

On Fintech and Financial Inclusion

Thomas Philippon*

September 2019

Abstract

The cost of financial intermediation has declined in recent years thanks to technological progress and increased competition. I document this fact and I analyze two features of new financial technologies that have stirred controversy: returns to scale, and the use of big data and machine learning. I argue that the nature of fixed versus variable costs in robo-advising is likely to democratize access to financial services. Big data is likely to reduce the impact of negative prejudice in the credit market but it could reduce the effectiveness of existing policies aimed at protecting minorities.

JEL: E2, G2, N2

*Stern School of Business, New York University; NBER and CEPR. This paper was prepared for the 2019 BIS Annual Research Conference. I am grateful to my discussants Manju Puri and David Dorn, to Hyun Shin, Marina Niessner, and participants at the 2019 BIS Annual Research Conference. I thank Marcos Sonnervig for outstanding research assistance.

Fintech covers digital innovations and technology-enabled business model innovations in the financial sector. Such innovations can disrupt existing industry structures and blur industry boundaries, facilitate strategic disintermediation, change how existing firms create and deliver products and services, provide new gateways for entrepreneurship, and democratize access to financial services. On the other hand, they create significant privacy, regulatory and law-enforcement challenges and they could increase the scope for some forms of discrimination. Examples of innovations that are central to Fintech today include various application of blockchain technologies, new digital advisory and trading systems, artificial intelligence and machine learning, peer-to-peer lending, equity crowdfunding and mobile payment systems. In this paper I offer some preliminary evidence and theoretical analysis about the impact of technological progress in the finance industry.

The first question is whether there has been any material change in financial intermediation in recent years. To shed some light on this question, I update the work of Philippon (2015) with post-crisis U.S. data. The puzzle emphasized in previous work was that the unit cost of financial intermediation had remained stubbornly close to 200 basis points for more than a century, despite advances – and large investments – in computers and communication technologies. The post-crisis data suggests that this puzzle might be diminishing. I find that the unit cost of financial intermediation has declined over the past 10 years.

I then study two issues that are at the heart of the Fintech debate: access to finance and discrimination. If we accept the fact that Fintech brings efficiency gains to financial intermediation, the next question is: how will these gains be shared? Will Fintech democratize access to financial services or will it increase inequality?

I highlight two forces that will shape the answer to these questions. The first is increasing returns to scale brought by technology. I illustrate how this force plays out in the context of asset management and robo-advising. Investors select advisors and advisors help them choose their investments (Pedersen, 2015). Search costs imply that wealthy households have access to better advice. I argue that robo-advising and related technologies will change the nature of fixed costs in a way that is likely to improve access to financial services. It may not, however, reduce inequality among all groups.

The second force is the use of big data and machine learning (BDML for short). I illustrate how this force plays out in the context of consumer credit. I argue that this technology is likely to reduce unwarranted human biases against minorities, but it will probably decrease the effectiveness of existing regulations. The tentative conclusion is that Fintech can bring widely-shared welfare benefits but changes in existing policies and regulations are necessary to achieve its full potential.

Recent literature Philippon (2016) discusses the literature up to 2016 so I will mention here some recent papers. Petralia et al. (2019) discuss the impact of new technologies on the banking industry. Focusing on residential mortgages, Buchak et al. (2018) study the growth in the market share of shadow bank and Fintech lenders, arguing that it can be explained by differences in regulation and technological advantages. They find that Fintech lenders

serve more creditworthy borrowers (relative to shadow banks) but charge higher interest rates (14-16 basis points), which is consistent with the idea that consumers are willing to pay for better user experience and quick decisions. Fuster et al. (2019) study the differences between Fintech and traditional lenders in the mortgage market and find that the former is quicker in processing applications (20% faster), without increasing loan risk. They also provide evidence that Fintech lenders adjust supply more elastically to demand shocks and increase the propensity to refinance, especially among borrowers that are likely to benefit from it. Their results suggest that Fintech firms have improved the efficiency of financial intermediation in mortgage markets.

The advent of Fintech is often seen as a promising avenue for reducing inequality in access to credit. Bartlett et al. (2018) study this issue, analyzing the role of Fintech lenders in alleviating discrimination in mortgage markets. They find that all lenders, including Fintech, charge minorities more for purchase and refinance mortgages but that Fintech algorithms discriminate 40% less than face-to-face lenders. Regarding the use of new technologies in credit markets, Berg et al. (2019) analyse the information content of the “digital footprint” (an easily accessible information for any firm conducting business in the digital sphere) for predicting consumer default. Berg et al. (2019) find that the predictive power of these new data equals or exceeds that of traditional credit bureau scores. Their results suggest that new technologies and new data might bring a superior ability for screening borrowers.

Fintech firms are also competing in the market for wealth management. The United States is the leading market for robo-advisors. In 2017, it accounted for more than half of all investments in robo-advisors (Abraham et al., 2019). Nevertheless, the amount of assets managed by robo-advisors is still a small portion of total assets under management, with average client wealth much smaller than the average in the industry (Economist, 2017). Abraham et al. (2019) argues that because they save on fixed costs (such as salaries of financial advisors or maintenance of physical offices), robo-advisors can reduce minimum investment requirements and charge lower fees. Finally, one area that is not well developed is the study of the interaction between Fintech and incentives within organizations. This is likely to be relevant for discrimination and other biases. For instance, Dobbie et al. (2018) show that the misalignment of incentives inside financial firms can lead to biased lending decisions.

1 (In)efficiency of the Existing System

The main finding in Philippon (2015) is that the unit cost of financial intermediation in the U.S. has remained around 200 basis points for the past 130 years. Improvements in information technologies have not been passed through to the end users of financial services. This section offers an update of this work.

1.1 Financial Expenses and Intermediated Assets

To organize the discussion I use a simple model economy consisting of households, a non-financial business sector, and a financial intermediation sector. The details of the model are in the Appendix. The income share of finance,

shown in Figure 1, is defined as¹

$$\frac{y_t^f}{y_t} = \frac{\text{Value Added of Finance Industry}}{\text{GDP}}.$$

The model assumes that financial services are produced under constant returns to scale. The income of the finance industry y_t^f is then given by

$$y_t^f = \psi_{c,t} b_{c,t} + \psi_{m,t} m_t + \psi_{k,t} k_t, \quad (1)$$

where $b_{c,t}$ is consumer credit, m_t are assets providing liquidity services, and k_t is the value of intermediated corporate assets. The parameters $\psi_{i,t}$'s are the unit cost of intermediation, pinned down by the intermediation technology. The model therefore says that the income of the finance industry is proportional to the quantity of intermediated assets, properly defined. The model predicts no income effect, i.e., no tendency for the finance income share to grow with per-capita GDP. This does not mean that the finance income share should be constant, since the ratio of assets to GDP can change. But it says that the income share does not grow mechanically with total factor productivity. This is consistent with the historical evidence.²

Measuring intermediated assets is complicated because these assets are heterogenous. As far as corporate finance is concerned, the model is fundamentally a user cost model. Improvements in corporate finance (a decrease in ψ_k) lower the user cost of capital and increase the capital stock, which, from a theoretical perspective, should include all intangible investments and should be measured at market value. A significant part of the growth of the finance industry over the past 30 years is linked to household credit. The model provides a simple way to model household finance. The model also incorporates liquidity services provided by specific liabilities (deposits, checking accounts, some form of repurchase agreements) issued by financial intermediaries. One can always write the RHS of (1) as $\psi_{c,t} \left(b_{c,t} + \frac{\psi_{m,t}}{\psi_{c,t}} m_t + \frac{\psi_{k,t}}{\psi_{c,t}} k_t \right)$. Philippon (2015) finds that the ratios $\frac{\psi_{m,t}}{\psi_{c,t}}$ and $\frac{\psi_{k,t}}{\psi_{c,t}}$ are close to one.³ As a result one can define intermediated assets as

$$q_t \equiv b_{c,t} + m_t + k_t. \quad (2)$$

The principle is to measure the instruments on the balance sheets of non-financial users, households and non-financial firms. This is the correct way to do the accounting, rather than looking at the balance sheet of financial


¹Philippon (2015) discusses various issues of measurement. Conceptually, the best measure is value added, which is the sum of profits and wages. Whenever possible, I therefore use the GDP share of the finance industry, i.e., the nominal value added of the finance industry divided by the nominal GDP of the U.S. economy. One issue, however, is that before 1945 profits are not always properly measured and value added is not available. As an alternative measure I then use the labor compensation share of the finance industry, i.e., the compensation of all employees of the finance industry divided by the compensation of all employees in the U.S. economy. Philippon (2015) also explains the robustness of the main findings to large changes in government spending (because of wars), the rise of services (finance as a share of services displays a similar pattern to the one presented here), globalization (netting out imports and exports of financial services).

²The fact that the finance share of GDP is the same in 1925 and in 1980 makes it already clear that there is no mechanical relationship between GDP per capita and the finance income share. Similarly, Bickenbach et al. (2009) show that the income share of finance has remained remarkably constant in Germany over the past 30 years. More precisely, using KLEMS for Europe (see O'Mahony and Timmer (2009)) one can see that the finance share in Germany was 4.3% in 1980, 4.68% in 1990, 4.19% in 2000, and 4.47% in 2006.

³This is true most of the time, but not when quality adjustments are too large. Philippon (2015) provides calibrated quality adjustments for the U.S. financial system.

intermediaries. After aggregating the various types of credit, equity issuances and liquid assets into one measure, I obtain the quantity of financial assets intermediated by the financial sector for the non-financial sector, displayed in Figure 1.

Figure 1: Finance Income and Intermediated Assets



Notes: Both series are expressed as a share of GDP. Finance Income is the domestic income of the finance and insurance industries, i.e., aggregate income minus net exports. Intermediated Assets include debt and equity issued by non financial firms, household debt, and various assets providing liquidity services. Data range for Intermediated Assets is 1886 - 2012. See Philippon (2015) for historical sources and details about the underlying data.

1.2 Unit Cost and Quality Adjustments


I can then divide the income of the finance industry by the quantity of intermediated assets to obtain a measure of the unit cost

$$\psi_t \equiv \frac{y_t^f}{q_t}. \tag{3}$$

Figure 2 shows that this unit cost is around 2% and relatively stable over time. In other words, I estimate that it costs two cents per year to create and maintain one dollar of intermediated financial asset. Equivalently, the annual rate of return of savers is on average 2 percentage points below the funding cost of borrowers. The updated series are similar to the ones in the original paper. The unit costs for other countries are estimated by Bazot (2013) who finds convergence to US levels.

The raw measure of Figure 2, however, does not take into account changes in the characteristics of borrowers. These changes require quality adjustments to the raw measure of intermediated assets. For instance, corporate finance involves issuing commercial paper for blue chip companies as well as raising equity for high-technology start-ups. The monitoring requirements per dollar intermediated are clearly different in these two activities. Similarly, with household finance, it is more expensive to lend to poor households than to wealthy ones, and relatively poor

Figure 2: Unit Cost of Financial Intermediation




Notes: The raw measure is the ratio of finance income to intermediated assets, displayed in Figure 1. The 2012 data is from Philippon (2015), while the new data was accessed in May 2016. Data range is 1886 - 2015.

households have gained access to credit in recent years.⁴ Measurement problems arise when the mix of high- and low-quality borrowers changes over time.

Following Philippon (2015), I then perform a quality adjustment to the intermediated assets series. Figure 3 shows the quality adjusted unit cost series. It is lower than the unadjusted series by construction since quality adjusted assets are (weakly) larger than raw intermediated assets. The gap between the two series grows when there is entry of new firms, and/or when there is credit expansion at the extensive margin (i.e., new borrowers). Even with the adjusted series, however, we do not see a significant decrease in the unit cost of intermediation over time.

⁴Using the Survey of Consumer Finances, Moore and Palumbo (2010) document that between 1989 and 2007 the fraction of households with positive debt balances increases from 72% to 77%. This increase is concentrated at the bottom of the income distribution. For households in the 0-40 percentiles of income, the fraction with some debt outstanding goes from 53% to 61% between 1989 and 2007. In the mortgage market, Mayer and Pence (2008) show that subprime originations account for 15% to 20% of all HMDA originations in 2005.

Figure 3: Unit Cost and Quality Adjustment



Notes: The quality adjusted measure takes into account changes in firms' and households' characteristics. Data range is 1886 - 2015.

As I have argued in the past, the puzzle is why we have not seen substantial productivity gains in financial intermediation. The good news is that, however late, these improvements might be happening now.

2 A Simple Model of Robo-Advising

I consider a simple model of imperfect competition in asset management services. The model emphasizes the role of technology, and fixed costs in particular. The key point is that there are two types of fixed costs: fixed costs to set up a business or a system or a platform; and then fixed cost per relationship with each client.

There is a continuum of mass 1 of households whose wealth w is distributed according to the (cumulative) distribution $G(w)$. Households are risk neutral (or, equivalently, returns are risk adjusted) and have access to an investment technology with gross return r . The reservation utility of a household is thus rw . Households also have the option to hire an asset manager in order to earn higher returns.

2.1 Traditional Asset Management Equilibrium

There are N asset managers with access to an investment technology with return $R > r$. To be active they need to pay the fixed cost Φ (per active firm). To work with a household they need to pay the relationship cost ϕ (per client). The asset management industry is oligopolistic and asset managers charge a fee $f(w)$ for their services. I restrict attention to linear fees of the form $f(w) = \phi + \mu w$, where the intercept covers the fixed cost and where μ is a markup. For now the markup is simply a parameter but, as discussed later, it could be a decreasing function of the number of active intermediaries. The parameters are such that $\mu < R - r$.

The key decision for a household is whether to hire an asset manager or not. Household w chooses to hire an

asset manager if $Rw - f(w) > rw$, which happens if and only if

$$w > w_o = \frac{\phi}{R - r - \mu}. \quad (4)$$

A fraction $1 - G(w_o)$ of households hire intermediation services, while the remaining fraction $G(w_o)$ invest by themselves at the low rate. I consider a symmetric equilibrium where intermediaries have the same number of clients. The net profit of any intermediary is therefore

$$\pi(N) \equiv \frac{\mu}{N} \int_{w_o}^{\infty} w dG(w),$$

where w_o is defined in equation (4). Note that N should be interpreted as the number of asset managers per capita since the population is normalized to one. Finally, free entry requires

$$\pi(N) \geq \Phi,$$

with equality if entry is positive. This pins down the number of firms entering the market.

Definition 1. Given the cost structure (Φ, ϕ) , an equilibrium with positive entry solves w_o from equation (4) and

$$N = \frac{\mu}{\Phi} \int_{w_o}^{\infty} w dG(w).$$

Welfare is given by the following expression:

$$\begin{aligned} \mathcal{W} &= \int_0^{w_o} rwdG(w) + \int_{w_o}^{\infty} (Rw - \phi) dG(w) - N\Phi \\ &= \int_0^{w_o} rwdG(w) + \int_{w_o}^{\infty} ((R - \mu)w - \phi) dG(w) \end{aligned}$$

Let us briefly discuss the first best allocation. The planner's solution would set $N = 1$ to save on fixed entry costs, and $\mu = 0$ to price at marginal cost. This implies $w_o^* = \frac{\phi}{R-r}$. In this equilibrium all asset managers loose money so the planner would need to impose lump-sum taxes on households to subsidize the financial intermediaries. Asset management clearly improves welfare compared to a situation where all households earn the low rate r , but, as expected, the decentralized equilibrium does not achieve the planner's outcome because asset managers need to charge a markup to cover their entry cost, and $w_o > w_o^*$.

2.2 Robo-Advisors: A Tale of Two Fixed Costs

Let us now introduce a new asset management technology, characterized by stronger returns to scale. More precisely, we assume that robo advisors have access to the investment technology $r < R' \leq R$, and have a higher fixed entry cost $\Phi' > \Phi$ but a lower cost per client $\phi' < \phi$. This captures the core idea that it is costly to program the robo ($\Phi' > \Phi$), but then it can easily manage a large number of clients ($\phi' < \phi$). In the limit, we can even imagine that $\phi' = 0$. The existing literature has failed to capture the difference between these two types of fixed costs. I will show that they have different welfare implications.

Households now have three options. They can invest by themselves and earn r . They can hire a traditional manager. Or they can hire a robo advisor. There are now two cutoffs to calculate w_1 and w_2 , one for the participation decision (as before) and one for the choice of the type of asset manager. I assume that both types of asset managers charge a markup μ , so robo advisory fees are $\phi' + \mu w$ while traditional management fees are $\phi + \mu w$. Since $R' \leq R$ and $\phi' < \phi$ we know that relatively poor households will choose between autarky and robots. Hence the first participation cutoff is

$$w_1 = \frac{\phi'}{R' - r - \mu}. \quad (5)$$

The second cutoff is between the robo and traditional advisors: $(R' - \mu)w_2 - \phi' = (R - \mu)w_2 - \phi$. Hence

$$w_2 = \frac{\phi - \phi'}{R - R'}. \quad (6)$$

One can relatively easily accommodate $R' = R$ if we introduce horizontal or vertical differentiation between advisors as in Pagnotta and Philippon (2018), or search costs between households and advisors as in Pedersen (2015) and Garleanu and Pedersen (2018). For simplicity I consider here the case $R' < R$ that ensures an interior solution without getting lost in the details of strategic interactions between various types of advisors. This could also capture a convenience yield that investor perceive from interaction with a human being instead of a robo. Finally, it is consistent with the stylized fact that very high net worth households still use mostly human advisors. The condition for profitable entry by robots is $w_1 < w_2$, which leads to the following Lemma.

Lemma 1. *Robo advisors are active if and only if $\phi R' > \phi' R + (\phi - \phi')(r + \mu)$.*

We can now characterize the equilibrium when both types of advisors are active.

Proposition 1. *Under the condition of Lemma 1, the equilibrium with robo advisors is characterized by the cutoffs (w_1, w_2) in equations (5,6) such that $G(w_1)$ poor households save by themselves, $G(w_2) - G(w_1)$ middle-class households hire robo advisors, and $1 - G(w_2)$ rich households hire traditional managers. The number of robo advisors is*

$$N_0 = \frac{\mu}{\Phi'} \int_{w_1}^{w_2} w dG(w),$$

and the number of traditional managers is

$$N_1 = \frac{\mu}{\Phi} \int_{w_2}^{\infty} w dG(w).$$

Let us now analyse the welfare implication of robo advisors. Welfare is given by

$$\begin{aligned} \mathcal{W}' &= \int_0^{w_1} rwdG(w) + \int_{w_1}^{w_2} (R'w - \phi') dG(w) - N_0\Phi' + \int_{w_2}^{\infty} (Rw - \phi) dG(w) - N_1\Phi, \\ &= \int_0^{w_0} rwdG(w) + \int_{w_1}^{w_2} ((R' - \mu)w - \phi') dG(w) + \int_{w_2}^{\infty} ((R - \mu)w - \phi) dG(w). \end{aligned}$$

Compared to the planner's allocation the participation cutoff w_1 is still distorted by the markup μ exactly as in the traditional equilibrium. The second cutoff w_2 , however, is the same one that the planner would choose as one can readily see from equation (6). There is thus no distortion at the robo/traditional advisor frontier.

The following proposition establishes the key result that if robo entry is profitable, i.e., if we see any robo advising at all, then there are more households using asset management services in the Fintech equilibrium than in the traditional equilibrium.

Proposition 2. *Under the condition of Lemma 1, robo advising improves access to asset management services, i.e., $w_1 < w_0$.*

Proof. We need to show that $w_1 = \frac{\phi'}{R' - r - \mu} < w_0 = \frac{\phi}{R - r - \mu}$. This is true as long as $\phi'R + (\phi - \phi')(r + \mu) < \phi R'$ which is exactly the condition in Lemma 1. \square

The key take-away from this analysis is that, by lowering the fixed cost per relationship, Fintech allows more households to benefit from advisory services. The fact that Fintech might require higher upfront cost $\Phi' > \Phi$ does not matter because in equilibrium, the rich subsidize the poor. The rich pay the lion's share of fees that serve to cover the fixed cost Φ' of setting up the robo. Once this cost is paid, poor households benefit from cheaper services. An important lesson here is that the nature of fixed costs matter a great deal for welfare. The welfare properties of fixed "coding" costs are fundamentally different from those of fixed costs per client.

The fact that Fintech increases participation does not mean that Fintech reduces inequality among all groups, however. If we think that the cutoff w_0 used to be above the middle class and that the cutoff w_1 is now below the middle class, then inequality between the poor and the middle class increases as the middle class joins the rich in having access to asset management services. Education may also become a stronger determinant of participation than before if education allows households to understand better the important features of the robo-advising technology.

3 Big Data and Discrimination

Let me now discuss the consequences of the increasing use of “non-traditional” data in consumer credit. For most of the post-war era, banks have relied on the credit history of consumers to make lending decisions. Exactly how these borrowing histories are compiled can differ from country to country, but the broad principles are similar. In the U.S. for example, lenders (banks, mortgage lenders, credit card companies, financing companies) feed data on their customers to credit-reporting firms (credit bureaus), such as Equifax, Experian, and TransUnion. These firms aggregate the data and create files (credit reports) on consumers. The same firms (e.g. Experian or Equifax), or others, can then create credit scores. In the U.S., another company, Fair Isaac Corporation, or FICO, has developed the most widely-used score, the FICO score, which ranges from 300 to 850. Some lenders base their credit decisions almost entirely on FICO score, while others also consider more detailed credit bureau reports.

In recent years credit companies have started using new sources of data on consumer behavior, such as phone bills, shopping histories, subscriptions, or browsing histories. These data are less structured than traditional ones and new lenders often rely on machine learning algorithms to exploit them effectively. I will therefore refer to the whole process as big data and machine learning, or BDML for short. BDML is useful since many potential borrowers are not properly scored by traditional methods. In the U.S. for example about 50 million adults do not have a FICO score, and another 50 million have subprime scores that are notoriously noisy. One can hope that BDML will help bring credit to deserving borrowers who lack reliable scores. On the other hand, regulators worry that BDML could inadvertently introduce new forms of discrimination or reactivate old ones. In the US, the Equal Credit Opportunity Act and the Fair Housing Act make it illegal for a creditor to discriminate based on race, color, religion or gender. I will therefore study the interaction between new technologies and existing regulations.

3.1 Setup

As in the case of robo-advising, I describe a simple model and then I ask how Fintech might affect the equilibrium. The basic model follows the classic paper of Aigner and Cain (1977). There is a one-dimensional variable q that captures the credit quality of an individual. The distribution of q is normal with mean \bar{q} and variance $\bar{\sigma}^2$. I use ‘bars’ to denote population averages and also, under rational expectations, unconditional beliefs. It is convenient to work with precisions instead of variances so I define

$$\bar{\tau} \equiv \frac{1}{\bar{\sigma}^2}.$$

All lenders observe a quantitative signal y_1 about the credit quality of the borrower:

$$y_1 = q + \epsilon_1,$$

where ϵ_1 is normal with mean 0 and precision τ_1 . We can think of y_1 as a standard credit score. After observing the signal y_1 , the conditional distribution of q is normal with mean

$$\mathbb{E}[q | y_1] = \frac{\bar{\tau}\bar{q} + \tau_1 y_1}{\bar{\tau} + \tau_1}$$

and precision $\bar{\tau} + \tau_1$. Now imagine that the population includes two groups, denoted by z : a majority group $z = A$, and a minority group $z = B$. If a lender can condition on group membership the conditional expectation becomes

$$\mathbb{E}[q | y_1, z] = \frac{\bar{\tau}\bar{q}_z + \tau_{1z}y_1}{\bar{\tau} + \tau_{1z}}.$$

I assume that the minority group has a weakly lower average credit quality, $\bar{q}_B \leq \bar{q}_A$. For simplicity I assume that $\bar{\tau}_A = \bar{\tau}_B$ is the same for both groups but it is easy to allow for different population variances. A classic argument in the literature on discrimination is that standard data favor the majority, in the sense that $\tau_{1B} < \tau_{1A}$. The lesser precision of traditional signals could arise from a higher prevalence of missing data in the minority. For instance, if they are less likely to be employed, minorities have fewer payroll statements. If they are less likely to get a loan, they have shorter histories of repayments.

Discrimination can take several forms. Statistical discrimination refers to the fact that, as long as the signal y_1 is noisy, the posterior puts a positive weight on the prior. As a consequence, for a given signal y_1 , a member of the majority will receive a better score than a member of the minority.⁵ Under statistical discrimination, however, there is no average bias since the average of a conditional expectation is correct and equal to the true population mean: $\mathbb{E}[\mathbb{E}[q | y_1, z]] = \bar{q}_z$ for all z . Aigner and Cain (1977) and others have argued that statistical discrimination is unlikely to explain all – or even most – of the discrimination that we observe empirically. Dobbie et al. (2018) argue that the misalignment of incentives inside the firm can also lead to biases. They find that the short-term incentives given to loan officers create a long-run bias against immigrants. I therefore consider the stronger issue of prejudice in lending decisions.

3.2 Traditional Lenders

Traditional lenders meet face-to-face with borrowers. This has two consequences: they observe directly the type z , and they generate another signal $u = q + \epsilon_u$ about credit quality with precision τ_u . With two signals the conditional distribution of q is normal with mean

$$\mathbb{E}[q | y_1, z] = \frac{\bar{\tau}\bar{q}_z + \tau_u u + \tau_{1z}y_1}{\bar{\tau} + \tau_u + \tau_{1z}},$$

⁵Suppose for simplicity that $\tau_{1B} = \tau_{1A}$. Then $\mathbb{E}[q | y_1, A] - \mathbb{E}[q | y_1, B] = \frac{\tau_q(\bar{q}_A - \bar{q}_B)}{\tau_q + \tau_1}$.

and precision $\bar{\tau} + \tau_1 + \tau_u$. I assume that loan officers have a biased belief about the minority. They perceive the average quality to be $\hat{q}_B = \bar{q}_B - \delta$. We can think of δ as arising from prejudice or negative stereotyping. The conditional expectation of the lender is then

$$\hat{\mathbb{E}}(y_1, u, B) = \frac{\bar{\tau}(\bar{q}_B - \delta) + \tau_u u + \tau_{1B} y_1}{\bar{\tau} + \tau_u + \tau_{1B}}.$$

Note that the average statistical bias is simply $\mathbb{E}[\hat{\mathbb{E}}(y_1, u, B) | B] - \bar{q}_B = -\frac{\bar{\tau}}{\bar{\tau} + \tau_u + \tau_{1B}}\delta$. The harm potentially imposed on the minority is stronger if the prejudice δ is large and if the signal is noisy (τ_{1B} is small).⁶

Regulations Regulators often impose constraints on the use of group membership in lending decisions. For instance it is illegal in the U.S. to make lending decisions based on race or gender. We can capture this idea by assuming that regulations prevent lenders from conditioning *explicitly* on z . The score of the traditional lender then becomes

$$T(y_1, u; B) = \frac{\bar{\tau}(\bar{q} - \delta) + \tau_u u + \tau_1 y_1}{\bar{\tau} + \tau_u + \tau_1},$$

which does not depend explicitly on $z = B$. The subjective bias remains and on average we have

$$\mathbb{E}[T(y_1, u; B) | B] = \bar{q}_B - \frac{\bar{\tau}(\delta - \bar{d})}{\bar{\tau} + \tau_u + \tau_1},$$

where $\bar{d} \equiv \bar{q} - \bar{q}_B$ reflects the regulatory constraint that prevents statistical discrimination based on group status. Consistent with the empirical literature I assume that the negative bias has not been fully eliminated by regulations, i.e., $0 \leq \bar{d} < \delta$.

3.3 Univariate Unbiased Fintech Lending

Fintech lenders do not meet their clients face to face, but they have access to another quantitative signal

$$y_2 = q + \epsilon_2,$$

where ϵ_2 is normal with mean zero and precision τ_2 . We can think of this signal as coming from non-standard data sources, such as social media footprint or internet browsing history. The key point is that new data sources are less likely to favor the majority. Everyone has a browsing history and almost everyone is on social media, or at least there is no reason to believe that minorities are less likely to be on social media than the majority. I will thus assume that $\tau_{2B} = \tau_{2A}$.

⁶I consider for simplicity the case of risk neutral lenders but it is worth pointing out that the results only get stronger if lenders are risk averse. Risk averse lenders dislike lending to minorities because the residual conditional variance is higher. If we assume for example that lenders have mean-variance preferences and hence care about the score per unit of uncertainty $\mathbb{E}[q | y_1, z] / \mathbb{V}[q | y_1, z] = \bar{\tau}\bar{q}_z + \tau_u u + \tau_{1z} y_1$, then we see that a low value of τ_{1z} directly hurts the minority.

I refer to lending based on y_2 as *univariate unbiased Fintech lending*. By *unbiased* I mean that the algorithms do not suffer from the prejudice of humans in face to face interactions. As a result y_2 is an unbiased estimate of q . By *univariate* I mean that the signal is only about q . In particular it contains no *direct* information about z . It is of course correlated with z because of differences in population-averages \bar{q}_z , but conditional on q , it conveys no information about z . I will relax these assumptions in the next sections. With two signals the conditional distribution of q is normal with mean $\mathbb{E}[q | y_1, y_2] = \frac{\bar{\tau}\bar{q} + \tau_1 y_1 + \tau_2 y_2}{\bar{\tau} + \tau_1 + \tau_2}$ and precision $\tau_q + \tau_1 + \tau_2$. On average, conditional on $z = B$, we have

$$\mathbb{E}[\mathbb{E}[q | y_1, y_2] | B] = \bar{q}_B + \frac{\bar{\tau}\bar{d}}{\bar{\tau} + \tau_1 + \tau_2}.$$

We can finally compare Fintech lending to traditional lending from the perspective of the average minority borrower:

$$\begin{aligned} \mathbb{E}[\mathbb{E}[q | y_1, y_2] | B] - \mathbb{E}[T(y_1, u; B) | B] &= \frac{\bar{\tau}\bar{d}}{\bar{\tau} + \tau_1 + \tau_2} + \frac{\bar{\tau}(\delta - \bar{d})}{\bar{\tau} + \tau_u + \tau_1}, \\ &= \frac{\bar{\tau}}{\bar{\tau} + \tau_u + \tau_1} \left(\delta - \bar{d} \frac{\tau_2 - \tau_u}{\bar{\tau} + \tau_1 + \tau_2} \right). \end{aligned}$$

The term $\frac{\bar{\tau}\bar{d}}{\bar{\tau} + \tau_u + \tau_1}$ represents the gains from avoiding prejudice arising from face-to-face lending. It is unambiguously positive and follows directly from the assumption that algorithms do not suffer from prejudice. The second term can be positive or negative, depending on the quality of big data analytics. Since the evidence suggests that alternative data sources improve signal quality we can assume that $\tau_2 > \tau_u$, in which case the term is negative. Because Fintech lending is more precise it puts less weight on priors and it reduces the regulatory subsidy that the minority enjoys. Notice however that the sum of the two effects, $\delta - \bar{d} \frac{\tau_2 - \tau_u}{\bar{\tau} + \tau_1 + \tau_2}$, is always positive. This leads us to the following Proposition

Proposition 3. *Unbiased univariate Fintech lending reduces biases arising from prejudice against the minority. Even though it can undermine the effectiveness of existing policies the net effect is always positive for the minority.*

The proposition highlights the fundamental benefits of new lending technologies which come from two sources: no prejudice and more precise signals. Note that the higher precision of the signal would be helpful in any case, even if it was used by traditional lenders.

Remark 1. Even in the absence of Fintech lenders, giving traditional lenders access to the signal y_2 decreases biases against the minority.

If traditional lenders also observe y_2 they will form the posterior $\frac{\bar{\tau}(\bar{q} - \delta) + \tau_u u + \tau_1 y_1 + \tau_2 y_2}{\bar{\tau} + \tau_u + \tau_1 + \tau_2}$ and the bias will be reduced to $\frac{\bar{\tau}(\delta - \bar{d})}{\bar{\tau} + \tau_u + \tau_1 + \tau_2}$. If τ_2 is high the bias becomes small. We turn next to the case where big data contains multiple signals.

3.4 Multivariate Unbiased Fintech Lending

The model above does not capture one pervasive worry that policy makers have about big data, namely that it can inadvertently discover proxies for group membership. As Barocas and Selbst (2016) write “Approached without care, data mining can reproduce existing patterns of discrimination, inherit the prejudice of prior decision makers, or simply reflect the widespread biases that persist in society. It can even have the perverse result of exacerbating existing inequalities by suggesting that historically disadvantaged groups actually deserve less favorable treatment.”

Let us then assume that, in addition to y_2 as a proxy for q , big data analysis also generates a proxy for z . The important point here is that this second signal is a byproduct of the machine learning approach. The BDML system is looking for information, and to the extent that z is informative and that proxies for z can be constructed in the data, then the system will find it and use it. I assume that the information takes the form of a signal s about z . It is easy to model an imperfect signal but the intuition is the same if the signal is perfect and the derivation is simpler. I therefore assume that the ML system constructs a perfect, albeit indirect proxy for z . The Fintech lender’s conditional expectation is then

$$\mathbb{E}[q \mid y_1, y_2; B] = \frac{\bar{\tau}\bar{q}_B + \tau_1 y_1 + \tau_2 y_2}{\bar{\tau} + \tau_1 + \tau_2}$$

which is now an unbiased estimator of \bar{q}_B . Therefore $\mathbb{E}[q \mid y_1, y_2, B] = \bar{q}_B > \mathbb{E}[T(y_1, u; B) \mid B]$.

Proposition 4. *Even if Big Data and Machine Learning lead to indirect proxies for group membership, Fintech lending still reduces biases against minorities.*

Even in the most extreme case, unbiased Fintech lending produces pure statistical discrimination. To the extent that minorities suffered from prejudice in the pre-fintech era, their welfare should increase when Fintech lenders enter a market. This is consistent with the evidence in Bartlett et al. (2018) who find that Fintech algorithms discriminate 40% less than face-to-face lenders.

These results assume that $0 \leq \bar{d} < \delta$, i.e., that policy did not over-shoot its targets and some negative bias remained before the introduction of the new lending technology. While this seems to be empirically plausible on average, it does not rule out the possibility that in some location j the minority is doing so poorly that $\bar{d}_j \equiv \bar{q}_j - \bar{q}_{Bj}$ is larger than δ , even though the opposite is true on average across all locations. In that particular location credit market regulations lead to positive redistribution. Since the use of big data makes the credit market more neutral, it reduces redistribution and it could potentially hurt the minorities in some specific locations.

Broadly speaking, however, the analysis suggests that Fintech lending is likely to reduce discrimination as long as the algorithms themselves do not suffer from prejudice. Let us therefore consider next the case of biased algorithms.

3.5 Biased Multivariate Fintech lending

Suppose that Fintech engineers suffer from the same prejudice as loan officers and export their bias into their algorithms. Once the BDML system constructs a proxy for z it applies a penalty to the true quality of minority borrowers just as face-to-face officers did. This is arguably an extreme and unrealistic assumption because it is more difficult to induce biases in an algorithm than during face-to-face meetings, and the evidence in Bartlett et al. (2018) suggests that algorithms discriminate less. But this extreme assumption is helpful to make the point and highlight the key issue. The conditional – and now biased – expectation becomes

$$\hat{\mathbb{E}}[q \mid y_1, y_2, B] = \frac{\bar{\tau}(\bar{q}_B - \delta) + \tau_1 y_1 + \tau_2 y_2}{\bar{\tau} + \tau_1 + \tau_2},$$

and on average we have

$$\mathbb{E}\left[\hat{\mathbb{E}}[q \mid y_1, y_2, B] \mid B\right] = \bar{q}_B - \frac{\bar{\tau}\delta}{\bar{\tau} + \tau_1 + \tau_2}.$$

In this example Fintech lenders also suffer from prejudice. If we compare with traditional lenders we get

$$\begin{aligned} \mathbb{E}\left[\hat{\mathbb{E}}[q \mid y_1, y_2, B] \mid B\right] - \mathbb{E}[T(y_1, u; B) \mid B] &= \frac{\bar{\tau}(\delta - \bar{d})}{\bar{\tau} + \tau_u + \tau_1} - \frac{\bar{\tau}\delta}{\bar{\tau} + \tau_1 + \tau_2}, \\ &= \frac{\bar{\tau}}{\bar{\tau} + \tau_u + \tau_1} \left(\delta \frac{\tau_2 - \tau_u}{\bar{\tau} + \tau_1 + \tau_2} - \bar{d} \right). \end{aligned}$$

Even though $\bar{d} < \delta$ the term $\delta \frac{\tau_2 - \tau_u}{\bar{\tau} + \tau_1 + \tau_2} - \bar{d}$ can be negative if τ_2 is small: Fintech lending can now potentially harm the minority. On the other hand, if τ_2 is significantly larger than τ_u Fintech lending helps the minority despite its prejudice. Why is that? As before, Fintech lending increases precision. The impact of the prejudice depends on the precision of the credit scoring signal y_2 . Even if there is prejudice, its impact is small if τ_2 is large because the Bayesian part of the mechanism carries more weight.

Proposition 5. *Biased multivariate Fintech lending can decrease the welfare of the minority, but this becomes increasingly less likely as Fintech algorithms become more precise.*

There are two key lessons here. The first lesson is that for Fintech to hurt the minority two issues must interact: (i) BDML needs to build a direct proxy for group membership; and (ii) the algorithm itself must contain prejudice. If only one issue is present, Fintech improves welfare for the minority. It is only in the biased multivariate case that Fintech can be detrimental. The second lesson is that even in the pessimistic case, when Fintech lending becomes precise enough it always improves welfare for the minority. Even when engineers suffer from prejudice and somehow embed this prejudice into their algorithms, the prejudice decreases with the precision of the credit scoring model.

4 Conclusion

Fintech is likely to decrease the costs of financial intermediation, but also to create new regulatory issues. In this short note I have highlighted two relevant forces that will shape the impact of Fintech on inequality. In the case of robo-advisors, I have argued that the new pattern of fixed costs is likely to improve participation by relatively less wealthy household. This may not lower inequality across all groups, however. In the credit market, I have argued that alternate data sources are likely to reduce non-statistical discrimination. To the extent that minorities were hurt by prejudice or negative stereotyping, alternate data sources will have a positive impact. On the other hand new data can reduce the effectiveness of existing regulations.

References

- Abraham, F., S. L. Schmukler, and J. Tessada (2019). Robo-advisors: Investing through machines. *World Bank Policy Research Working Paper* (134881).
- Aigner, D. J. and G. G. Cain (1977). Statistical theories of discrimination in labor markets. *ILR Review* 30(2), 175–187.
- Barocas, S. and A. Selbst (2016). Big data’s disparate impact. *California Law Review* 104, 671–732.
- Bartlett, R., A. Morse, R. Stanton, and N. Wallace (2018). Consumer-lending discrimination in the era of fintech. Working paper.
- Bazot, G. (2013). Financial consumption and the cost of finance: Measuring financial efficiency in europe (1950-2007). Working Paper Paris School of Economics.
- Berg, T., V. Burg, A. Gombović, and M. Puri (2019). On the rise of fintechs – credit scoring using digital footprints. Working paper.
- Bickenbach, F., E. Bode, D. Dohse, A. Hanley, and R. Schweickert (2009, October). Adjustment after the crisis: Will the financial sector shrink? Kiel Policy Brief.
- Buchak, G., G. Matvos, T. Piskorski, and A. Seru (2018). Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics* 130(3), 453 – 483.
- Dobbie, W., A. Liberman, D. Paravisini, and V. Pathania (2018). Measuring bias in consumer lending.
- Economist, T. (2017). Silicon speculators.
- Fuster, A., M. Plosser, P. Schnabl, and J. Vickery (2019). The role of technology in mortgage lending. *The Review of Financial Studies* 32(5), 1854–1899.
- Garleanu, N. and L. H. Pedersen (2018). Efficiently inefficient markets for assets and asset management. *The Journal of Finance* 73(4), 1663–1712.
- Mayer, C. and K. Pence (2008). Subprime mortgages: What, where, and to whom? Staff Paper Federal Reserve Board.
- Moore, K. B. and M. G. Palumbo (2010, June). The finances of american households in the past three recessions: Evidence from the survey of consumer finances. Staff Paper Federal Reserve Board.
- O’Mahony, M. and M. P. Timmer (2009). Output, input and productivity measures at the industry level: The eu klems database. *The Economic Journal* 119(538), F374–F403.
- Pagnotta, E. and T. Philippon (2018, May). Competing on speed. *Econometrica* 86.
- Pedersen, L. H. (2015). *Efficiently Inefficient: How Smart Money Invests and Market Prices Are Determined*. Princeton University Press.
- Petralia, K., T. Philippon, T. Rice, and N. Véron (2019). Banking disrupted? financial intermediation in an era of transformational technology. Technical Report 22, Geneva Reports on the World Economy, ICMB and CEPR.
- Philippon, T. (2015). Has the us finance industry become less efficient? on the theory and measurement of financial intermediation. *The American Economic Review* 105(4), 1408–38.
- Philippon, T. (2016). The fintech opportunity. NBER Working Paper.