

Financial inclusion, economic development, and inequality: Evidence from Brazil

Julia Fonseca^{a,*}, Adrien Matray^{b,c,d}

^a University of Illinois at Urbana-Champaign, Gies College of Business, Champaign, 61820, United States of America

^b Stanford GSB, Stanford, 94305, United States of America

^c National Bureau of Economic Research (NBER), United States of America

^d Centre for Economic Policy Research (CEPR), United Kingdom

ARTICLE INFO

JEL classification:

G20

J24

O10

Keywords:

Financial inclusion

Economic development

Inequality

ABSTRACT

We study a financial inclusion policy targeting Brazilian cities with low bank branch coverage using data on the universe of employees from 2000–2014. The policy leads to bank entry and to similar increases in both deposits and lending. It also fosters entrepreneurship, employment, and wage growth, especially for cities initially in banking deserts. These gains are not shared equally and instead increase with workers' education, implying a substantial increase in wage inequality. The changes in inequality are concentrated in cities where the initial supply of skilled workers is low, indicating that talent scarcity can drive how financial development affects inequality.

1. Introduction

The presence of bank branches at fine geographical levels has long been considered a key determinant of financial inclusion and an important driver of economic growth. Bank branch proximity mitigates the transaction costs of mobilizing savings from many agents, which can increase the capital available to entrepreneurs. Branch proximity also lowers the cost to banks of screening and monitoring entrepreneurs, and of providing access to liquidity services. However, setting up bank branches is not only costly but also risky, as demand for deposits and loans can only be observed after the creation of the branch, which can lead to an under-provision of financial services.

For this reason, policymakers across the world have sought to promote financial inclusion by implementing large-scale reforms to expand the physical networks of bank branches.¹ The popularity of these financial inclusion policies raises multiple questions: do they succeed in promoting financial and economic development and, if so, how? And what are the distributional consequences of such policies?

In this paper, we trace out the dynamic effects on both economic development and wage inequality of a government program that improved access to mainstream financial services. We use the introduction of the “Banks for All” program (“*Banco para Todos*”) by the Brazilian

federal government in 2004, which explicitly targeted underbanked cities by introducing branches of government-owned banks. This policy constitutes a unique natural experiment featuring a large, plausibly exogenous shock to financial access and capital deepening at the level of entire labor markets.

Our empirical analysis combines Brazilian administrative employer-employee data covering the universe of formal employees in Brazil with detailed bank branch balance sheets from 2000–2014. In a difference-in-differences research design, we compare the evolution of various outcomes in cities benefiting from this policy (those with no government-owned banks prior to the reform) relative to unaffected cities. We use a parsimonious matching procedure to select control cities for each treated city, where we match on the pre-reform population quintile and Gini growth, and we estimate the effect of financial development on employment, entrepreneurship, firm growth, average wage, and wage inequality.

Our identification strategy exploits ex-ante differences in the presence of government-owned banks across cities, but it does *not* require the initial presence of government-owned banks to be random. It only requires that outcomes of treated and control cities would have evolved similarly absent the reform. While, by definition, this identifying assumption is untestable, we provide a battery of tests that are supportive

¹ Examples include China in the 1970s, India in the 1980s, Thailand in the 1980s and 1990s.

of it, which we discuss in detail after summarizing our results. Specifically, we: (i) show evidence of pre-reform parallel trends for our key city-level outcomes, (ii) show that our matching procedure leads to covariate balance across a wide array of city-level characteristics not targeted by the matching, (iii) directly control for city characteristics pre-reform, (iv) employ a city-by-industry difference-in-differences estimator to control for sector-specific shocks that could differentially impact cities exposed to the policy. The stability of point estimates across the different strategies implies our results cannot be explained by differential exposure to aggregate or city-specific shocks.

We start by showing that the reform leads to financial inclusion and financial development. After 2004, the number of bank branches increases in treated cities, leading to an inflow of local deposits and an increase in credit supply of similar magnitude. The increases in total bank branches, deposits, and credit do not mean revert but instead shift to a new, higher steady state. This implies better access to external funding and liquidity services, both of which can foster economic development.² Consistent with the policy driving this shift in the steady state of financial development, these increases are driven by government-owned banks, whose expansion only modestly crowds out private banks. The absence of an effect on private credit can be seen as a placebo test, showing that our results are not driven by differential exposure of treated cities to economy-wide shocks experienced by Brazil during this period.

Our second set of results is about the average effect of the reform on economic development. We show that the reform leads to an increase in employment by 10%, mostly driven by an expansion of smaller firms. Increased labor demand pushes up the average wage per worker by 4.1%. Looking at firm dynamics, we find that the reform-induced bank branch expansion fosters entrepreneurship, as the number of firms increases by roughly 10%. This increase masks an even higher acceleration of underlying firm dynamism, as both firm entry and firm exit rate increase. While these results only reflect patterns in the formal sector, entrepreneurial activity and firm growth in the formal sector are of first-order importance for economic development, as the transition from subsistence to transformational entrepreneurs and the integration of workers in the formal sector foster economic growth (e.g., [Dix-Carneiro et al., 2021](#)). Nonetheless, we show in robustness checks that our results are not driven by firms and workers moving from the informal sector to the formal sector.

The richness of our data allows us to examine the mechanisms that link financial inclusion and economic development. Financial inclusion could foster growth by increasing aggregate demand by improving households' access to credit and by allowing households to better smooth their consumption, thereby reducing the need for precautionary saving. We rule out this local demand channel as the main driver of our results by showing that employment growth is mostly driven by firms in the tradable sector, which are by definition less dependent on local demand.

So why would financial inclusion promote entrepreneurship and firm growth? We show that the policy-induced increase in financial inclusion works by reducing the physical distance between banks and entrepreneurs. Indeed, the positive effect of the reform increases proportionally with the distance between treated cities and the closest city with a bank branch prior to the reform and is larger for smaller firms, which likely face bigger informational frictions. These gains are the same whether the nearest bank is private or public, suggesting that our results are not driven by particulars of how public banks operate (e.g., by access to subsidized or politically-motivated loans). Instead, these results are consistent with the distance to bank branches affecting access to liquidity, and with models in which the distance between borrowers and lenders affects the cost of credit either because it reduces

screening and monitoring costs (e.g., [Greenwood and Jovanovic, 1990](#); [Townsend and Ueda, 2006](#)), in particular when soft information is prevalent (e.g., [Petersen and Rajan, 1994](#); [Hombert and Matray, 2017](#)).

Our third set of results is about the distributional effects of the reform. We find that the policy leads to a sizable increase in wage inequality within treated cities. This is explained by the fact that, although all workers are better off after the reform, the magnitude of wage gains rises monotonically with the position of workers in the wage distribution. Our detailed panel data of workers allow us to show that this increase is not driven by a change in the sample composition, but instead reflects an increase in wages holding fixed individuals' sex, age, education, occupation, and sectoral specialization. We also show that our results are quantitatively unchanged when we restrict our sample to workers that we observe throughout the sample period and to firms already in the data prior to the reform. These sample restrictions allow us to show that our results are not driven by workers entering the formal sector after the reform or more general changes in the sample composition of worker characteristics.

We then explore two explanations that can account for the rise in inequality. First, financial development could increase the relative demand for skilled labor, either because of a large fixed component to the cost of skilled labor (e.g., [Benmelech et al., 2021](#); [Schoefer, 2021](#)) or because the relative productivity of skilled workers increases with financial development ([Fonseca and Doornik, 2022](#)). Models that assume that financial development increases the relative productivity of skilled workers or loosens constraints on the demand for skilled workers generally predict that the equilibrium skill mix should change, with firms increasing the share of skilled workers in their workforce. However, when looking at the effect of the policy on the average skill composition of firms, we find that the share of skilled workers does not increase in treated cities.

Instead, we find support for another explanation: skills are scarce, especially in developing countries, which means that the supply of skilled workers is more inelastic than that of unskilled workers in the short run. We show that cities in our setting are characterized by high internal migration costs and that the reform does not induce worker migration to treated cities. This lack of inter-city mobility implies that an increase in labor demand can only be served by the supply of local workers. Consistent with skilled workers being in short supply, we find that all the increase in inequality is concentrated in cities where a lower fraction of the population is educated prior to the reform.

We consider a wide range of robustness checks. We start by showing that our results are quantitatively unchanged when we use different matching procedures. We then discuss threats to identification. Our strategy faces two key threats. First, even in the absence of pre-trends, treated cities may be ex-ante different in ways that differentially expose them to aggregate shocks post-2004. That would be the case if, for instance, treated cities are ex-ante more exposed to the commodity boom of the mid-2000s. Second, our policy might have coincided with shocks that specifically affected treated cities, such as idiosyncratic shocks to banks entering treated cities or a targeted expansion of welfare programs.

We address the threat produced by ex-ante differences in three ways. (i) We show that our matched treatment and control groups are similar over a rich array of city characteristics that were not included in the matching process, including exposure to the commodity sector, skilled employment, political affiliation, size of the informal sector, or the co-movement of local GDP with aggregate fluctuations. While common support in *levels* is not required for differences-in-differences designs, such similarity makes the common-trend assumption more plausible, as these similarities in the level of characteristics make it less likely that they reacted differently to broader macroeconomic shocks post-2004.

In addition, (ii) we show that our results are quantitatively unchanged after directly controlling for a wide range of pre-reform controls interacted with year fixed effects. Estimating all possible combinations of pre-reform controls across the hundreds of different specifications yields very similar point estimates. Finally (iii), we exploit

² See [Xu and Yang \(2022\)](#) for an example of how access to liquidity services thanks to bank branch expansion promoted development in 19th-century US.

the granularity of our data to build a city-by-industry difference-in-differences estimator. This allows us to include industry-by-year fixed effects and non-parametrically control for any unobserved time-varying sector-specific shocks (e.g., commodity booms or trade shocks). Our coefficients of interest are estimated in this case by comparing the *same* sector across treated and control cities, and therefore this strategy does not require that treated and control cities are similarly exposed to sector-specific shocks. We show that point estimates at the city-industry level are quantitatively similar to city-level estimates.

Our setting also addresses a wide array of potential ex-post treated-specific shocks because, by construction, control cities already have a government-owned bank. Therefore, any shocks specific to government-owned banks (such as an overall increase in lending by public banks) will affect both treated *and* control cities at the same time and will be absorbed by our difference-in-differences specification. This design also addresses the possibility that welfare program expansions might differentially benefit treated cities, since some of the largest welfare programs are distributed by one of the public banks already present in control cities (including the largest at the time, *Bolsa Família*). Therefore, an expansion of these programs would also affect both treated and control cities. In additional robustness tests, we show that results are robust to including state-by-year, which controls for state-level shocks such as differences in state-administered welfare programs. At the municipality level, our effects are unchanged when we compare treated and control cities with the same political affiliation or directly control for the observed changes in expenditures.

We end the paper by discussing how our reduced-form identified coefficients can provide useful causal moments for the macro-finance development literature and speak to potentially important frictions or sources of heterogeneity that future models could incorporate. Our paper shows the importance of explicitly linking distance to the nearest bank to the cost and availability of credit as in the structural model of [Ji et al. \(2023\)](#). We provide causal estimates of how changes in distance can affect credit supply and savings in interest-bearing products, as well as their impact on employment, firm growth, and firm entry. We also provide moments linking changes in the supply of credit and real outcomes, which relate to key parameters in macro-development models in which a reduced-form collateral constraint affects economic growth. The considerably larger effects we find for treated cities in banking deserts point toward the existence of a non-linearity around very low levels of external finance, something that is usually not explicitly modeled and could help to reconcile different results in the literature. Finally, our paper highlights the importance of worker heterogeneity and constraints on the supply of human capital in accounting for the dynamics of wage inequality as a result of financial development in macro-finance models.

Literature. Our paper contributes to several strands of literature. The closest one is on the role of financial inclusion on financial intermediation costs and economic development, which plays a central role in macro-development models, but for which we have limited causally estimates elasticities that can be tightly linked with model parameters.³

Our setting allows us to better understand why financial inclusion fosters economic growth and provide new insights on the distributional effects of such policies. First, our high-quality administrative data allow us to study many new margins through which financial inclusion policies can foster development. Because we can observe the universe of formal economic activity and track workers and firms over time, we can study a host of new outcomes: the process of creative destruction (firm entry and exit rate), the evolution of firm size distribution and the differential benefit for small vs. large firms measured pre-reform, as well as the evolution of sectoral composition (and in particular

the relative growth of tradable vs. non-tradable sectors). Because we observe the universe of workers and can track them over time, we can also shed light on the evolution of the skill premium and wage inequality, accounting for changes in sample composition, and the interaction between financial inclusion policies and constraints on the accumulation of human capital.

The literature on the effect of bank branch presence has so far mostly focused on developed economies and studied financial outcomes such as credit or wealth accumulation. The literature on developing countries has mostly focused on the introduction of specific bank branches in localized markets and used short-run cross-sectional survey data to study real outcomes and find small, short-lived positive effects or even negative effects.⁴ By contrast, our administrative panel data allow us to track the long-run effect of the policy on the *full causal chain* linking financial inclusion to financial development (credit supply, access to liquidity services); real outcomes (employment, sectoral growth, average wages, entry and exit rates of firms, firm size distribution) and the distributional consequences of this financial-inclusion driven economic growth.

Second, the policy-driven expansion of bank branches happens through publicly owned commercial banks that, as a whole, are profitable rather than development-focused state banks with specific sectoral and poverty reduction mandates as in the Indian experiment ([Burgess and Pande, 2005](#); [Cole, 2009](#)). The policy also induced large variation at a sufficiently large geographical level to estimate “local general equilibrium effects”, but sufficiently small to provide precise estimates.⁵ This level of analysis coupled with data on the universe of bank branch networks in Brazil allows us to provide the first causal estimate of the role of distance to financial services, a key parameter in macro-development models that analyze the role of banks, and to highlight that financial inclusion is not a dichotomous concept but should instead be thought of as continuous.

Taken together, our results have several policy implications. First, financial inclusion policies can generate large effects on economic development, even while banks operate profitable branches. Second, these policies can have large distributive effects, in particular as they interact with constraints on human capital accumulation. Third, the initial distance to existing bank branches matters for the expected gains in terms of economic development. Therefore, financial inclusion policies can reap larger benefits at lower costs by carefully taking into account the initial network of bank branches, which is important given the large number of policies around the world attempting to foster economic development by promoting financial inclusion.

This paper also contributes to the empirical literature using natural experiments to show how financial frictions, broadly defined, affect economic development.⁶ Most of the evidence for developing countries

⁴ See [Bruhn and Love \(2014\)](#); [Burgess and Pande \(2005\)](#); ([Cramer, 2022](#)) for positive effects and [Kochar \(2011\)](#) for negative effects. A complementary approach exploits randomized controlled trials to study the implications of access to microcredit and savings products in developing countries. The literature on microcredit is surveyed in [Banerjee et al. \(2015\)](#), which concludes that microcredit has “modestly positive, but not transformative, effects”.

⁵ By contrast ([Bruhn and Love, 2014](#)) studies the opening of outlets of a bank specialized for low-income in supermarkets, while ([Burgess and Pande, 2005](#)) uses state-level variations across Indian states. More generally, most of the literature on developing countries has used randomized controlled trial and shocks to specific banks with a focus on directly affected bank clients. The positive effects on non-clients are potentially a key driver of multiplier effects, which can account for why we find large positive effects on economic development while most papers find limited effects.

⁶ An earlier literature looks at how financial frictions relate to economic development using cross-country evidence. This literature is for instance reviewed in [Anon \(2018\)](#). See [Buera et al. \(2011\)](#), [Buera and Shin \(2013\)](#), or ([Midrigan and Xu, 2014](#)) for macro-models linking financial frictions and development and the survey in [Buera et al. \(2015\)](#). See also ([Xu, 2022](#)) and [Xu and Yang \(2022\)](#) and references therein for the importance of financial frictions in cross-country trade and growth, and long-run historical contexts.

³ See for instance ([Greenwood and Jovanovic, 1990](#)); [Greenwood et al. \(2010\)](#); [Ji et al. \(2023\)](#), where geographical distance to a bank branch governs financial intermediation costs and the returns to saving.

studies short-run capital injections that originate outside the city and focuses on changes in credit, *holding fixed the network of banks*.⁷ Our paper highlights a fundamentally different mechanism and our object of interest is the promotion of financial inclusion by *expanding the network of bank branches*. This branch expansion fosters the mobilization and pooling of local savings to start a virtuous circle between increased deposits, higher credit, and economic development, originating from within the city as the distance between depositors, lenders, and borrowers is reduced. Therefore, our results imply that *how* credit is distributed across places can matter as much as *how much* credit is distributed.

Third, we contribute to the literature that studies the effect of financial development on wage inequality. Theoretical work in this literature focuses mostly on wealth inequality or total income inequality (which includes capital income) and derives ambiguous predictions. The effect of financial development depends on whether that development is concentrated on the intensive or the extensive margin (Greenwood and Jovanovic (e.g., 1990), Townsend and Ueda (e.g., 2006), Greenwood et al. (e.g., 2010), how it alters the aggregate demand of workers and investment returns (e.g., Giné and Townsend, 2004; Buera et al., 2021), and whether individuals can accumulate human capital (e.g., Mestieri et al., 2017). These models generally conclude that capital income pushes inequality upward, as it mostly benefits the wealthy and entrepreneurs, while rising wages push inequality downward (e.g., Buera et al., 2021; Ji et al., 2023). These theoretical and quantitative results that wage inequality should go down as financial development increases rely on the assumption that labor is a *homogeneous* input to production. Therefore, higher labor demand in more-productive sectors will benefit more lower-paid workers who reallocate away from less-productive sectors.

Our contribution to this literature is twofold. First, we provide rare empirical evidence on the effect of financial inclusion on wage inequality, as empirical evidence focuses on developed countries and studies credit rather than financial inclusion.⁸ Second, we show that financial inclusion leads to higher wage inequality in our setting due to skill differentials. Therefore, we show that taking into account labor heterogeneity and limits to human capital accumulation in macro-development models is crucial to better understanding and predicting how policies promoting financial development will affect inequality.

Fourth, this paper contributes to our understanding of how financial frictions impact capital and entrepreneurial talent misallocation and thereby economic development.⁹ More broadly, we relate to the literature studying how financial frictions affect firm labor demand and employment outcomes.¹⁰ We contribute to the specific subset of the literature that studies how financial frictions affect the demand for skilled workers and the skill premium in developing countries (Fonseca and Doornik, 2022).¹¹

⁷ For instance shocks to the liquidity of lenders coming from targeted lending programs in India (Banerjee and Duflo, 2014) and Brazil (Bazzi et al., 2023), deposit volatility (Choudhary and Limodio, 2022), large government grants in Thai villages (e.g., Kaboski and Townsend, 2011, 2012), or broader financial market reforms such as bankruptcy reforms (e.g., Fonseca and Doornik, 2022, for Brazil), collateral laws (e.g., Vig, 2013), or financial liberalization (e.g., Crescenzi and Limodio, 2022; Bau and Matray, 2023, and references therein).

⁸ See Beck et al. (2010) for evidence from the U.S., and a discussion of their results once accounting for staggered D-i-D corrections by Baker et al. (2022).

⁹ See, among many others: (Giné and Townsend, 2004); Townsend and Ueda (2006); Banerjee and Moll (2010); Buera et al. (2011); Kaboski and Townsend (2011); Buera and Shin (2013); Midrigan and Xu (2014); Moll et al. (2017); Bau and Matray (2023).

¹⁰ See among many others: (Peek and Rosengren, 2000); Chodorow-Reich (2014); Hombert and Matray (2017); Bai et al. (2018); Berton et al. (2018); Caggese et al. (2019); Greenstone et al. (2020); Baghai et al. (2021); Doornik et al. (2021).

¹¹ For recent works on financial frictions and the demand for skills in developed countries, see Quincy (2023) and Jasova et al. (2021).

Finally, because the reform we explore relies on the expansion of government-owned banks, we relate to the broad literature studying the economic effects of government ownership of banks (e.g., Sapienza, 2004; Dinç, 2005; Cole, 2009; Carvalho, 2014; Delatte et al., 2022). Most of this literature emphasizes the risk of political capture and the creation of politically motivated credit cycles. We show that such forms of ownership can have positive effects on economic development when the private sector is unable or unwilling to serve underprivileged areas, even in countries where corruption can be high.

2. Institutional background and data

2.1. The Brazilian banking landscape

2.1.1. Situation pre-reform

Brazil has three types of public banks: government-owned banks controlled by the federal government (Banco do Brasil, Caixa Econômica Federal, Banco do Nordeste, and Banco da Amazônia), government-owned banks controlled by state governments, and a national development bank (BNDES).

Government-owned banks, in particular the ones controlled by the federal government, differ from most public banks in developing countries and are better described as “government-owned commercial banks”. They are profitable and their performance is comparable to that of both foreign and domestic private banks (Mettenheim, 2010).

BNDES, the national development bank, differs substantially from federal and state-owned banks and is much closer to public banks in other developing countries that have been studied previously such as India (e.g., Burgess and Pande, 2005). This bank provides subsidized loans to targeted sectors or even “grants”, in the form of loans that are often not reimbursed. Two things are important to highlight about BNDES. First, this is the only bank studied by previous papers that have documented the existence of political influence in banks’ behavior in Brazil (e.g., Carvalho, 2014). Second, this bank, as well as state-controlled banks, was not part of the policy and not involved in the branch expansion we study, and its lending was not specifically targeted at treated or control cities. Therefore, in the rest of the paper, we use the term “public banks” to refer to federally-owned banks, which excludes BNDES.

Public banks differ from private banks in some dimensions, such as in having a legal mandate to provide earmarked credit.¹² However in practice, public banks are similar to private banks in their lending practices and the sectors to which they lend. For instance, public and private banks have similar portfolio compositions across credit products and borrowers, charge similar interest rates and face similar delinquency rates (Coelho et al., 2013). These banks are also similar along standard balance sheet and income statement measures. As we show in Fig. 2, there are no significant differences between branches of public and private banks across a wide range of covariates, including measures of lending, profitability, loan performance, and size.

2.1.2. The banks for all program

Banks for All (*Banco para Todos*) was a federal government program announced in 2004 as part of the government’s 2004–2007 multi-year plan (*Plano Plurianual*). The program was under the purview of the Finance Ministry (*Ministério da Fazenda*) and aimed to provide Brazil’s unbanked population with access to financial services and products through the actions of federal government banks, particularly Caixa Econômica Federal and Banco do Brasil.

To achieve the goal of reaching underserved communities, the federal government promoted the physical presence of public banks

¹² There are also incentives for private banks to provide earmarked credit and, in fact, nearly 40% of outstanding indirect earmarked loans to firms in 2016 were originated by private banks (Ornelas et al., 2021).




Fig. 1. Share of municipalities without bank branches.

This figure plots the evolution of the share of municipalities without at least one government-owned bank branch in red and the share of municipalities without any bank branches in blue.




Fig. 2. Public vs. Private bank covariate balance.

This figure shows coefficient estimates and 95% error bands of the difference between public and private banks along different variables, using ESTBAN data at the branch-year level between 2000 and 2004. All variables are normalized to have a mean of zero and a standard deviation of one. Standard errors are clustered at the bank level.

throughout the country, focusing on cities with no presence of government banks. Fig. 1 plots the evolution of municipalities without a public bank branch since 2000 (the dashed red line). Consistent with the effect of the reform, this share is stable until 2004 at 60%, then drops abruptly in 2005 and keeps declining such that in 2014, only 44% of municipalities have no government-owned banks. Fig. 1 also reports the share of municipalities without any bank branch (the solid blue line), and shows that the expansion of public banks resulted in a drop in the share of cities without any bank branches.¹³

¹³ Between 2004 and 2014 (the end of our sample period), 1,262 new bank branches were open in cities eligible to the reform.

The program succeeded in reaching unbanked cities and under-banked populations. According to an evaluation of the program by the federal government, public banks opened 7.8 million accounts and banked 1.46 million low-income, previously unbanked individuals between 2004–2007 (Ministério da Fazenda, 2007).¹⁴ In Section 4, we formally show that cities without public bank branches prior to 2004 saw a sharp increase in credit and deposits following the introduction of the program.

¹⁴ For comparison, in 2007, there were approximately 16 million individuals residing in the cities that compose our treatment group.

Banking correspondents. In order to reach unbanked households, the program also relied on correspondent banking outlets. These arrangements consist of banks hiring commercial entities—typically lottery retailers, post offices, pharmacies, and other retailers—to serve as distribution outlets for financial services. Since 2003, financial services offered by correspondents include the opening of accounts, deposits and withdrawals, payments, and loan applications.¹⁵ The number of correspondents went from approximately 20,000 in 2000 to over 150,000 in 2010 (Loureiro et al., 2016) and, taking into account partnerships with correspondents, government-owned banks were present in 100% of municipalities by 2007 (Ministério da Fazenda, 2007).

While we do not observe the precise location of correspondents, we know that, prior to the reform, they were concentrated in areas already served by bank branches (Loureiro et al., 2016). One notable network of correspondents was that of Banco Postal, which emerged from a 2001 partnership between Bradesco and the Brazilian Post and Telegraph Company to provide financial services at post offices. However, despite its goal of reaching all unbanked municipalities, it was only present in roughly half of unbanked municipalities prior to the reform (Gual and Ansón, 2008). Moreover, prior to 2003, correspondents were not allowed to provide most of the services they offer today. Thus, to the extent that they were available, they mostly provided bill payment services and were not a meaningful substitute for financial institutions and were particularly ill-suited for the needs of firms (Bittencourt et al., 2005). Even after 2003, the vast majority of Brazilian households and an even larger fraction of Brazilian businesses did not use correspondent banking outlets for deposits, withdrawals, account openings, or borrowing, and did not view correspondents as safe or trustworthy according to survey evidence (Sanford, 2013).

Taken together, these findings suggest that banking correspondents were present in many unbanked municipalities prior to the reform but mostly provided bill payment services. While having some access to financial services is certainly better than no access at all, the fact that these services were limited implies that the expansion of branch networks we study led to financial inclusion by improving the quality and scope of financial products and services to which households and businesses have access.

2.2. Data

We use data from four distinct sources. Matched employer-employee data come from the *Relação Anual de Informações* (RAIS), a mandatory annual survey containing information on the universe of tax-registered firms in Brazil. There are severe penalties associated with incomplete or late information, which leads to a high degree of compliance and essentially complete coverage of all employees in the formal sector. RAIS contains time-invariant identifiers for workers and firms, as well as information on where the firm is located. We also observe data on workers average gross monthly earnings, occupation and several socio-demographic characteristics such as their education, race, age, and gender.

Using geographical information on firms, we build a city-level panel from 2000 to 2014 with information on average wages, wage inequality, employment, and skill-specific wages. Because municipality borders have changed over time, we use as our level of aggregation minimum comparable areas (*Área Mínima Comparável*, or AMC), which can be consistently tracked throughout our sample period. This reduces the number of cities from over 5,000 to 4,260. In the rest of the text, we use the term “city” to refer to an AMC.

The number of bank branches, lending activity, and deposits come from the ESTBAN database maintained by the Central Bank of Brazil. The data provides branch-level balance sheet information that we aggregate to the city level, which allows us to decompose the number of

branches, credit, and deposits between public and private banks. Note that these data do not include correspondent banking outlets (such as the outlets of Banco Postal), which means that we do not observe the full impact of the program on financial inclusion. We discuss this point further in Section 4.

Finally, we use city-level aggregate data. We obtain time-varying outcomes from the Brazilian Institute of Geography and Statistics (*Instituto Brasileiro de Geografia e Estatística*, or IBGE), and cross-sectional demographic and economic characteristics in 2000 from the Census, such as population distribution across years of schooling and share of workers in informality.

3. Empirical strategy

The reform promoted financial inclusion by targeting cities with no government-owned banks, so we identify treated cities as those that did not have a public bank prior to 2004. This implies that all control cities had a public bank prior to the reform.¹⁶ We can identify the effect of the financial inclusion reform by comparing the evolution of multiple economic outcomes for treated and control cities, before and after the reform, in a difference-in-differences setting. The key identifying assumption is that absent the reform, treated and control cities would have evolved in close parallel. While this assumption is untestable, we discuss in this section the conditions under which it is plausibly satisfied.

This strategy raises two natural challenges. First, even if the reform can be considered as “quasi exogenous” from the point of view of each city, there is no reason to believe that the decision of a public bank to set up a branch in a specific city within the pool of eligible ones, and the timing of this decision, is random. Instead, it could be the result of systematic differences in behavior or characteristics among cities eligible for the reform.

This situation would bias our estimates (mostly likely upward) since the eligible cities that do receive a bank branch would be the ones where having a branch is the most profitable. Such a bias would emerge if we were to consider that only the eligible cities that do receive a bank branch are treated and use the actual date of the opening to define the treatment, because the latter would imply defining treated and control units based on the *ex-post realization* of the policy, rather than an *ex-ante assignment* rule.¹⁷

Our approach to address this problem is to instead consider that all cities without a public bank branch are treated immediately after 2004, whether or not they end up with a branch post-reform. In this case, it is possible to recover an unbiased estimate of the true effect of having potential access to a bank branch by focusing on the assignment to the treatment, rather than its actual take-up, as long as we can assume that the policy (the assignment) is exogenous based on the *ex-ante* characteristic of not having a government bank prior to 2004. In other words, the identifying assumption is not whether cities in which a public bank branch enters would have trended similarly than other cities absent of the reform, but is whether all the cities that become eligible to receive a bank branch would have trended similarly.

The validity of our research design therefore depends on whether we believe that cities that did not have a public bank branch prior to 2004 would have trended similarly as cities that did have a public bank branch prior to 2004. This raises the second challenge: the location of banks prior to 2004 is not random and, consequently, the average treated city in Brazil does not look like the average untreated city. Note that similarity in levels between treated and control units across covariates prior to a shock is not a necessary condition for identification

¹⁵ CMN Resolution 3,110 of July 31, 2003.

¹⁶ Cities with no public bank prior to 2004 represent 43% of Brazilian cities.

¹⁷ In Appendix A.3, we show that defining treatment as the actual entry of a public bank branch does lead to estimates that are overall larger than our baseline.

in a difference-in-differences research design, as level differences will be absorbed by the unit fixed effects. However, similarity in levels, though not a necessary condition, makes the parallel-trends assumption more plausible.

Since the reform targeted unbanked cities, these tended to be smaller and less developed, and it is possible that they evolved and grew in different ways after the reform relative to other untreated cities for reasons not directly tied to the reform. For instance, they could have disproportionately benefited from the period of sustained growth, partially fueled by a commodity boom, that Brazil entered into during our sample period. Fig. 3 plots a covariate balance test and shows that the unconditional difference in levels between treated and untreated cities (green coefficients) is large and significant for most city characteristics.

In order to address this challenge and to strengthen our empirical strategy, we use a parsimonious matching approach to construct a control group of untreated cities that is observably similar to treated cities on a wide set of characteristics.

Matching. Our matching strategy first targets city size. We start with all 4,260 cities and compute quintiles of population. We then match each treated city with all control cities in the same population quintile with replacement.¹⁸ This parsimonious approach addresses a large part of the heterogeneity. The red dots in Fig. 3 show that the treatment and control groups are now similar over a rich array of city characteristics constructed pre-reform that were not targeted in the matching process. These characteristics include proxies for economic development (GDP, employment, skilled employment, size of informal sector); propensity to receive social transfers (local government expenditures, political affiliation of the mayor); economic integration (distance to the state capital, share of population born elsewhere, exports and imports made by local firms); exposure to aggregate shocks and the commodity sector (local GDP co-movement with aggregate GDP, total employment in commodities, exposure to commodity prices post-reform); and development of the private banking sector (private loans and private deposits). In addition, while some of the point estimates are not exactly zero, the standardized difference between both groups remains well below the threshold of 0.20 suggested by Imbens and Rubin (2015).

After matching on population quintile, the only remaining large and statistically significant difference between treated and control cities is the change in the Gini index during the pre-period. Since we are interested in understanding how financial development affects inequality, and because Brazil experienced large changes in inequality during this period (e.g., Lopez and Perry, 2008) we also match on changes in inequality pre-reform. We do so by selecting the three control cities in the same population quintile with the closest pre-reform Gini growth. The blue dots in Fig. 3 show differences between treated and control cities after we further restrict our matches to this criterion. The difference in Gini growth becomes much closer to the 0.2 threshold and later in the paper, we show that the Gini index of treatment and control units evolved in close parallel prior to the reform, and that there is no evidence of pre-trends (Fig. 8).

After our baseline matching procedure, we are left with 1,415 treated cities and a total of 3,918 control cities. We report the summary statistics of our final sample in Table 1, and we display the spatial distribution of treated and control cities in Fig. 4. Treated and control cities are spread out across Brazil and do not show geographical clustering.

Econometric specification: city level. We analyze the effect of an increase in bank coverage on economic development and inequality by estimating a series of matched difference-in-differences (D-i-D) specifications of the form:

$$Y_{c,g,t} = \beta \text{Treated}_c \times \text{Post}_{t \geq 2004} + X_{c,t} + \theta_c + \delta_{g,t} + \epsilon_{g,c,t} \quad (1)$$

¹⁸ Because the same city can be used multiple times as a control, we follow (De Chaisemartin and Ramirez-Cuellar, 2024) and adjust the standard errors by clustering at the matching pair level

Table 1
Summary statistics.

	Mean	Med.	St. Dev.	N
Loans/GDP	0.18	0.14	0.17	79,995
Public Loans/GDP	0.16	0.11	0.16	79,995
Private Loans/GDP	0.02	0.01	0.04	79,995
Total branches	1.77	1.00	4.60	79,995
Public branches	0.93	1.00	1.29	79,995
Private branches	0.84	1.00	3.43	79,995
Deposits/GDP	0.13	0.10	0.11	79,995
Public deposits/GDP	0.09	0.07	0.09	79,995
Private deposits/GDP	0.04	0.01	0.06	79,995
Wage	926.41	893.98	273.93	79,995
Total employment	1,056.30	620.00	5,653.74	79,995
Share skilled	0.09	0.09	0.05	79,995
Skill premium	2.28	2.14	0.69	79,901
Gini index	0.31	0.31	0.06	79,995
Population	12,347.19	8,635.00	24,640.18	79,995
GDP per capita	13,478.91	9,630.87	16,685.64	79,995
Share manufacturing	0.20	0.12	0.20	79,995
Share agriculture	0.14	0.09	0.14	79,995

This table reports summary statistics of average city-level characteristics our final sample. Monetary values are in 2010 BRL. Number of bank branches, lending activity and deposits are from the ESTBAN database. Wage, employment, and other labor market variables are from the RAIS database. Local GDP per capita, population, and the share of manufacturing and agriculture in local value added are from the Brazilian Institute of Geography and Statistics.

where $Y_{c,g,t}$ are various city outcomes for city c at year t that belongs to a matched treated-control group g , and Treated_c is a dummy variable that takes the value one if city c had no government-owned banks prior to 2004. θ_c are city fixed effects that remove time-invariant heterogeneity across cities, and $\delta_{g,t}$ are matched group-by-year fixed effects that controls for time-varying unobserved heterogeneity across groups. Because we select our groups using pre-reform population size and inequality growth, the inclusion of matched group-by-year fixed effect implies that we are absorbing unobserved correlated shocks that might exist between these characteristics and the reform.

For example, concerns that smaller cities may have grown for reasons unrelated to the reform will be addressed because the parameter of interest β is identified solely by comparing cities within the *same* group, i.e., within the same size quintile. Similarly, unobserved shocks to places with larger changes in their Gini prior to 2004 will also be differenced out by these fixed effects. $X_{c,t}$ is a collection of city-level controls that we include in the robustness analysis.¹⁹ We cluster our standard errors at the matching-pair level to account for serial correlation and weight the regression by population size at the beginning of the period to estimate the aggregate effect of the reform on inequality and economic development.

Identifying assumptions and potential threats to identification.

Our identification strategy faces two main threats: (i) Even if treated and control cities are perfectly similar ex-ante, unobserved ex-post shocks might specifically affect the cities that are treated by our financial inclusion policy. (ii) Despite the use of a matching procedure, the variable we use to sort cities into treatment and control groups – the presence of a government-controlled bank – might still be correlated with other city-level characteristics that make treated units more sensitive to aggregate shocks post 2004. We discuss how we address both of these concerns below.

(i) Treated-specific ex-post shocks. Even with perfect ex-ante covariate balance between treated and control cities, the estimated effect of promoting financial inclusion on city-level outcomes could be biased if this policy correlates with other unobserved shocks that specifically affect cities that received the treatment. This is a concern in a setting where no city has a bank and identification is achieved by bank entry

¹⁹ We use the value pre-reform and interact with year fixed effects to avoid the classic problem of “bad controls”.



Fig. 3. Covariate balance.

This figure shows coefficient estimates and 95% error bands of the difference between treated and control cities along different variables. All variables are normalized to have a mean of zero and a standard deviation of one in the full sample. “Unconditional” refers to the sample where we compare treated cities to all untreated cities. “Population” refers to the sample where we match treated cities with untreated cities in the same population quintile pre-reform. “Population + Gini” refers to the sample where we select the three control cities in the same population quintile with the closest pre-reform Gini growth.




Fig. 4. Geographical distribution of treated and control cities.

This figure shows the geographical distribution of treated and control cities. “Not in sample” refers to cities that are neither treated nor part of the matched control group, and thus not in our final sample.

in some cities and not others. However, it is important to emphasize that our setting is conceptually different. By construction, public banks are present in all control cities prior to the reform. Therefore, any bank-specific shock after 2004 will affect both treated and control cities.

This setting therefore directly addresses two standard concerns about empirical designs featuring bank entry. First, if public banks

experience idiosyncratic shocks that affect their credit supply after 2004 (either due to shocks to their cost of funding or because they face political pressure to extend credit), both treated and control cities will benefit from a credit expansion, and our coefficient of interest will not be biased.²⁰

Second, potential correlations between financial inclusion policies and the expansion of other social welfare programs after 2004 are also addressed due to a specificity of the Brazilian institutional context. Most of the large-scale welfare programs, and in particular *Bolsa Família*, are distributed via public banks. Therefore, all cities (including cities in the control group) would benefit from the creation or expansion of such programs. We also conduct more detailed tests about this specific concern in Section 7.2 and find no evidence that it biases our results.

(ii) **Covariate balance and ex-ante differences.** The second main concern is that ex-ante differences lead treated cities to respond differentially to aggregate shocks.

We address this problem in four ways. First, as we show in Fig. 3, using a parsimonious matching estimator allows us to obtain covariate balance across a wide range of proxies for exposure to commodity-driven aggregate growth, economic integration of the city, and exposure to welfare programs promoted by left-leaning governments. Second, we show that treated and control cities were on a similar trend before the reform for a host of outcomes (credit, employment, number of firms, inequalities) in Sections Section 4, 5, and 6. The parallel-trends pre-reform indicate that any remaining unobserved differences that could drive the estimated effects would need to have not mattered before 2004 and only mattered afterward.

Third, we directly control for a collection of additional city-level characteristics. We show in Appendix Tables A4 and A5 that point estimates are very stable to the inclusion of controls such as GDP

²⁰ To be precise, the coefficient is not biased under the assumption that this bank-specific shock affects all branches of the bank in the same way. This would not be the case if this bank shock affects the bank’s branches differentially, for instance, due to age differences across branches.

per capita, employment in the commodity sector, skilled employment, political affiliation of the mayor, trade, distance to the state capital, and the co-movement of local GDP with aggregate fluctuations, as well as to all the different combinations of such controls.²¹

Finally, we exploit the granularity of our data and adapt Eq. (1) into a D-i-D estimator at the city-by-industry level, which allows us to relax the assumptions needed to identify the effect of the reform. Because we can now include industry-by-year fixed effects and therefore non-parametrically control for time-varying unobserved industry shocks, the effect of the reform remains unbiased even if treated and control cities are unbalanced in their exposure to sector-specific shocks (for instance because treated cities have more employment in the commodity sector).

Specifically, we estimate the regression:

$$Y_{i,c,g,t} = \beta \text{Treated}_c \times \text{Post}_{t \geq 2004} + X_{i,c,t} + \gamma_{i,c} + \delta_{i,g,t} + \varepsilon_{g,c,t} \quad (2)$$

The key difference in Eq. (2) relative to our city-level D-i-D is that we can include $\delta_{i,g,t}$, i.e., matched group-by-industry-year fixed effects. These fixed effects mean that β is estimated by comparing the *same* industry across treated and control cities that belong to the same matched group. This implies in particular, that sector-specific level shocks post 2004, such as commodity booms or productivity shocks specific to certain sectors, cannot bias the estimation of β . We report the results and details of the estimation in Section 7.1.

4. Effect on financial inclusion

In this section, we start by describing how the policy affected financial inclusion by fostering bank entry and increased deposits and lending in similar proportion. We then discuss the possible frictions that can explain the patterns we observe in the data and how the reform could affect real outcomes.

4.1. Higher access to bank branches

We start by showing that the reform increased access to bank branches, as the entry of government-owned banks did not crowd out private banks. To do so, we use as the LHS variable in Eq. (1) a dummy that equals one if the city has any bank branch (private or public), which allows us to estimate the change in likelihood for treated cities to have access to a bank branch.

We also decompose this dummy between a dummy that equals one if the city has a public bank branch, and a dummy that equals one if the city has a private bank branch. Fig. 5 reports the event study coefficients of our difference-in-differences estimation. Panel (a) shows results when the LHS is a dummy that equals one if the city has any bank branch (public or private), while panel (b) decomposes the change in likelihood to have a bank branch (the gray circles) into the change coming from access to a public bank (the blue diamonds) and access to a private bank (the green triangles).

Two facts are noteworthy. First, the probability of having a branch from a private bank in treated and control cities evolve in close parallel prior to the reform. This result indicates that private banks in treated and control cities evolved in the same way during the large credit boom that Brazil experienced prior to the reform, and remain on similar trends even after the reform.

Second, the expansion of public banks only modestly crowds out private banks, resulting in a large increase in *overall* financial development for treated cities. The probability of having a public bank branch or any bank branch increases sharply after 2004, in line with the aggregate pattern reported in Fig. 1, and it continues to increase progressively throughout the period with no mean reversion post reform.

²¹ Given that the reform may have a direct impact on many city characteristics, using time-varying controls would potentially bias our coefficients of interest. This is commonly referred to as the problem of “bad controls” (e.g., Angrist and Pischke, 2008). We address this problem by using the pre-reform value of these controls interacted with year fixed effects.

4.2. Increased deposits and credit

We then study how this higher presence of bank branches affects deposits and credit supply. In order to measure financial development net of any mechanical wealth effects, we compute total credit and total deposits scaled by city GDP. Because treated cities often have zero credit and deposits before the reform (since many have no bank branch) and, by construction, have zero credit and deposits from government-owned banks, we estimate the effect of the policy using the inverse hyperbolic sine transformation of the log function.²²

We show the event study coefficients of our difference-in-differences estimation for total deposits and total credit in Figure A1. The event studies show that credit and deposits for treated and control units evolve in close parallel prior to the reform and that the expansion in credit and deposits is entirely driven by public banks, with minimal crowding out of private banks. Panel (a) of Figure A1 shows that the initial increase in credit after the reform continues throughout the period and is driven entirely by public credit. There is a modest decline in private credit after 2010, but the total amount of credit still rises substantially after the reform, implying that overall, treated cities benefit from an increase in credit. Panel (b) of Figure A1 reports analogous results for deposits, and shows that deposits increase sharply in 2005 and continue to rise throughout the post-reform period.

We report pooled estimates in Table 2 and confirm the results of Figure A1. For all variables, the reform has a strong and significant effect on financial inclusion, driven by government-owned banks. The probability of having a bank branch increases by 18.6 p.p. (column 1), which fosters the accumulation of total deposits and total credit at the city level (columns 4 and 7).

These results are suggestive evidence that financial inclusion policies might be successful at starting a virtuous circle between savings and credit, since the average expansion of credit in a city matches the average increase in deposits flowing into the branches of the same city. In this case, bank branches in treated cities would be able to increase their credit supply without having to use capital coming from other cities, but instead could engage in a self-sustained expansion of credit and deposits.²³

5. Effect on economic development

5.1. Average effect

We start by estimating the effect of the reform on aggregate outcomes at the city level. Standard models of macro-finance development emphasize that financial frictions hamper economic development because talented but poor individuals are unable to start a firm (misallocation of talent) and existing productive but cash-poor firms are unable to expand their business (misallocation of capital). As financial development progresses, more firms are created and existing firms grow, generating higher demand for labor that translates into higher wages.

We test how the financial inclusion policy affects the different elements of this causal chain by estimating Eq. (1) with the total number of firms, total employment, employment growth at small firms

²² The inverse hyperbolic sine transformation of the log function is defined as $\log[X + (X^2 + 1)^{1/2}]$. Except for very small values of X , the inverse sine is approximately equal to $\log(2X)$ or $\log(2) + \log(X)$, and so it can be interpreted in exactly the same way as a standard logarithmic dependent variable. But unlike a log variable, the inverse hyperbolic sine is defined at zero and is less sensitive to jumps around zero than the more widely used $\log(X + 1)$ transformation.

²³ While the SUTVA assumption is, by definition, untestable, we also show in Appendix Figure A2 the evolution of total credit separately for treated and control cities. While the growth of credit in treated cities accelerates after 2004, it does not appear to be at the expense of a slowing down of credit in control cities.




Fig. 5. Effect of the program on having a bank branch.

This figure plots the yearly coefficients and 95% confidence intervals of the difference-in-differences estimator in Eq. (1) of the 2004 bank reform. Dependent variables are dummies that equal one if the city has at least a branch of any bank, a public bank, or a private bank, respectively. Standard errors are clustered at the matching-pair level.

Table 2

Effect of the program on bank branches, credit, and deposits.

Dependent variable:	Has Bank Branch			Deposits			Credit		
	All (1)	Public (2)	Private (3)	All (4)	Public (5)	Private (6)	All (7)	Public (8)	Private (9)
Treated×Post	0.186*** (0.016)	0.425*** (0.017)	−0.022** (0.009)	0.880*** (0.070)	1.295*** (0.074)	−0.098*** (0.037)	0.935*** (0.072)	1.557*** (0.071)	−0.154*** (0.036)
City FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Match×Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	79,995	79,995	79,995	79,995	79,995	79,995	79,995	79,995	79,995

This table shows the effect of the reform on financial development outcomes at the city level using the difference-in-differences estimator in Eq. (1). Has Bank Branch variables are dummies that equal one if the city has a branch of any bank, a public bank, or a private bank, respectively. Credit and deposits are both scaled by city GDP, and are in arcsin logs (inverse hyperbolic sine transformation of the log function). Standard errors are clustered at the matching-pair level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels.

(less than 20 employees), large firms, average wage and the number of industries as outcomes. Table 3 reports the results of these different regressions. In column 1, we show that the number of firms increases by 9.8% and that total employment increases by 10%. This increase is concentrated at small firms (less than twenty employees), which expands twice as fast as large firms (column 3 vs. 4). This increase in the demand for labor explains why average wage increases by 4.1%.

In column 6, we study how the reform affected industry dynamics. Consistent with models emphasizing that economic development requires countries to diversify their industrial base and explore their comparative advantage (e.g., Hausmann and Rodrik, 2003; Imbs and Wacziarg, 2003), we find that financial development increases the number of industries, which we measure as the number of distinct 4-digit industries (column 6).²⁴

We reproduce this analysis in graphical form by estimating the event study version of Eq. (1) in Fig. 6. In all cases, we find that treated cities display no pre-trend relative to control cities. We also find that each outcome increases progressively over time after the reform and stabilizes at a new high after five years, consistent with the notion that the reform relaxed financial constraints and allowed the local economy to reach a new steady state with a higher level of development.

Discussion of magnitudes. While we show in Section 4 that the reform led to financial development, we do not observe the entire effect of the reform and therefore cannot use these results as a “first stage”. In particular, *all* treated cities that did not experience the entry of a bank branch after the reform obtained at least a banking correspondent (Ministério da Fazenda, 2007), but we unfortunately cannot observe it in

the data.²⁵ And while the evidence presented in Section 2.1 indicates that correspondents do not provide the quality and scope of financial services as bank branches, having a correspondent almost certainly provides more access than having neither a correspondent nor a branch, suggesting that all cities receive some level of treatment. Therefore, our estimates on financial outcomes under-estimate the true impact of the reform on financial development, and rescaling the coefficients on the economic development outcomes shown in Table 3 by the point estimates in Table 2 – as in a standard 2SLS approach – would inflate the true magnitude of the elasticities.

Due to this caveat, we think the more natural approach is to directly interpret the point estimates in Table 3 as the elasticity of economic development outcomes with respect to the introduction of formal financial services. In this context, the two closest experiments to our setting are (Barboni et al., 2023), which looks at the entry of bank branches in Indian villages, and Bruhn and Love (2014), which looks at the opening of bank branches in stores of a large retailer of consumer goods focused on underserved and low-income clients.

(Barboni et al., 2023) finds that a new bank branch leads to an 8% reduction in poverty and a 6% increase in average income. The latter is comparable to the 4.1% increase in average wages that we estimate in our setting. Bruhn and Love (2014) finds similar estimates, with income increasing by 7%, employment by 1.4%, and informal businesses by 7.6%, although formal business is unaffected. Our larger effects on employment and business creation can be explained by the fact that our experiment improved financial development at the city level, and

²⁵ As explained in Section 2.2, data on bank branches does not keep track of banking correspondents.

²⁴ There are 614 distinct industries and the definition is consistent over time.

Table 3
Effect of the reform on economic development.

Dependent variable	# Firms	Employment all	Employment small firms	Employment large firms	Average wage	# Industries
	(1)	(2)	(3)	(4)	(5)	(6)
Treated×Post	0.098*** (0.011)	0.100*** (0.015)	0.214*** (0.023)	0.116*** (0.018)	0.041*** (0.005)	0.088*** (0.008)
City FE	✓	✓	✓	✓	✓	✓
Match×Year FE	✓	✓	✓	✓	✓	✓
Observations	79,995	79,995	79,947	79,995	79,995	79,995

This table shows the effect of the reform on economic development at the city level using the difference-in-differences estimator in Eq. (1). All variables in columns 1–6 are in logs. In columns 3 and 4, employment is decomposed between firms with less than 20 employees (“small firms”) and more than 20 employees (“large firms”). The number of industries (column 6) is the number of distinct 4-digit industries in the city-year. Standard errors are clustered at the matching-pair level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.




Fig. 6. Effect of the program on firms, employment, and wage.

This figure plots the yearly coefficients and 95% confidence intervals of the difference-in-differences estimator in Eq. (1) of the 2004 bank reform. Dependent variables are logs of total number of firms, total employment, and average wage in panels (a), (b) and (c) respectively. Standard errors are clustered at the matching-pair level.

therefore is more likely to have positive “local GE effects”. The longer time period over which we can trace out the effect of the reform can also partly explain the difference since resources reallocate slowly, particularly in developing countries (e.g., Buera and Shin, 2013). These slow-moving changes underscore the importance of measuring and taking into account transitional dynamics when estimating the effect of reforms on economic development.

5.2. Mechanisms

There are two main channels through which financial development promotes economic growth in this setting. First, bank expansion can

foster *local demand* either because it relaxes individuals’ borrowing constraints and reduces their need for precautionary savings, or because the opening of a bank branch functioned as fiscal stimulus.²⁶

Second, bank expansion can foster *supply* by reducing investment frictions, thereby boosting investment of existing firms and facilitating the entry of new firms. In this case, the differential dependence on local demand should not matter and we expect both tradable and non-tradable industries to benefit from the reform.

²⁶ Financial inclusion will reduce the need of precautionary savings for instance because of limited insurance in developing countries.

Table 4
Employment in tradables and non-tradables.

Dependent variable	Employment		Value of traded goods		Geographic concentration	
	Manufacturing					
	Yes	No	Yes	No	Yes	No
Tradable definition						
Tradable	(1)	(2)	(3)	(4)	(5)	(6)
Treated×Post	0.200*** (0.037)	0.123*** (0.015)	0.190*** (0.042)	0.124*** (0.015)	0.146*** (0.025)	0.126*** (0.015)
Match×Industry×Year FE	✓	✓	✓	✓	✓	✓
Observations	106,574	414,795	84,772	436,597	129,728	391,641

This table reports the effect of the policy on employment at the city-by-(4-digit) industry level. Data are collapsed as an average “pre” ($t \leq 2004$) and the average “post” ($t > 2004$) periods, and each dependent variables are the midpoint growth rate $g_{j,c}^X = [(X_{j,c,t} + X_{j,c,t-1}) \times 0.5]$. Each cell is weighted by $g_{j,c}^X / (\sum_{j \in c} g_{j,c}^X) \times pop_{2000}$. See Section 7.1 for a detailed explanation of the construction. In columns 1–2, tradable is defined as firms in the manufacturing sector. In columns 3–4, we define tradable industries based on the value of exports and imports in the custom data aggregated at the sector level. In columns 5–6, tradable is defined using the geographical HHI of employment of each industry. “Tradable” (column 5) corresponds to an HHI in the top quartile (i.e., high level of geographic concentration). Standard errors are clustered at the matching-pair level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

5.2.1. Consumption vs. business development channel

To test if most of the effect is coming from a bank-expansion-induced increase in demand, we decompose growth in aggregate employment at the city level by firms in non-tradable vs. tradable industries. Indeed, since non-tradable industries are more dependent on local demand than tradable industries – because, by definition, tradable industries produce goods that can be sold across the whole country, if not worldwide – an increase in local demand driven by the reform should benefit non-tradable industries relatively more.

To do so, we estimate Eq. (2) and split the regression between tradable and non-tradable, which requires us to work at the city-industry level. This requires a slight modification to our specification since the reform had an impact on the entry and exit of industries at the city level (Table 3–column 6), implying that the baseline specification of Eq. (1) at the city-by-industry level will not match the aggregate results at the city level.

We explain in detail how we account for this adjustment in Section 7.1. Briefly, we create a balanced panel and compute the mid-point growth rate between the average pre period (before to 2004) and post period (after 2004). We show in Table 11 that this specification preserves the aggregate city level results and report that, for all variables, it produces very similar point estimates as the ones obtained with our baseline log panel specification.

We use three methods to identify industries that produce tradable goods and are therefore not dependent on local demand. First, we classify an industry as tradable if it is in the manufacturing sector, and non-tradable otherwise. Second, we flag industries as tradable if the imports plus exports equal to at least \$10,000 per worker, or if total exports plus imports for the industry four-digit industry exceeds \$500M.²⁷ Third, we compute the geographical dispersion (HHI) of employment at the industry level and classify tradable industries as those in the top quartile of the HHI distribution. The intuition behind this proposed measure is that, since non-tradable industries have to be consumed locally, they should be less geographically concentrated.

We report the results of the effect of the reform on tradable and non-tradable industries in Table 4. We find that employment growth is almost two times bigger in tradable industries depending on the exact definition (e.g., column 1 vs. 2). This implies that while the entry of bank branches in the cities might have had a direct effect on demand, a substantial part of the increase in economic development induced by the financial inclusion policy is coming from a relaxation of financing constraints for entrepreneurs.

The fact that firms in tradable industries benefited more from the reform than firms in non-tradable industries is also consistent with

financial inclusion policies promoting economic development by providing the liquidity needed for transactions. Indeed, firms and households need a place to safely store their liquidity and transact with others. This is particularly true for firms in tradable industries that need to pay suppliers and receive customer payments from firms that are physically far away. Banks provide such services, which facilitate trade and business transactions (see Xu and Yang (2022) for an example in 19th-century US).

5.2.2. Why do local branches matter? the role of distance

Next, we investigate why the presence of local banks appears to relax firms’ financial constraints. A classic assumption in the macro-development and finance literature is that geographical proximity reduces banks’ monitoring and screening costs (e.g., Greenwood et al., 2010; Ji et al., 2023) in particular in the presence of soft information (e.g., Rajan and Zingales, 2001; Hombert and Matray, 2017), which is prevalent in developing countries.

The main prediction of these models is that the effect of the policy should increase with the ex-ante distance to the nearest bank. We test this prediction in two ways. First, we compare cities with a local (private) bank before the reform with cities that did not have a private bank by interacting the *Treated* × *Post* variable with a dummy variable *No Private Bank_c* that takes the value one if the city did not have a private bank pre-reform. Second, within the set of treated cities that have no bank branches prior to the reform, we compute the distance to the nearest city with a bank (public or private). We estimate the conditional effect of distance in the full panel by interacting the main *Treated* × *Post* variable with a dummy equal to one if the distance is zero, and then with a continuous variable *Distance to the nearest bank* that is the (log) distance to the nearest bank.

This analysis also allows us to test whether our results are specific to the presence of government-controlled banks—for instance, because these banks extend subsidized credit. To do so, we separately interact our main *Treated* × *Post* variable with the distance to the nearest public bank and with the distance to the nearest private bank. Intuitively, if our results hinge on access to *subsidized* credit, the effect of the reform should be stronger for cities that were far from a government-controlled bank than for those that were far from a private bank.

We report these results in Table 5. In panel A, we show that our results are much stronger in cities with no bank presence prior to the reform. Cities without a local private bank before the reform experience a larger increase in the number of firms (10.4%, column 1), in employment (13.9%, column 2), and in average wages (7%, column 5). In panel B, we show that, conditional on not having a local private bank pre-reform, the real effects of financial inclusion increase with the distance to the nearest bank.

Both sets of results are consistent with the reform promoting economic development by reducing the distance between borrowers and lenders, which lowers the monitoring and screening costs of financial

²⁷ Numbers are in 2006 values. We manually build a crosswalk between US NAICS codes and the Brazilian industry classification.

Table 5
Financial frictions, the role of distance.

Dependent variable	# Firms	Employment all	Employment small firms	Employment large firms	Average wage
	(1)	(2)	(3)	(4)	(5)
Panel A: Private bank before the reform					
Treated×Post	0.054*** (0.014)	0.042*** (0.016)	0.111*** (0.026)	0.050*** (0.020)	0.012* (0.006)
Treated×Post×No private bank	0.104*** (0.023)	0.139*** (0.031)	0.246*** (0.048)	0.158*** (0.038)	0.070*** (0.011)
Panel B: Distance to nearest bank					
Treated×Post	0.054*** (0.014)	0.042*** (0.016)	0.111*** (0.026)	0.050*** (0.020)	0.012* (0.006)
Treated×Post×Distance to nearest bank	0.068*** (0.012)	0.108*** (0.013)	0.164*** (0.034)	0.131*** (0.022)	0.022*** (0.007)
Panel C: Distance to nearest public bank					
Treated×Post	0.054*** (0.015)	0.042** (0.017)	0.111*** (0.028)	0.050** (0.021)	0.012* (0.007)
Treated×Post×Distance nearest public bank	0.069*** (0.012)	0.110*** (0.013)	0.166*** (0.030)	0.132*** (0.020)	0.024*** (0.006)
Panel D: Distance to nearest private bank					
Treated×Post	0.054*** (0.015)	0.042** (0.017)	0.111*** (0.028)	0.050** (0.021)	0.012* (0.007)
Treated×Post×Distance nearest private bank	0.070*** (0.011)	0.113*** (0.013)	0.174*** (0.027)	0.136*** (0.019)	0.026*** (0.007)
City FE	✓	✓	✓	✓	✓
Match×Year FE	✓	✓	✓	✓	✓
Observations	79,995	79,995	79,947	79,995	79,995

This table shows the effect of the policy on multiple outcomes interacted with various measures of ex-ante distance between treated cities and existing banks using the difference-in-differences estimator in Eq. (1). In Panel A, *No private bank* is a dummy equal to one if the treated cities did not have a private bank prior to the reform. In Panel B, *Distance to the nearest bank* is the distance to the nearest city with a bank, for treated cities with no bank prior to the reform. Panel C and D compute this distance for the nearest private and the nearest public bank. Standard errors are clustered at the matching-pair level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

intermediaries. This idea is also consistent with the result that small firms (fewer than 20 employees) benefit more from the reduction in distance than large firms (column 3 vs. 4), in line with the idea that small firms are more opaque and more intensive in soft information.

The fact that the effect of the reform on economic development increases with the initial distance of treated cities to existing bank branches is also consistent with the “bank liquidity channel”. Indeed, a reduction in distance to a branch may matter simply because it makes it less costly to access liquidity services.

In panels C and D, we report similar results for the distance to the nearest public and private bank, respectively. Across all variables, we find that the conditional effect of distance is quantitatively similar whether we measure distance to public banks (panel C) or private banks (panel D). These findings suggest that our results are not specific to services provided by government-controlled banks, such as access to subsidized credit. Indeed, if access to government-owned banks mattered in itself, we should find a larger effect of distance to a public bank relative to the distance to a private bank.

5.3. Discussion of the reform

Interpretation: improved access to financial services. The permanent shift in financial development, the similar increase in deposits and credit, the sustained economic development principally driven by the growth of the tradable sector, and the fact that these effects increase

with the ex-ante distance to a bank branch imply that the financial policy we study fostered access to mainstream financing services for firms.

This policy operates through a distinct mechanism that is inconsistent with two alternative interpretations: (i) the reform was a one-time capital infusion coming from outside the city that potentially generated excessive borrowing, and (ii) the policy was a generic government spending policy.

On (i), the permanent change in the steady state of financial development, together with the sustained economic development over ten-year with no sign of mean reversion, rules out the concern that the policy was fueled by bad loans, as this should trigger a boom-bust cycle.

On (ii), the long-lived effect on economic development is also inconsistent with a short-term stimulus. In addition, the fact that employment in the tradable sector expands more than employment in the non-tradable sector (Table 4) and that the effect increases proportionally with the distance to the nearest bank are also inconsistent with these results being driven by higher spending in treated cities overall.

Is the reform beneficial in net? While the reform was successful in terms of fostering substantial economic development, it does not necessarily mean that this success is unambiguous. Indeed, the cost of operating bank branches in these unbanked cities might be substantial. Doing a clear cost–benefit analysis of the program is beyond the scope of this paper and would require access to data on the total costs of setting up and operating local bank branches, which do not exist to the best of our knowledge.




Fig. 7. Average ROA and expense-to-revenue ratio.

This figure shows the evolution of ROA (panel a) and total operating expenses divided by total operating revenues (panel b) for new branches in treated and control cities, defined as branches that were opened after the reform, as well as for all branches in control cities.

Nonetheless, ESTBAN contains information about expenses and profitability at the branch level. This allows us to look at the yearly average expense-to-revenue ratio and ROA for the public branches created in treated cities after the reform, and to compare them to the average for all private branches in control cities and for new private branches in control cities.

Fig. 7 plots the evolution between 2005 and 2014, and shows that branches opened in treated cities are not less profitable or more expensive to run than those opened in control cities. Panel (a) displays that the average ROA of new branches in treated cities is, if anything, higher than that of new branches in control cities and comparable to the average ROA of all branches, new and existing, in control cities. Panel (b) displays that the average expense-to-revenue ratio of new branches in treated and control cities is similar, but higher than that of all branches in control cities. This suggests that new branches might be less profitable than existing branches, but does not point to differences in the profitability of new branches in treated and control cities.

Of course, it is still possible that the fixed costs associated with setting up a bank branch are much higher in treated cities. But, conditional on paying this fixed cost, there is no evidence that the bank branches of public banks in treated cities are particularly subsidized.

Why do banking deserts exist in the first place? The permanent shift in deposits and credit, together with the sustained effect on economic activity that the financial inclusion policy triggered and the similarity in bank branch ROA in treated and control cities, imply that treated cities that received a bank branch were profitable markets. Then why did banks not serve these markets prior to the reform? This question is beyond the scope of our paper, but we discuss possible explanations in Appendix A.2. Briefly, the lack of coverage can be explained by banks being in a situation of monopolistic competition and facing downward-sloping demand curves, which implies that banks can maximize their markups and hence their profits by restricting quantities, *i.e.*, limiting their branch expansion.

On the cost side, banks might be able to minimize the cost of their branch network by extending their network in a capillary way and only up to a certain limit. In addition, the existence of sunk costs in setting up a branch might imply that, while branches are profitable once they operate, the total cost (variable plus fixed cost to set up) might still be too high, particularly if set-up costs are larger in more remote cities. As such, a possible rationale for why bank branches are not present in treated cities prior to the reform, but their entry has such a strong positive effect on economic development, is that using public

banks to open up branches in unbanked cities is akin to the government subsidizing the set-up cost of a bank branch.

6. Effects on inequality

6.1. Aggregate results


To study how the aggregate economic gains produced by the financial inclusion policy are distributed in each local labor market, we estimate Eq. (1) using the wage Gini at the city level as an outcome, as well as the average wage per worker in each bin of the city-level wage distribution. We graphically report the result for the evolution of Gini and the change in average wage for each quartile of the wage distribution in Fig. 8. Fig. 8(a) shows the effect of the reform on the Gini coefficient. As before, treated cities display no differential pre-trend prior to the reform. Following the reform, we find a continuous increase in Gini, implying an increase in wage inequality. The magnitude is substantial, with treated cities having a Gini index that is two points higher ten years after the reform relative to control cities, which represents an increase of 7% relative to the pre-reform mean.

While this result shows that higher financial development leads to higher inequality, it does not tell us why the Gini is increasing in treated cities.


We therefore unpack the evolution of Gini by computing the average wage for each quartile of the city wage distribution to better understand the source of the overall change in inequality. To do so, we estimate the distribution of wage within each city-year cell, split the sample into quartiles, and take the mean wage in each cell.

Fig. 8(b) reports the evolution of each wage quartile. Consistent with the idea that economic development is a “tide that lifts all boats”, we find that all workers benefit from the reform. However, workers in the first quartile of the distribution (the purple line) gain far less than workers in the last quartile (the red line), and wage gains increase monotonically with the initial position in the wage distribution.

Table 6 reports estimates of Eq. (1). The point estimates tend to underestimate the effect of the reform on inequality since, as Fig. 8(b) shows, inequality rises steadily over time, while these regression results show the average over the whole post-reform period. The Gini increases on average by 1.2 points (column 1), which is driven by larger wage gains at the top of the income distribution. Individuals in the bottom quartile of the wage distribution experience an increase in their average wage of 1% (column 2), while individuals in the top quartile see their wages increase by 5.5% (column 5), five time more.



(a) Gini index



(b) Wage distribution

Fig. 8. Effect of the program on wage inequality.

This figure plots the yearly coefficients and 95% confidence intervals of the difference-in-differences estimator in Eq. (1) of the 2004 bank reform on city-level wage Gini (Fig. 8(a)). In Fig. 8(b), the wage distribution is computed every year at the city level. Standard errors are clustered at the matching-pair level.

Table 6

Effect of the program on wage inequality.

Dependent variable:	Gini	Wage			
	(1)	[0–25th] (2)	[25th–50th] (3)	[50th–75th] (4)	[75th+] (5)
Treated×Post	0.012*** (0.002)	0.010*** (0.003)	0.024*** (0.004)	0.034*** (0.006)	0.055*** (0.007)
City FE	✓	✓	✓	✓	✓
Match×Year FE	✓	✓	✓	✓	✓
Observations	79,995	79,995	79,995	79,995	79,995

This table reports the effect of the policy on earnings inequality at the city level. In columns 2–4, the dependent variable is the (log) average wage for each bin of the wage distribution in a city-year cell. Standard errors are clustered at the matching-pair level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

6.2. Mechanisms

We explore three channels that can account for the increase in inequality following a reduction in financial frictions: better matching, skilled-labor demand and constrained skilled-labor supply. First, financial development might lead to better employer-employee matching. This could happen either because looser financial constraints on individuals allow them to search longer and to find a better job match, or because less severe financial frictions can allow productive firms to front-load wages and attract more productive workers, resulting in a reduction in labor misallocation and higher wages at the top of the distribution.

Second, financial development can foster higher labor demand for skilled workers relative to unskilled workers. Financial frictions can directly impact labor demand if there is a mismatch between payments to labor and the generation of cash-flows or if labor has a fixed-cost component due to hiring and firing costs (Schoefer, 2021; Benmelech et al., 2021). Since skilled workers require higher wages and are arguably more expensive to recruit and train, financial frictions disproportionately constrain the demand for skilled labor and, when lessened by the reform, lead to an increase in the demand for skilled labor relative to unskilled labor.

Alternatively, if capital and skilled labor are relative complements, looser financial constraints can increase capital investment and, consequently, increase the marginal productivity of skilled workers relative to unskilled workers, also leading to an increase in the relative demand for skilled workers (Fonseca and Doornik, 2022). A testable implication of either version of the skilled labor demand hypothesis is that, as the relative demand for skilled workers rises, both the relative price and the relative quantity of skilled workers should rise, leading to an increase in the skill premium *and* in the share of skilled workers in treated cities.

Third, labor demand might go up uniformly across the skill distribution, but the supply of unskilled workers could be more elastic than the supply of skilled workers. In this case, the skill composition of firms remains stable, but the price of skilled workers goes up, particularly so in cities facing higher shortages of skilled workers.

Better matching. To test if the matching between workers and firms improves following the reform, we build on (Eeckhout and Kircher, 2011) and Lopes de Melo (2018), which give a structural interpretation to the firm fixed effects in Abowd et al. (1999) regressions and show that better matching should reduce the dispersion of worker ability within the firm.²⁸

We proxy for worker type with the average log wage over all job spells. We compute the standard deviation of worker types at the firm-year level, residualize the variable from firm fixed effects to account for changes in industry-city composition over time, and take the mean of the residualized dispersion in worker types at the 2-digit-industry-by-city level for each year. We can then test whether the average dispersion declines as a consequence of the reform.²⁹

Table 7 shows results of this exercise. Across all specifications, we find that if anything, the within-firm dispersion in worker type increases (by a small amount relative to the pre-reform average of 0.34). This is the opposite of what we would expect from an improvement in employer-employee matching, which should lead to lower within-firm dispersion in worker types.

Increase in demand for skilled workers. To test whether a change in the relative demand for skilled workers can explain the rise in wage inequality, we need an ex-ante, time-invariant definition of skill. We

Table 7

Dispersion in worker type.

Dependent variable:	Std. Dev. worker type			
	(1)	(2)	(3)	(4)
Treated×Post	0.026* (0.015)	0.027~ (0.017)	0.026* (0.016)	0.027 (0.020)
City FE	✓	—	—	—
City×Industry FE	—	✓	✓	✓
Match×Year FE	✓	✓	✓	—
Industry×Year FE	—	—	✓	—
Match×Industry×Year FE	—	—	—	✓
Observations	1,286,478	1,286,478	1,286,478	1,286,478

This table shows the effect of the reform on the change in the average within-firm standard deviation of worker type at the city-by-(2 digit) industry level. Worker type is measured as the average log wage over all job spells of a given worker. We then compute the standard deviation of worker types at the firm-year level and residualize this variable from firm fixed effects. Standard errors are clustered at the matching-pair level. ***, **, *, ~ indicate statistical significance at the 1%, 5%, 10% and 11% levels, respectively.

leverage the fact that the Brazilian matched employer-employee data allow us to observe education and classify workers as skilled if they have at least some college education and unskilled otherwise.³⁰

In Table 8, we start by showing that this measure tracks the evolution of inequality well. In column 1, we show that the skill premium increases by 8.3% (column 1) and that this increase is driven by a much faster increase in the wage of skilled workers (+11.8%, column 2) than unskilled workers (+2.8%, column 3). These magnitudes are actually bigger than the wage increase in the top quartile of the distribution (+5.5%, column 5- Table 6) relative to first quartile (+1%, column 2- Table 6), which suggests that the increase in inequality reflects an increase in the returns to skill.

Absent labor supply constraints or other frictions, a credit-fueled rise in the relative demand for skilled labor increases the relative quantity of skilled labor (e.g., Fonseca and Doornik, 2022). While the coefficient for the share of skilled workers is positive and significant at 10%, the magnitude (+0.2%) is very small compared to the 8% increase in the skill premium. This suggests that other frictions, such as labor supply constraints, are necessary in order to explain the bulk of our results. In Appendix Table A9, we show that we find similar results at the industry-by-city level controlling for time-varying industry shocks.

Constraints in the supply of skilled workers. To argue that a city's own supply of skilled workers is a driver of higher wage inequality, we first need to establish that worker mobility across cities is limited. To do so, we exploit the panel dimension of our data to decompose the number of workers in a given city-year into "local", defined as workers who are already in the city prior to the reform, "movers", defined as workers who were living in a different city prior to the reform, and "new", defined as workers who appear for the first time in labor-market data in a given city and did not come from another city.

Table 9 estimates the effect of the reform on the composition of workers across these three groups for all workers (columns 1–3) and skilled workers only (columns 4–6). We find that the reform has no effect on the share of workers coming from other cities in general (column 2), and that it has a positive but very small effect (+0.7%) when we focus on skilled workers (column 5). This implies that the reform had a limited effect on domestic migration and that cities that benefited from the financial inclusion policy did not experience an important inflow of skilled workers.

While the low domestic migration of skilled workers following the reform might seem surprising given the skill premium increase in

²⁸ Another potential way of testing for sorting would be to study the correlation between firm and worker fixed effects, but, as (Eeckhout and Kircher, 2011) and Lopes de Melo (2018) show, this correlation does not measure the strength of sorting in a general setting.

²⁹ See Bombardini et al. (2019) for an application of this method in a trade context.

³⁰ This is a less stringent definition than studies looking at developed countries who use college education as a proxy, since we include college dropouts in our definition of skilled.

Table 8
Demand for skilled workers.

Dependent variable	Skill premium (1)	Wage skilled (2)	Wage unskilled (3)	Share skilled (4)
Treated×Post	0.083*** (0.010)	0.118*** (0.012)	0.028*** (0.006)	0.002* (0.001)
City FE	✓	✓	✓	✓
Match×Year FE	✓	✓	✓	✓
Observations	79,901	79,901	79,995	79,995

This table shows the effect of the reform on the skill premium (column 1), the average wage of skilled and unskilled workers (columns 2 and 3), and the share of workers that are skilled (column 4) at the city level. Skilled workers are defined as workers with at least some college education. All dependent variables are in arcsin-logs. Standard errors are clustered at the matching-pair level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 9
Worker migration.

Sample:	All workers			Skilled workers		
Dependent variable:	Share local (1)	Share movers (2)	Share new (3)	Share local (4)	Share movers (5)	Share new (6)
Treated×Post	−0.019*** (0.004)	0.000 (0.001)	0.021*** (0.004)	−0.021*** (0.006)	0.007*** (0.002)	0.020*** (0.005)
City FE	✓	✓	✓	✓	✓	✓
Match×Year FE	✓	✓	✓	✓	✓	✓
Observations	79,995	79,995	79,995	79,901	79,901	79,901

This table shows the effect of the reform on the share of workers by migration status at the city level. Skilled workers are defined as workers some college education. “Local” workers are workers observed in the city before the reform. “Movers” are workers that we observe in a different city before the reform. “New” are workers that appear in the city for the first time. All dependent variables are in arcsin-logs. Standard errors are clustered at the matching-pair level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

treated cities, this can be explained by the existence of very large migration costs in Brazil, particularly for residents of poor cities (e.g., Porcher, 2022). We provide evidence for this hypothesis by estimating how the migration response varies as a function of migration cost. We proxy for migration cost using the share of movers during the pre-reform period and split the data into deciles of migration cost. We then estimate the effect of the reform on the share of within-country migrants for each decile of the migration cost distribution. Figure A3 in the Appendix reports the results. Consistent with outsiders being attracted by a higher skill premium when migration costs are low, we find an increase in the share of migrant workers in the first decile of migration cost, with an increase of 1%. However, this effect sharply drops to zero at the second decile and remains around zero afterwards.

Given the low rate of internal migration, an increase in the demand for labor (skilled and unskilled) can only be met by local workers. To proxy for the potential supply of skilled labor, we use the share of the local population with 11 years or more of education from the 2000 Demographic Census. The intuition behind this measure is that if a treated city faces a shortage of skilled workers, we should observe an abnormally large skill premium. In order to determine what is abnormally large, we compare the skill premium in treated cities with the skill premium in the same industry-by-firm-size category in control cities.

Our measure of the relative supply of skilled labor is based on the population census, which has the advantage of neither being affected by the fraction of workers in the informal sector, nor reflecting the equilibrium outcomes in the formal labor market. As a robustness check, we supplement this measure by computing a measure of the “skill gap” at the city level to construct that measure. We split firms into employment size quartiles according to the city-year distribution and, for each year in the pre-reform period, we compute the skill premium in each city-industry-firm-size cell for both treated and control cities. We then take the ratio of treated to control skill premium at the industry-firm-size level and define the skill gap as the city-level mean of all industry-firm-size ratios in a given city.

We split both measures along the sample median and interact each dummy with all the variables, including the fixed effects. Table 10 reports the results. The increase in Gini (column 1) is entirely explained

by the increase in inequality in cities where the fraction of skilled workers is low (column 2). Since we use an interaction term, the coefficient on the variable *Treated* × *Post* shows the result for the subsample of cities that are below the median of the supply of skilled labor. The total effect for cities with high supply skilled labor is obtained by adding the coefficient of *Treated* × *Post* with the marginal interaction term. Irrespective of the proxy (columns 2 and 3), we find that the total effect of the policy on inequality for cities with a high supply of skilled workers is much smaller (column 2), and it is close to zero when we measure the supply of skilled workers with the share of population with some college education (column 3). In Appendix Table A10, we show that these results are robust to using continuous versions of these skill supply measures and adding a wide range of control variables.

7. Robustness

7.1. City-industry level estimation

Controlling for industry-dynamics. Even though pre-reform covariates are balanced across treated and control cities (Fig. 3) and we show in Appendix Table A5 that our results are robust to directly controlling for these levels, it is possible that industry-specific shocks post 2004 might affect our results. Alleviating these concerns requires us to work at the city-by-industry level. However, since we observe entry and exit of industries at the city level, the baseline specification of Eq. (1) does not guarantee that aggregate results at the city level (which capture the extensive margins by construction) are preserved when we disaggregate the data at the city-by-industry level.

We are able to provide an alternative estimation that does ensure this aggregation property. We modify our baseline specification in two ways. First, we create a balanced panel by assuming that each industry we observe at any point in a given city is present during the whole sample period, and we fill observations without firms in an industry with zero. Second, we collapse the data into two periods: the average “pre” ($t < 2004$) and the average “post” ($t \geq 2004$). We then compute the mid-point growth rate for all our different outcomes, that we define for a variable X as: $g_{j,c}^X = (X_t - X_{t-1}) / [(X_{j,c,t} + X_{j,c,t-1}) \times 0.5]$.

Table 10
Effect on Gini, heterogeneity in skill supply.

Dependent variable:	Gini			
	(1)	(2)	(3)	(4)
Treated×Post	0.012*** (0.002)	0.016*** (0.002)	0.018*** (0.002)	0.020*** (0.003)
Treated×Post×Low skill gap		−0.008*** (0.003)		−0.006** (0.003)
Treated×Post×High share skilled population			−0.014*** (0.003)	−0.013*** (0.003)
City×Industry FE	✓	✓	✓	✓
Match×Year FE	✓	✓	✓	✓
Observations	79,995	79,995	79,995	79,995

This table shows the effect of the reform the Gini index at the city-by-(2 digit) industry level. In column 2, we split treated cities based on whether their fraction of population with at least 11 years of education is above or below the median of the sample distribution. In column 3, we estimate the ratio of skilled workers in treated cities relative to the national average, and split along the sample median. Standard errors are clustered at the matching-pair level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Specifically, we estimate the following equation at the city c , industry j , period t level:

$$\Delta Y_{c,j,t} = \beta_1 \text{Treated}_c \times \text{Post}_t + \delta_{j,t} + \varepsilon_{c,j,t} \quad (3)$$

Since $\Delta Y_{c,j,t}$ is the change between the pre and post period, we do not need to include city×industry fixed effects as they are already differenced out, but we do include industry-by-pair-by-time fixed effects $\delta_{j,t}$.

This specification has two appealing properties. First, it handles entry and exit of industries without relying on transformations of the log function (such as “x+1”), that are always sensitive to small variations around zero. Second, it ensures that the coefficient at the city-industry level aggregates exactly to the coefficient at the city level when using the correct weights, which is not possible with the non-linear log function. The weights are defined as the share of the denominator in the total city-period cell. For each industry j in city c , we use the mid-point growth rate for a variable X in city c and industry j $g_{j,c}^X$, and compute the weight as $g_{j,c}^X / (\sum_{j \in c} g_{j,c}^X)$.³¹

This specification allows us to include industry-by-year fixed effects (or even industry-by-match-by-year fixed effects), which ensures that the effect of the reform is now estimated by comparing the *same* industry across treated and control cities in the same matched pair. This implies, for instance, that even if treated cities are more dependent on the commodity sector in the midst of a commodity boom, these industry-specific dynamics will not bias our estimates.

In Table 11, we start by reproducing the baseline results at the city-by-industry level. In columns (1) and (4), we report results at the city level and show that they are very close to the baseline city-level results of Table 3. In columns (2) and (5), we show that the point estimates are *identical* at the city-industry level with the weighting described above. Finally, in columns (3) and (6), we show that the inclusion of match×year×industry fixed effects yield, if anything, larger point estimates. In this case, the identification relies solely on comparing outcomes in the same industry within a given group of treated-control cities. These additional fixed effects ensure that our baseline effects are not driven by industry shocks that might correlate with the reform and the sectoral composition of treated cities.

Estimating the underlying change in firm dynamics. An additional advantage of this industry-level specification that both accommodates zeros and uses a linear estimator is that we can exactly decompose the change in the number of firms in the cross section of cities into the evolution of entry and exit. To measure firm entry and exit, we count the number of firms entering or leaving the city each year and

set the year 2000 to zero. This allows us to decompose the change in the number of firms in a industry-city dell c, j as:

$$\Delta Firms_{c,j} = \frac{Firms_{c,j,2014} - Firms_{c,j,2000}}{Firms_{c,j,2000}} = \frac{\sum_{t=2001}^{t=2014} Entry_{c,j,t} - \sum_{t=2001}^{t=2014} Exit_{c,j,t}}{Firms_{c,j,2000}}$$

We report the results in columns 7 and 8 of Table 11. We find that the reform leads to a substantial increase in both firm entry *and* exit, supporting the view that financial development fosters a process of creative destruction. In terms of magnitude, the increase in the number of new firms created is almost two time bigger than the increase in number of firms (column 7 vs. column 6), which highlights the importance of having panel data rather than cross-sectional data in order to fully grasp changes in firm dynamics as positive local shocks accelerate churn.

7.2. Government programs

One potential concern is that the expansion of government and social welfare programs might be correlated with the entry of government-owned banks in treated cities after 2004. Of special concern is the far-reaching cash transfer program Bolsa Família, which was introduced one year before our reform, in 2003. We think that this concern is unlikely to explain our results for four reasons.

First, this mechanism is inconsistent with some of our results: (i) additional income from government programs could serve as a positive income shock, fostering growth by driving up local demand. This would imply that non-tradable sectors grow faster than tradable sectors, which is the exact opposite of what we find in Table 4. (ii) While government transfers can affect *income* inequality, there is no reason it should affect *wage* inequality *a priori*. A possible connection would be that higher government transfers increase the reservation wage of workers, but this would imply an increase in workers’ bargaining power that should mostly benefit workers at the bottom of the wage distribution, thereby predicting a *reduction* in wage inequality rather than the *increase* we find in Fig. 8.

Second, the most ambitious programs, such as Bolsa Família, are distributed directly by government-owned banks. Since our control cities have a branch of a public bank by design, this implies that control cities have the same access as treated cities to government programs disbursed through government-controlled banks.

Third, we show in Fig. 3 that treated and control cities are similar in their government expenditures, as well as in the likelihood that the mayor is affiliated with Lula’s party (the Worker’s Party). Therefore, even if, post 2004, the Lula government decided to expand social transfers particularly to places with more left-leaning voters, both treated and control cities would benefit from such an expansion in the same way.

Fourth, we test if the point estimates for our main outcomes are affected when we directly control for total local government expenditures

³¹ In our case, we multiply this weight by the population in 2000 in order to be able to exactly reproduce the city level results, which does not affect the aggregation property.

Table 11
Effect on economic development, city-industry level.

Dependent variable Unit of analysis	Employment			# Firms			Entry	Exit
	City	City×Ind.		City	City×Ind.		City×Ind.	City×Ind.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated×Post	0.098*** (0.015)	0.098*** (0.013)	0.126*** (0.016)	0.104*** (0.012)	0.104*** (0.011)	0.167*** (0.014)	0.303*** (0.028)	0.136*** (0.020)
Match×Year FE	✓	✓	✓	✓	✓	–	–	–
Match×Industry×Year FE	–	–	–	–	–	✓	✓	✓
Observations	5,333	521,369	521,369	5,333	511,058	511,058	511,058	511,058

This table reports the effect of the policy on economic development at the city-by-(4-digit) industry level. Data are collapsed as an average “pre” ($t < 2004$) and the average “post” ($t \geq 2004$) periods, and each dependent variable is the midpoint growth rate $g_{j,c}^x = [(X_{j,c,t} + X_{j,c,t-1}) \times 0.5]$. Each cell is weighted by $g_{j,c}^x / (\sum_{j \in c} g_{j,c}^x) \times pop_{2000}$. In columns (1)–(4)–(7), the sample is at the city-by-year level. In all other columns, the sample is at the city-by-(4 digit) industry-by-year level. Standard errors are clustered at the city level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

or the political affiliation of the mayor. We show in Appendix Table A5 that our results remain quantitatively the same when controlling only for local government expenditures (column 14), whether the mayor is affiliated to Lula’s party (column 15) or both at the same time (column 16). The inclusion of these controls imply that the effect of the reform is estimated by comparing cities that have similar political inclinations and welfare spending. We also show in Appendix Table A6 that results are similar when we include state-year fixed effects, implying that differences in state-level welfare programs or differences in political incentives at the state level cannot explain our results.

Finally, we provide additional evidence that our results are not driven by government programs by exploiting the fact that some of the largest government programs, like Bolsa Família, are distributed by a specific government bank: Caixa Econômica Federal. If our results were driven by Bolsa Família or other welfare programs, they would be strongest when treated cities are compared with control cities that did *not* have a branch of Caixa, as, in this case, treated cities would benefit from the welfare program expansion and control cities would not since, by construction, control cities do not have access to the distributor of the program. We report results of this exercise in Table A8 in Appendix. Unlike what we would expect if results were driven by government programs, we find that, if anything, results are weaker when no Caixa branches were present in control cities prior to the reform.

7.3. Sample composition

Our results on changes in inequality might be partially driven by a change in the worker composition in treated cities. Inequalities might increase for instance because following the reform, more low productivity workers enter the sample, pushing the mean wage of low-skill workers downward. We investigate this possibility in Table 12, in which we measure inequality using the city-level standard deviation of log wage.³² This allows us to measure wages as the residual of a Mincer equation including different worker characteristics. The inclusion of these characteristics is equivalent to holding fixed the sample composition along these dimensions.

In column 1, we report the result when we use the raw wage. In column 2, we add a third-order polynomial on age and fixed effects for sex and seven categories of race.³³ In column 3 we include 2-digit industry fixed effects and in column 4 we include 2-digit industry-by-2 digit occupation fixed effects (4,479 distinct dummies). Finally, in columns 5 and 6, we use the unfiltered wage, but restrict to the sample of workers present from 2004 to 2014 (column 5) and to firms present

prior to the reform (column 6) to estimate whether our effect are driven by a change in the entry/exit of workers or firms.³⁴ Across all the different level of controls, we find an overall stable effect of the reform, with higher financial development leading to more inequality.

7.4. Informality and exposure to commodity sector

Note that columns 5 and 6 of Table 12 show that our results are robust to restricting to workers and firms already in the formal sector, and thus suggest our findings are not driven by workers and firms moving into or out of the informal sector. We complement these results by controlling for the city-level employment in the informal sector from the 2000 Census, which we include as one of controls in column (6) of Table A5 in Appendix. This confirms once again that our results are not driven by the informal sector. These results are in line with the fact that treated and control cities have the same level of informality prior to the shock, as shown in the covariate balance test of Fig. 3.

We also directly test if exposure to the commodity sector could explain our results by controlling for employment in the commodity sector (column 9 of Table A5), or for the change in commodity prices post reform. We construct this variable as the weighted sum of prices across the fourteen main commodities in Brazil, similar to the measure developed by Benguria et al. (2023).³⁵ Our results remain quantitatively the same, consistent with the analysis of Table 11, where we non-parametrically control for sector-specific shocks.

7.5. Other robustness checks

As we discuss in Section 3, we conduct a number of other robustness checks relating to our matching procedure and empirical specification.

We show in Table A3 in the appendix that results are robust to using different numbers of control cities. In Table A4, we show results are not sensitive to the matching procedure. In panel A we replicate our results in the baseline sample. In panel B, we additionally exact match on quintiles of the share of skilled workers pre-reform. In panel C, we exact match on quintiles of the share of manufacturing pre-reform and, in panel D, we exact match on quintiles of the level of inequality pre-reform. In all cases, the point estimates of all the outcomes are quantitatively very similar.

In Table A5, we include a collection of additional city-level controls, such as GDP, employment, skilled employment, political affiliation of the mayor, trade, distance to the state capital, and the comovement of local GDP with aggregate fluctuations. These results help rule out the

³² We use the standard deviation instead of the Gini here because the Gini requires only positive values, but residualizing wages leads to potential negative values. By contrast, the standard deviation is always well defined.

³³ There are six race categories in RAIS: Indigenous, White, Black, Asian, multiracial, and not reported. We also include missing race values as a seventh category so as not to exclude those observations from this analysis.

³⁴ Results are similar when we require firms to be present throughout the period. We only condition on firms exiting pre-reform because the increase in firm exit post reform and workers losing their firm-specific human capital or firm-specific shared rent could be a channel through which financial development affects inequality.

³⁵ We would like to thank the authors for generously sharing their measure with us.

Table 12
Standard deviation of wages.

Dependent variable	Std[log(Wage)]					
Fixed effects	None	Age×Sex ×Race	Industry	Industry ×Occupation	Workers 2004–2014	Firms 2004
	(1)	(2)	(3)	(4)	(5)	(6)
Treated×Post	0.015*** (0.002)	0.014*** (0.002)	0.011*** (0.002)	0.010*** (0.002)	0.021*** (0.003)	0.013*** (0.002)
City FE	✓	✓	✓	✓	✓	✓
Match×Year FE	✓	✓	✓	✓	✓	✓
Observations	79,995	79,995	79,995	79,995	79,980	79,995

This table shows the effect of the reform on the change in the standard deviation of log (wage) at the city level. From columns 2 to 5, we use as the wage the residual of a Mincerian regression, after we have filtered a polynomial of age (age, age-square, age-cube) and fixed effects for gender and seven race categories (column 2), added 2-digit industry fixed effects (column 3), and 2-digit industries×2 digit occupation fixed effects (column 4)). In columns (5) and (6), we use the unfiltered wage, but restrict to the sample of workers present from 2004 to 2014 (column 5) and to firms present prior to the reform (column 6). Standard errors are clustered at the city level. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

possibility that our results are driven by differential exposure of treated cities to aggregate shocks or by political connections. In particular, including for state-by-year fixed effects implies that we control for any state-level political cycles, and controlling for the political affiliation of the mayor implies that we estimate our effects by comparing cities where mayors have the same political affiliation.

Given that the reform may have a direct impact on many city characteristics, we control for the pre-reform value of these characteristics interacted with year fixed effects. Finally, we also show in Table A6 in the Appendix that results are robust to adding state-by-year fixed effects to control for time-varying unobserved variation across regions of Brazil.

Finally, in Appendix A.3 we adopt an alternative identification strategy and use the actual entry of a public bank branch to identify treated cities. Specifically, we consider that a city is treated on the first year that we observe a public bank branch in the city. In this case, we correct for the staggered nature of the design using the estimators developed by [De Chaisemartin and D'Haultfoeuille \(2020\)](#), [Sun and Abraham \(2021\)](#), and [Borusyak et al. \(2024\)](#), which can accommodate high dimensional fixed effects. Across all outcomes, we find that the point estimate of our baseline specification is usually slightly smaller.

8. Conclusion

In this paper, we show that the expansion of financial access and capital deepening promoted by the Brazilian government led to a permanent increase in economic development, driven both by an expansion of existing businesses and an increase in firm creation, which accelerated the exit of existing firms. These effects materialize over time, underlying the need to study a long-enough period to capture the true effect of one-time reforms on long-run development.

This important economic development triggered a substantial rise in wage inequality, which is mostly explained by the limited supply of skilled labor in some cities. This result raises the question of whether governments should also implement simultaneous labor-oriented policies in order to reap the full benefit of formal financial market policies.

The importance of financial inclusion also has potential implications for current and future policy on digital banking. Such policies are already underway in some developing countries with the goal of expanding financial access, including in Brazil with the launch of an instant payment platform (Pix) and its mandatory use by all financial institutions and payment institutions that are licensed by the Central Bank of Brazil. This initiative, by improving deposit provisions and lowering costs of access to liquidity services, should boost development and firm growth, particularly in tradable industries. However, whether this initiative alone will achieve its intended goals or if, instead, it would be complementary with the expansion of physical bank branches remains an open debate. It might well be the case that the expected

gains from Pix will materialize mostly in places where households and businesses are already banked if the initial connection to the formal finance sector requires the physical presence of a bank branch.

Our results also imply that policies such as digital banking that increase financial inclusion for retail customers and for small and medium-sized enterprises could be a source of substantial increase in inequality in the future if they interact with other frictions, such as the limited supply of human capital, that are prevalent in many developing countries.

CRedit authorship contribution statement

Julia Fonseca: Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Adrien Matray:** Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Philipp Schnabl was the editor for this article. We thank Victor Duarte and Chenzi Xu for their unfailing support and numerous discussions. We also thank Daniel Carvalho (discussant), Shawn Cole (discussant), Tatyana Deryugina, Pascaline Dupas, Melanie Morten, Chad Jones, Dean Karlan (discussant), Ben Moll, Tarun Ramadorai, Rebecca de Simone (discussant), Amit Seru, Yongseok Shin, Tracy Wang (discussant), and seminar participants at Stanford Economics, Stanford GSB, SITE Financial Regulation, the Finance-Organization-and-Markets Conference at Dartmouth, WashU Olin, Boston College, Queen Mary University, University of Pittsburgh, FGV-EPGE, Insper, University of São Paulo, HEC-Paris, the Brazilian Econometric Society Seminar, Cheung Kong University GSB, the Bank of Lithuania, University of Georgia, University of Arizona, Imperial College, Australian National University, the Women in Applied Microeconomics conference, the Wabash River Finance conference, the Central Bank of Brazil, NBER-SI, Central Bank of Italy, Tuck Dartmouth, McDonough-Georgetown, FED-board, John

Hopkins Carey business school and Berkeley University for helpful comments. Filipe Correia, Thomás Gleizer Feibert and Pablo Enrique Rodriguez provided superb research assistance.

References

- Abowd, J., Kramarz, F., Margolis, D., 1999. High wage workers and high wage firms. *Econometrica* 67 (2), 251–333.
- Angrist, J., Pischke, J.-S., 2008. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Anon, 2018. In: Beck, T., Levine, R. (Eds.), *Handbook of Finance and Development*. Edward Elgar Publishing, Northampton, MA, xvi, 597 pages.
- Baghai, R., Silva, R., Thell, V., Vig, V., 2021. Talent in distressed firms: investigating the labor costs of financial distress. *J. Finance* 76 (6), 2907–2961.
- Bai, J., Carvalho, D., Phillips, G., 2018. The impact of bank credit on labor reallocation and aggregate industry productivity. *J. Finance* 73 (6), 2787–2836.
- Baker, A.C., Larcker, D.F., Wang, C.C., 2022. How much should we trust staggered difference-in-differences estimates? *J. Financ. Econ.* 144 (2), 370–395.
- Banerjee, A., Duflo, E., 2014. Do firms want to borrow more? Testing credit constraints using a directed lending program. *Rev. Econ. Stud.* 81 (2), 572–607.
- Banerjee, A., Karlan, D., Zinman, J., 2015. Six randomized evaluations of microcredit: introduction and further steps. *Am. Econ. J. Appl. Econ.* 7 (1), 1–21.
- Banerjee, A., Moll, B., 2010. Why does misallocation persist? *Am. Econ. J.: Macroecon.* 2 (1), 189–206.
- Barboni, G., Field, E., Pande, R., 2023. Rural banks can reduce poverty: evidence from 870 Indian villages. Working Paper.
- Bau, N., Matray, A., 2023. Misallocation and capital market integration: evidence from India. *Econometrica* 91 (1), 67–106.
- Bazzi, S., de Freitas Oliveira, R., Muendler, M.-A., Rauch, J., 2023. Credit supply shocks and firm dynamics: evidence from Brazil. NBER Working Paper.
- Beck, T., Levine, R., Levkov, A., 2010. Big bad banks? The winners and losers from bank deregulation in the United States. *J. Finance* 65 (5), 1637–1667.
- Benguria, F., Saffie, F., Urzúa, S., 2023. The transmission of commodity price super-cycles. *Rev. Econ. Stud.* rdad078.
- Benmelech, E., Bergman, N., Seru, A., 2021. Financing labor. *Rev. Finance* 25 (5), 1365–1393.
- Berton, F., Mocetti, S., Presbitero, A., Richiardi, M., 2018. Banks, firms, and jobs. *Rev. Financ. Stud.* 31 (6), 2113–2156.
- Bittencourt, G., Magalhães, R., Abramovay, R., 2005. Informação de crédito: um meio para ampliar o acesso dos mais pobres ao sistema financeiro. *Pesquisa Debate* 16 (28), 203–248.
- Bombardini, M., Orefice, G., Tito, M., 2019. Does exporting improve matching? Evidence from french employer-employee data. *J. Int. Econ.* 117, 229–241.
- Borusyak, K., Jaravel, X., Spiess, J., 2024. Revisiting event study designs: robust and efficient estimation. *Rev. Econ. Stud.* (Forthcoming).
- Bruhn, M., Love, I., 2014. The real impact of improved access to finance: evidence from Mexico. *J. Finance* 69 (3), 1347–1376.
- Buera, F., Kaboski, J., Shin, Y., 2011. Finance and development: a tale of two sectors. *Amer. Econ. Rev.* 101 (5), 1964–2002.
- Buera, F., Kaboski, J., Shin, Y., 2015. Entrepreneurship and financial frictions: a macrodevelopment perspective. *Annu. Rev. Econ.* 7 (1), 409–436.
- Buera, F., Kaboski, J., Shin, Y., 2021. The macroeconomics of microfinance. *Rev. Econ. Stud.* 88 (1), 126–161.
- Buera, F., Shin, Y., 2013. Financial frictions and the persistence of history: a quantitative exploration. *J. Polit. Econ.* 121 (2), 221–272.
- Burgess, R., Pande, R., 2005. Do rural banks matter? Evidence from the Indian social banking experiment. *Amer. Econ. Rev.* 95 (3), 780–795.
- Caggese, A., Cunat, V., Metzger, D., 2019. Firing the wrong workers: financing constraints and labor misallocation. *J. Financ. Econ.* 133 (3), 589–607.
- Carvalho, D., 2014. The real effects of government-owned banks: evidence from an emerging market. *J. Finance* 69 (2), 577–609.
- Chodorow-Reich, G., 2014. The employment effects of credit market disruptions: firm-level evidence from the 2008–9 financial crises. *Q. J. Econ.* 129 (1), 1–59.
- Choudhary, A., Limodio, N., 2022. Liquidity risk and long-term finance: evidence from a natural experiment. *Rev. Econ. Stud.* 89 (3), 1278–1313.
- Coelho, C., de Mello, J., Rezende, L., 2013. Do public banks compete with private banks? Evidence from concentrated local markets in Brazil. *J. Money Credit Bank.* 45 (8), 1581–1615.
- Cole, S., 2009. Fixing market failures or fixing elections? Agricultural credit in India. *Am. Econ. J. Appl. Econ.* 1 (1), 219–250.
- Cramer, K.F., 2022. Bank presence and health. Working Paper.
- Crescenzi, R., Limodio, N., 2022. The impact of Chinese FDI in africa: evidence from ethiopia.
- De Chaisemartin, C., D'Haultfoeulle, X., 2020. Two-way fixed effects estimators with heterogeneous treatment effects. *Amer. Econ. Rev.* 110 (9), 2964–2996.
- De Chaisemartin, C., Ramirez-Cuellar, J., 2024. At what level should one cluster standard errors in paired and small-strata experiments? *Am. Econ. J. Appl. Econ.* 16 (1), 193–212.
- Delatte, A.L., Matray, A., Pinardon Touati, N., 2022. Political quid pro quo in financial markets. Working Paper.
- Ding, S., 2005. Politicians and banks: political influences on government-owned banks in emerging markets. *J. Financ. Econ.* 77 (2), 453–479.
- Dix-Carneiro, R., Goldberg, P., Meghir, C., Ulyssea, G., 2021. Trade and informality in the presence of labor market frictions and regulations. NBER Working Paper.
- Doornik, B.F.N.V., Gomes, A., Schoenherr, D., Skrastins, J., 2021. Financial access and labor market outcomes: evidence from credit lotteries. Working Paper.
- Eeckhout, J., Kircher, P., 2011. Identifying sorting—in theory. *Rev. Econ. Stud.* 78 (3), 872–906.
- Fonseca, J., Doornik, B.V., 2022. Financial development and labor market outcomes: evidence from Brazil. *J. Financ. Econ.* 143 (1), 550–568.
- Giné, X., Townsend, R., 2004. Evaluation of financial liberalization: a general equilibrium model with constrained occupation choice. *J. Dev. Econ.* 74 (2), 269–307.
- Greenstone, M., Mas, A., Nguyen, H.-L., 2020. Do credit market shocks affect the real economy? Quasi-experimental evidence from the great recession and “normal” economic times. *Am. Econ. J. Econ. Policy* 12 (1), 200–225.
- Greenwood, J., Jovanovic, B., 1990. Financial development, growth, and the distribution of income. *J. Polit. Econ.* 98 (5), 1076–1107.
- Greenwood, J., Sanchez, J., Wang, C., 2010. Financing development: the role of information costs. *Amer. Econ. Rev.* 100 (4), 1875–1891.
- Gual, L.B., Anson, J., 2008. Financial access and inclusion through postal networks: evaluating the experience of Brazil's banco postal. In: Anson, J., Toledano, J. (Eds.), *Postal Economics in Developing Countries*. Universal Postal Union, Berne, pp. 139–174.
- Hausmann, R., Rodrik, D., 2003. Economic development as self-discovery. *J. Dev. Econ.* 72 (2), 603–633.
- Hombert, J., Matray, A., 2017. The real effects of lending relationships on innovative firms and inventor mobility. *Rev. Financ. Stud.* 30 (7), 2413–2445.
- Imbens, G., Rubin, D., 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, Cambridge.
- Imbs, J., Wacziarg, R., 2003. Stages of diversification. *Amer. Econ. Rev.* 93 (1), 63–86.
- Jasova, M., Mendicino, C., Panetti, E., Peydró, J.-L., Supera, D., 2021. Monetary policy, labor income redistribution and the credit channel: evidence from matched employer-employee and credit registers. Working Paper.
- Ji, Y., Teng, S., Townsend, R.M., 2023. Dynamic bank expansion: spatial growth, financial access, and inequality. *J. Polit. Econ.* 131 (8), 2209–2275.
- Kaboski, J., Townsend, R., 2011. A structural evaluation of a large-scale quasi-experimental microfinance initiative. *Econometrica* 79 (5), 1357–1406.
- Kaboski, J., Townsend, R., 2012. The impact of credit on village economies. *Am. Econ. J. Appl. Econ.* 4 (2), 98–133.
- Kochar, A., 2011. The distributive consequences of social banking: a microempirical analysis of the Indian experience. *Econom. Dev. Cult. Chang.* 59 (2), 251–280.
- Lopes de Melo, R., 2018. Firm wage differentials and labor market sorting: reconciling theory and evidence. *J. Polit. Econ.* 126 (1), 313–346.
- Lopez, H., Perry, G., 2008. *Inequality In Latin America : Determinants And Consequences*. The World Bank, Policy Research Working Papers.
- Loureiro, E.R., Madeira, G.d., Bader, F.L.C., 2016. Expansão dos correspondentes bancários no brasil: uma análise empírica. Central Bank of Brazil Working Paper Series Working Paper, No. 433.
- Mestieri, M., Schauer, J., Townsend, R., 2017. Human capital acquisition and occupational choice: implications for economic development. *Rev. Econ. Dyn.* 25, 151–186.
- Mettenheim, K., 2010. *Federal Banking in Brazil: Policies and Competitive Advantages*. Routledge, London.
- Midrigan, V., Xu, D.Y., 2014. Finance and misallocation: evidence from plant-level data. *Amer. Econ. Rev.* 104 (2), 422–458.
- Ministério da Fazenda, 2007. *Plano Plurianual 2004–2007: Relatório Anual de Avaliação*. Technical Report.
- Moll, B., Townsend, R., Zhorin, V., 2017. Economic development, flow of funds, and the equilibrium interaction of financial frictions. *Proc. Natl. Acad. Sci. USA* 114 (24), 6176–6184.
- Ornelas, J.R.H., Pedraza, A., Ruiz-Ortega, C., Silva, T.C., 2021. Credit allocation when private banks distribute government loans. Working Paper.
- Peek, J., Rosengren, E., 2000. Collateral damage: effects of the Japanese bank crisis on real economic activity in the United States. *Amer. Econ. Rev.* 90 (1), 30–45.
- Petersen, M., Rajan, R., 1994. The benefits of lending relationships: evidence from small business data. *J. Finance* 49 (1), 3–37.
- Porcher, C., 2022. Migration with costly information. Working Paper.
- Quincy, S., 2023. “Loans for the little fellow.” credit, crisis, and recovery in the great depression.
- Rajan, R., Zingales, L., 2001. Financial systems, industrial structure, and growth. *Oxford Rev. Econ. Policy* 17 (4), 467–482.

- Sanford, C., 2013. Do agents improve financial inclusion? Evidence from a national survey in Brazil. Technical Report, Bankable Frontier Associates.
- Sapienza, P., 2004. The effects of government ownership on bank lending. *J. Financ. Econ.* 72 (2), 357–384.
- Schoefer, B., 2021. The financial channel of wage rigidity. NBER Working Paper.
- Sun, L., Abraham, S., 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J. Econometrics* 225 (2), 175–199.
- Townsend, R., Ueda, K., 2006. Financial deepening, inequality, and growth: a model-based quantitative evaluation. *Rev. Econ. Stud.* 73 (1), 251–280.
- Vig, V., 2013. Access to collateral and corporate debt structure: evidence from a natural experiment. *J. Finance* 68 (3), 881–928.
- Xu, C., 2022. Reshaping global trade: the immediate and long-run effects of bank failures. *Q. J. Econ.* 137 (4), 2107–2161.
- Xu, C., Yang, H., 2022. Real effects of supplying safe private money. NBER Working Paper.