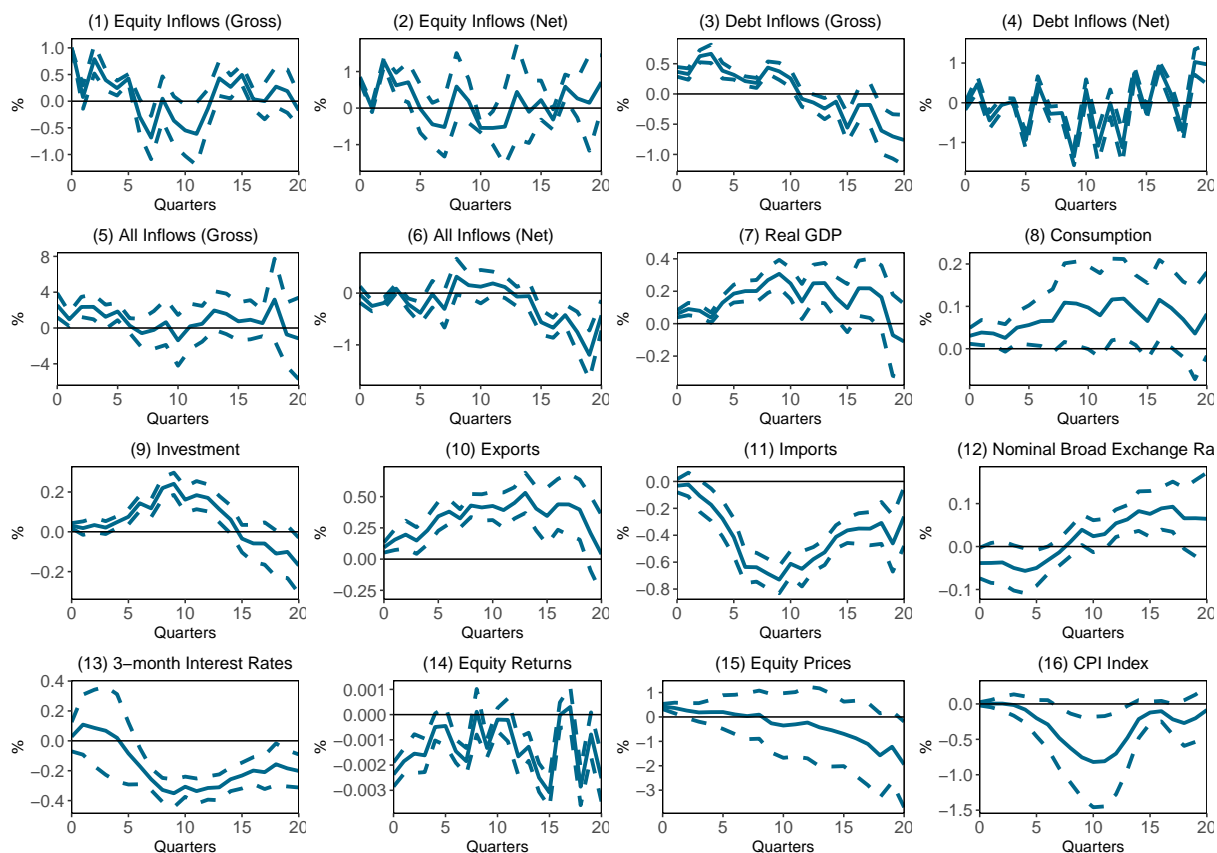


Notes: Each graph shows the distribution of funds investing in the labelled country during the Asia Crisis (1997 Q2). The panels with grey backgrounds show the inflow rate distributions of mutual funds investing and not investing in the country primarily affected by the crisis. The panels with white backgrounds show the distributions prior to and after the crisis date. The vertical lines show the respective means of each distribution with 95 percent confidence bands. Each row is the same number of quarters before the crisis and each column of columns 1 to 3 is a different country.

D Robustness of Second Stage

The purpose of this section is to present a set of robustness checks of the results presented in Section 5. In order to be able to present a large variety of robustness checks, I will limit attention to the robustness of the responses of activity variables. These responses are the more interesting of my findings and the most important.

Figure 25: Pooled impulse responses for all variables

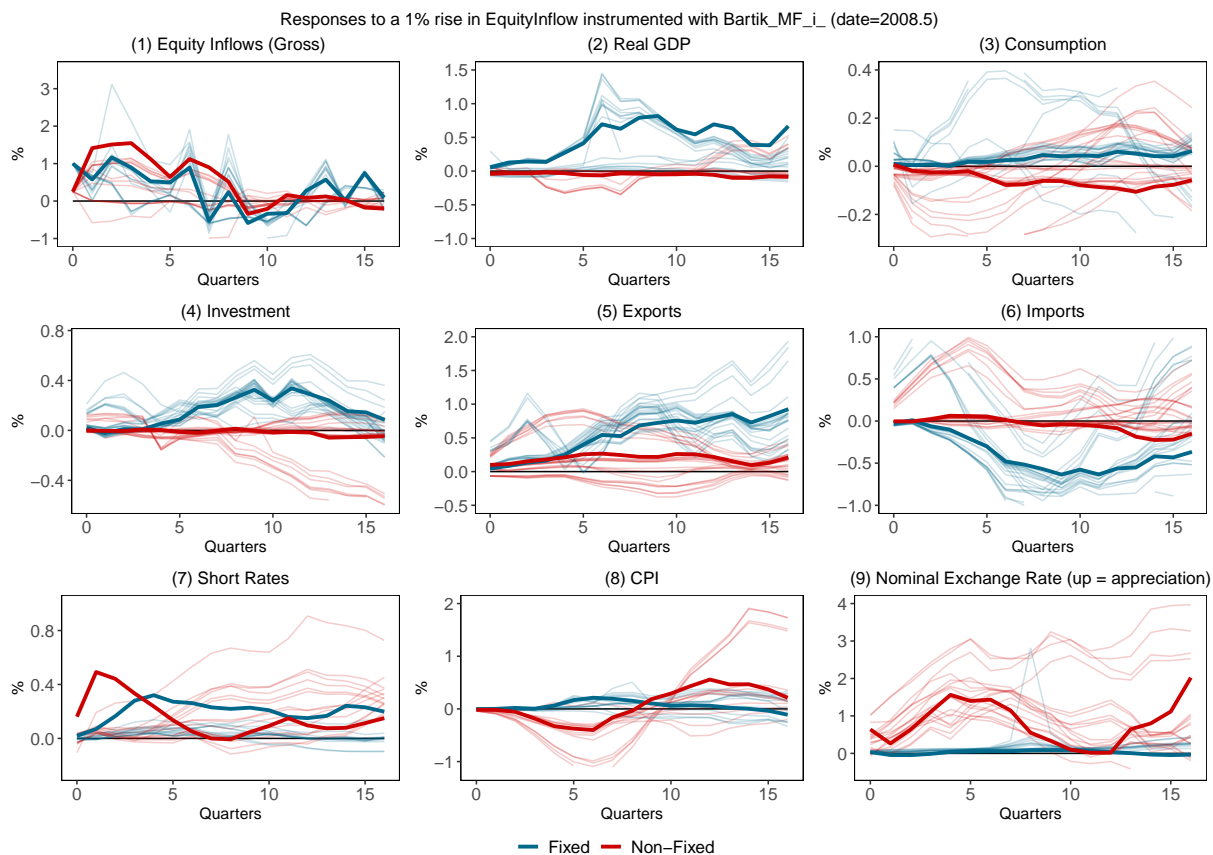


First, Figure 25 shows pooled impulse responses for all variables shown in the paper. Given the heterogeneity between fixing and non-fixing countries, the pooled results are not particularly interpretable but confirm that on average capital inflows tend to be expansionary on output.

Next, Figure 26 shows how different choices of the instrument base period affect the main results. Each grey line in the panels represents an impulse response from a different base period (where I only show base periods where the F-statistic is above 10). While the estimates are somewhat noisy, they tend to be similar qualitatively to the baseline results.

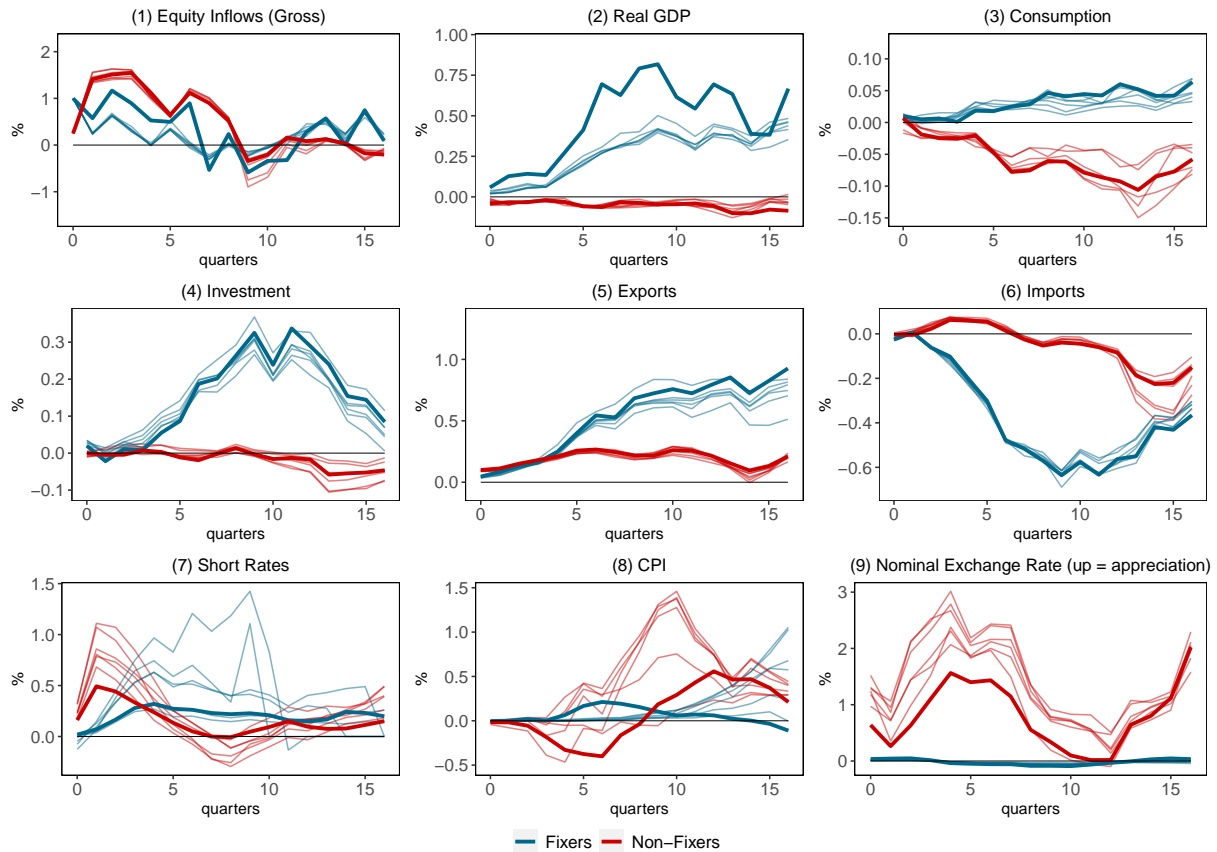
Next, Figure 27 repeats the previous graphs but varies the number of lags in the regressions

Figure 26: Robustness of results by exchange rate regime to different base period choices



from 2 to 8 lags.

Figure 27: Robustness of results to adjusting the lags



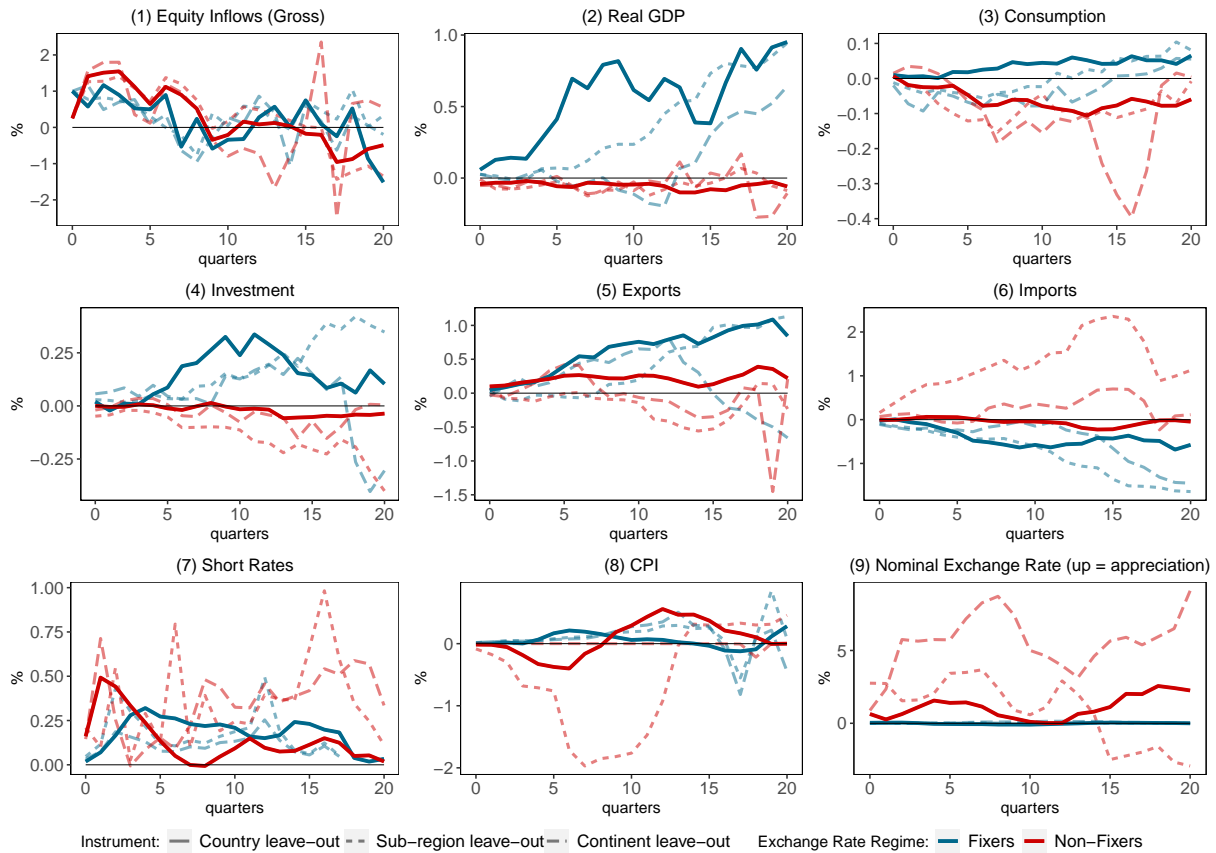
Notes: This figure shows the robustness of the main results to the number of lags chosen. The colors represent the exchange rate regimes of countries in the sample, the bold lines show the baseline impulse responses and lighter lines represent the results from adjusting the number of lags. All impulse responses are calculated with the same base period, 2008 Q3.

D.1 Further Refinements to the Instrument for Equity inflows

As was discussed in Section 2.2, with the instrument refinement the identification condition becomes that lagged interest rates in a country j cannot predict future economic events in a different country i conditional on controls. For countries close together, this might be violated if their economies are sufficiently integrated. For example, Eurozone countries are connected with a common monetary policy, and even controlling for trade connections this additional connection could be missed. Therefore, in this section I present results from a further refinement to the leave-out instrument that excludes all neighboring countries from the subregion each country i is in, and then as a further check I consider predicting inflows into mutual funds investing in country i with lagged returns from countries in a different UN subregion to i and then construct a third refinement where the instrument leaves out all returns earned in the same continent as country i . Readers should note that both refinements mean that the first stage is not robustly above 10 for either refinement and the standard errors

are often very large (and as such are not shown here). Nonetheless, the point estimates of the impulse responses shown in Figure 28 are qualitatively similar to the baseline results (except in the case of the continent-leave-out instrument's GDP responses for fixers).

Figure 28: Robustness to the leaving-out continent-level returns or subcontinent-level returns in the instrument construction



Notes: This figure shows the effects on the impulse responses of leaving out countries near i and countries in the same continent as i when constructing the instrument. The solid lines are the baseline results for fixers (blue) and non-fixers (red), whereas the short dashed lines show the results for fixers (again in blue) and non-fixers (again in red) when the instrument is constructed leaving out not only the return each mutual fund makes in country i but all countries in the same UN subregion as i , and the long dashed lines show the results when the instrument is constructed leaving out any returns mutual funds make in the same continent as i . All impulse responses are calculated with the same base period, 2008 Q3.