

## Blockchain analysis of the Bitcoin market. (Part 1)

Igor Makarov<sup>1</sup>, Antoinette Schoar<sup>2</sup>

<sup>1</sup>*London School of Economics Houghton Street London WC2A 2AE UK,  
i.makarov@lse.ac*

<sup>2</sup>*MIT Sloan School of Management 100 Main Street, E62-638 Cambridge, MA 02142  
and NBER, aschoar@mit.edu*

**Abstract.** The detailed analysis of the Bitcoin network and its main participants. The expert authors (Igor Makarov, London School of Economics, Antoinette Schoar, MIT Sloan School of Management) completed the study authorized by the National Bureau of Economic Research (NBER), the US-based private agency. The Bitcoin network is defined as a new database comprising many of public and proprietary sources to link bitcoin address to real object, and an extensive set of algorithms to extract information on market key players behavior. Three major pieces of analysis of the Bitcoin eco-system were conducted. First, the authors analyze the transaction volume and network structure of the main participants on the blockchain. Second, they document the concentration and regional composition of the miners which are the backbone of the verification protocol and ensure the integrity of the blockchain ledger. Finally, they analyze the ownership concentration of the largest holders of Bitcoin. The researchers found that 1/3 of all bitcoins issued were owned by 10,000 individual investors. They conclude that the high concentration makes the first cryptocurrency market vulnerable to hypothetical hacker attack. The translator notes that paraphrasing English text in Russian was rather challenging due to the newness of the financial agenda and introduction of the term *entity* extensively used in the Western countries though new to Russia. Nevertheless, it is necessary to introduce readers to the bitcoin technology which will be also practical and useful for the library and information community.

**Keywords:** cryptocurrency, bitcoin, blockchain, transaction, miner, multiple ownership

**Cite:** Makarov I., Shoar A. Blockchain analysis of the Bitcoin market. (Part 1) / I. Makarov, A. Shoar // Scientific and technical libraries. 2022. No. 9. P. 98–111. <https://doi.org/10.33186/1027-3689-2022-9-98-111>

## Introduction

Cryptocurrencies have seen a remarkable growth in value and public attention since their inception more than a decade ago. Opinions about the impact of cryptocurrencies range all the way from being a revolution in financial access to a threat to financial stability and monetary policy. A distinguishing feature of cryptocurrencies is the promise of a decentralized system of payments or store of value outside the traditional nexus of government scrutiny. The blockchain technology at the heart of cryptocurrencies replaces the reliance on a few centralized record keepers, such as banks or credit card networks, with a large set of decentralized and anonymous agents. The absence of centralized accountability and the anonymity of its users are often viewed as major benefits by crypto supporters, but it hinders the timely diagnosis of the health of the system, generates many challenges for regulators, and introduces new sources of systematic risk.

Bitcoin, the original cryptocurrency, is still the largest and most popular coin, with a market cap that is larger than all the other coins combined. It is often seen as a template or point of comparison for other new coins. Many industry participants are now calling for even wider Bitcoin adoption, either as a public investment vehicle or legal tender. These pressures put regulators who want to find the right balance between protecting the public interest and allowing innovation in a difficult position. There are still many open questions about the utilization of bitcoin, its ownership concentration as well as the structure of core entities that form the backbone of the Bitcoin ecosystem, despite being in existence for more than ten years. A better understanding of the Bitcoin network and its participants is required for any decision about how and whether to integrate Bitcoin into the traditional financial system.

In this paper, we aim to shed light on these open questions by developing a novel database that allows us to document the evolution of the Bitcoin market and its different participants over time. To build this database we use a large number of public and proprietary sources that link Bitcoin addresses to real entities and develop a suite of algorithms that

use the semi-public nature of the Bitcoin blockchain to extract information about the behavior of the main market participants. We believe that this is the most complete Bitcoin database used in academic research to date.

We conduct three major pieces of analysis that focus on the main participants of the blockchain eco-system. First, we analyze the transaction volume and network structure of the main participants on the Bitcoin blockchain. Second, we document the concentration and regional composition of miners which ensure the integrity of the blockchain ledger. Finally, we analyze the ownership concentration of the largest holders of Bitcoin.

**Transaction Volume and Network Structure.** We first document that 90% of transaction volume on the Bitcoin blockchain is not tied to economically meaningful activities but is the byproduct of the Bitcoin protocol design as well as the preference of many participants for anonymity. Because the Bitcoin blockchain is a public ledger all payment flows between addresses are perfectly observable. Therefore, many bitcoin users adopt strategies designed to impede the tracing of bitcoin flows by moving their funds over long chains of multiple addresses and splitting payments among them resulting in a large amount of spurious volume. We develop algorithms to filter out this spurious volume and trace economically meaningful payments between real entities on the Bitcoin network.

We show that the vast majority of Bitcoin transactions between real entities are for trading and speculative purposes. Starting from 2015, 75% of real bitcoin volume has been linked to exchanges or exchange-like entities such as on-line wallets, OTC desks, and large institutional traders. In contrast, other known entities are only responsible for a minor part of total volume. For example, illegal transactions, scams and gambling together make up less than 3% of volume<sup>1</sup>. The fraction of volume explained by miners is even smaller.

---

<sup>1</sup> Our estimates of illegal transactions are much smaller than the previous literature found, see for example Foley et al. (2019). One reason for this difference is that we have a much more detailed and comprehensive identification of participants on the blockchain. The prior work had to rely on an imputed network of illegal entities where any Bitcoin address recursively is classified as belonging to an illegal entity if the majority of its transactions is with addresses that themselves were previously classified as illegal. However, this method leads to significant overstatement of illegal volume, since it does not discriminate between real users and spurious volume.

Exchanges not only generate the most volume, but they are also the most connected nodes in the Bitcoin network. In particular, they have the highest measure of eigenvalue centrality<sup>2</sup>. Furthermore, a large fraction of exchange volume consists of cross-exchange flows. The high cross-exchange flows are the consequence of the current market structure. Different from traditional, regulated exchanges, cryptocurrency markets consist of many non-integrated and independent exchanges without any provisions to ensure that investors receive the best price when executing trades. As a result, the consistency of the Bitcoin price across exchanges depends on arbitrageurs and speculators who trade across them. In support of this idea, we show that exchanges that trade similar currency pairs have higher cross-exchange flows.

The strong interconnectedness of exchanges has important implications for the transparency and traceability of transactions, and especially the enforcement of KnowYour-Customer (KYC) norms, across the network. The current regulatory efforts focus on creating greater transparency through enforcement of KYC norms and capital gains tax reporting at the level of individual institutions, such as exchanges or payment processors. However, if users of Bitcoin can freely trade across regulated and unregulated exchanges or even countries with different enforcement levels, effective KYC regulation might not be possible at the level of individual institutions.

We use the example of Hydra Market, which is one of the largest dark net marketplaces, to study flows in this market. Our analysis shows that the highest volume entities interacting directly with Hydra Market users are non-KYC exchanges, including Binance and Huobi which are two of the largest exchanges worldwide. Once the flows arrive at these exchanges, they get mixed with other flows and become virtually untraceable, and so can be sent anywhere afterwards, even to exchanges that enforce KYC norms. In contrast, the direct interaction of KYC exchanges,

---

<sup>2</sup> See Section 3.3 for the definition and details. We show that the eigenvalue centrality can serve as a new and useful measure for ranking the volume and importance of exchanges because it is based on the cross-exchange Bitcoin flows on the blockchain, and therefore, is likely to be more resilient to manipulation than other measures.

such as Coinbase or Gemini, with Hydra Market users is modest. But their indirect interaction with flows originating from Hydra market is significantly larger, since these flows are channeled through a network of short-lived clusters, solely created for the purpose of obfuscating the origin of these funds.

These results highlight that non-KYC entities serve as a gateway for money laundering and other gray activities. The decentralized nature of the Bitcoin protocol makes it easy for these entities to operate – they only need to have their servers in a country where the authorities are willing to tolerate their existence. If KYC entities are allowed to accept flows from entities that are not following strict KYC norms (the current state), then the digital footprint has a very limited effect on preventing tainted flows from entering into wide circulation.

Even if KYC entities were restricted to deal exclusively with other KYC entities, preventing inflows of tainted funds would still be nearly impossible, unless one was willing to put severe restrictions on who can transact with whom and make every transaction subject to the approval of a blockchain “monitoring entity”, e. g. similar to what companies like Bitfury Crystal Blockchain<sup>3</sup> or Chainalysis are providing. Note that if this regime was to realize, the blockchain monitoring entities would become de facto trusted parties essential for the functioning of the Bitcoin network. But this is exactly what the Bitcoin protocol aims to overcome.

**Composition of Bitcoin Miners.** In a second major piece of analysis, we study the concentration and regional composition of Bitcoin miners, which are responsible for processing and verifying Bitcoin transactions and maintaining the integrity of the Bitcoin blockchain. For this service, miners are rewarded with newly created Bitcoins and transaction fees.

A proof of work protocol like Bitcoin requires a majority of decentralized miners to be honest for its record keeping function to work. If a single miner or a set of colluding miners were to command a majority of the mining power in the network, the ledger could become controlled by the colluding group and result in the infamous 51% attack, in which the group can alter the previously verified records. The possibility of such attacks creates systemic risks for financial stability and potentially even

---

<sup>3</sup> <https://crystalblockchain.com/>

for national security if a large fraction of citizens uses Bitcoin as a store of value.

It is therefore important to understand how concentrated the mining capacity is. The previous literature has mainly focused on mining pool concentration. By design, the probability of mining a block and obtaining a block reward in the Bitcoin blockchain is proportional to the hashing power spent on mining. This provides strong incentives for miners to pool their computing power and co-insure each other. As a consequence, mining in the Bitcoin blockchain is dominated by mining pools.

But while pools function like aggregators of hashing capacity and can therefore have substantial influence over the Bitcoin protocol, they do not necessarily control their miners. As Cong et al. (2020a) emphasize, the power that a pool operator has vis a vis the miner depends on the ease with which miners can shift capacity across pools, which in turn depends on the underlying size distribution of the miners.

Unlike information about mining pools, which is commonly available, information about individual miners is not readily available. We identify individual miners by tracking the distribution of mining rewards from the largest 20 mining pools to the miners that work for them. Since each pool uses its own algorithm to distribute rewards, we build separate algorithms for each pool. To the best of our knowledge, this is the first study that accurately links miners to their mining pools.

We show that the Bitcoin mining capacity is highly concentrated and has been for the last five years. The top 10% of miners control 90% and just 0.1% (about 50 miners) control close to 50% of mining capacity. Furthermore, this concentration of mining capacity is counter cyclical and varies with the Bitcoin price. It decreases following sharp increases in the Bitcoin price and increases in periods when the price drops or. Thus, the risk of a 51% attack increases in times when the Bitcoin price drops precipitously or following the halving events.

In addition, we show that there is significant geographic clustering of miners. While it has been previously discussed that a large majority of mining pools are registered in China, this does not automatically mean that miners have to be located in China. So far, the main data about miners' location has come from the analysis of miners' IP 4 addresses from a

few select pools. When a miner connects to a pool server, the pool operator can see the IP address of the miner. Unless a miner uses a VPN address, the pool operator can use this IP address to determine the geographical location.

Here, we utilize a new approach, which takes advantage of our ability to trace miners on the blockchain. Since we can trace miners' addresses and Bitcoin transactions, we can see at which exchanges they use to cash out their rewards. The idea is that miners in a particular region would most likely send their rewards to an exchange that is also in this region. Using our approach we show that starting in 2015 and until April 2020 a majority of mining capacity, between 60% to 80% is located in China, which is consistent with anecdotal evidence.

In order to verify the validity of our approach of identifying miner locations by looking at where miners cash out their Bitcoin rewards, we use a recent incidence in April 2021 in the Xinjiang province of China. After a devastating coal mining accident, the government shut down coal mining and electricity supply for the entire area. Many Chinese Bitcoin miners are located in this province due to the cheap supply of coal powered electricity. Of course, not all Chinese miners are located in this area and thus we do not use it as a test of the mining capacity in China. But the shutdown of electricity for more than two days allows us to identify a set of miners for which we can be sure that they are physically located in China since they had to stop their operations. Using this strategy, we confirm that these Chinese miners, indeed utilize the cashing out policies that we had conjectured.

**Ownership concentration.** Finally, we study the ownership and concentration of Bitcoin holdings. Since the inception of Bitcoin, there has been intense interest in the question of who are the largest owners of Bitcoin, and how much do they actually own. There are websites dedicated to tracking the addresses with the largest Bitcoin holdings, the so called "rich list," one of the most well-known and widely followed lists in the crypto community. But the question of ownership concentration is not only a matter of curiosity and intrigue. From a public policy perspective, it is important to understand who is positioned to benefit most from any price appreciation that would happen if regulators allow a broader adoption of Bitcoin. Are these a select few investors or the general public?

Determining the concentration of ownership is more complicated than just tracking the holdings of the richest addresses, since many of the largest addresses belong to cold wallets of exchanges and online wallets, which hold Bitcoin on behalf of many investors. We develop a suite of algorithms based on graph analysis to classify addresses into those belonging to individual investors or those belonging to intermediaries<sup>4</sup>.

We show that the balances held at intermediaries have been steadily increasing since 2014. By the end of 2020 it is equal to 5.5 million bitcoins, roughly one-third of Bitcoin in circulation. In contrast, individual investors collectively control 8.5 million bitcoins by the end of 2020. The individual holdings are still highly concentrated: the top 1000 investors control about 3 million BTC and the top 10,000 investors own around 5 million bitcoins.

The rest of the paper is structured to first discuss the data sources and the construction of the data set. The next section documents the evolution of volume to different participants on the blockchain, in particular, we develop algorithms to separate spurious volume from real volume and then map the network structure of participants. In the following section we analyze miners, their composition and geographic concentration. And finally we document the ownership concentration of Bitcoin participants.

## 1. Related Literature

Our paper contributes to a fast-growing literature on cryptocurrencies and blockchains. Raskin and Yermack (2016) and Hardle et al. (2020) provide a broad perspective on the economics of cryptocurrencies and the blockchain technology they are built upon.

Budish (2018), Abadi and Brunnermeier (2018), and Biais et al. (2019) study consensus mechanisms and limitations of the proof-of-work protocol, the core innovation of this new technology.

Athey et al. (2016), Cong et al. (2020b), Pagnotta and Buraschi (2018), Sockin and Xiong (2020), and Han and Makarov (2021) develop different theoretical frameworks to study bitcoin adoption and bitcoin pricing and highlight that beliefs about adoption are central for Bitcoin

---

<sup>4</sup> See Section 5 for a detailed description of the identification.



pricing. Schilling and Uhlig (2019) propose a model, in which a cryptocurrency such as Bitcoin coexists and competes with a traditional government-issued fiat money.

A number of papers study the economics of Bitcoin mining. Prat and Walter (2021) examines the relationship between the Bitcoin price and the investment in hashing capacity. Easley et al. (2019) and Huberman et al. (2021) develop equilibrium models of Bitcoin mining fees. Cong et al. (2020a) propose a theory of mining pools and suggest that mining pools escalate miners' arms race and significantly increase the energy consumption of proof-of-work-based blockchains. Ferreira et al. (2019) model the joint behavior of miners, mining pools, and firms producing specialized mining equipment. We contribute to this literature by developing a suite of algorithms to identify individual miners on the blockchain. This data is the first to trace individual miners and allows us to study their concentration and regional composition.

Similar to our paper, Foley et al. (2019) use the Bitcoin blockchain data to examine the prevalence of illegal transactions on the Bitcoin blockchain. Wallet-level blockchain data are also used by Griffin and Shams (2020) to study whether tether issuance affects bitcoin prices. In comparison to the earlier literature, we develop a novel database that not only has a much more comprehensive classification of participants on the blockchain, but also eliminates spurious volume. This granular data allows us to attribute economically meaningful transactions more precisely and to provide a detailed analysis of the evolution of the Bitcoin market.

## **2. Data**

All bitcoin transactions are recorded on a distributed public ledger, the so-called blockchain. Transactions are organized in blocks that are added to the ledger every 10 minutes on average. Each block contains a few thousand transactions. A typical Bitcoin transaction includes a list of senders and recipients represented by pseudonymous addresses, the number of bitcoins sent and received, and a time-stamp of the transaction.

We download the blockchain data using the open-source software of Bitcoin Core and use the BlockSci program to parse the raw data into in-

dividual transactions<sup>5</sup>. As of June 28, 2021, there have been 689,000 blocks of 652 million Bitcoin transactions and 896 million addresses organized in a blockchain database of more than 379 GB in size.

An address on the blockchain can be thought of as a bank account. Anyone can send bitcoins to any address. But to send bitcoins from a given address one needs to know a password associated with this address. Unlike bank accounts, Bitcoin addresses can be generated freely, so typically the same entity controls several addresses, and in some cases, even tens of millions of different addresses.

The Bitcoin community developed several heuristics to assign addresses to the same entity. As a starting point, we use the most conservative method to cluster addresses whereby all addresses that send bitcoins in any single transaction are deemed to belong to the same entity<sup>6</sup>.

This heuristic is justified by the Bitcoin protocol that requires the party that signs a transaction to have control of all output addresses

In practice, a user typically only needs to specify the destination addresses and the amounts to be transferred. A special piece of software, called a wallet, then decides which addresses to send bitcoins from to cover a given amount that the user wants to transfer. This process then allows the clustering algorithm to successfully group all user's addresses together. It should be stressed however, that with a little bit of effort, a user can deliberately conceal the connections between his different addresses by making sure that no two addresses are ever used in the same transaction. As a result, this clustering heuristics only produces a lower bound for the true number of distinct entities.

To link address clusters to real entities we scrape cryptocurrency blogs and websites, such as Reddit, Blockchain.info, bitinfocharts.com, bitcointalk.org, walletexplorer.com, and Matbea.com for all publicly avail-

---

<sup>5</sup> Bitcoin Core and BlockSci are available at <https://bitcoin.org/en/bitcoin-core/> and <https://github.com/citp/BlockSci>, respectively.

<sup>6</sup> See Ron and Shamir (2012) or Meiklejohn et al. (2013). Bitcoin mixing services, such as CoinJoin, let users mix their coins with other users, and are designed to confuse this heuristic. The BlockSci accounts for that and avoids CoinJoin transactions in its clustering algorithm. See BlockSci documentation for more details.

lable addresses of prominent Bitcoin entities such as exchanges, payment processors, gambling sites, and others. We supplement this information with the state-of-the-art database of crypto entities from Bitfury Crystal Blockchain. Bitfury Crystal Blockchain is one of the leading providers of anti-moneylaundering tools and analytic solutions in the crypto space.

To the best of our knowledge, we have the most complete information about crypto entities that have been used in academic research up to this point. Our data cover 1,043 different entities. These include 393 exchanges, 86 gambling sites, 39 on-line wallets, 33 payment processors, 63 mining pools, 35 scammers, 227 ransomware attackers, 151 dark net market places and illegal services.

### **3. Bitcoin Blockchain Volume**

#### **3.1. Spurious Volume**

The design of the Bitcoin blockchain and the preference of many of its users for anonymity creates a lot of spurious volume that is not tied to economically meaningful transactions. In this section, we describe how we identify and separate this volume from the real volume, i. e. payments for goods and services and other financial transfers between two parties. It is instructive to start by looking at a particular example, see the transaction depicted in

Figure 1<sup>7</sup>. In this transaction, the address “17A16Q...” sends its balance to the following three addresses “3QKAn2...”, “1F8fDp...”, and “17A16Q...”. The amount received is equal to the amount sent except for a small fee of 0.001 bitcoins, which is a part of the block reward. Notice that the last of the three addresses is the same as the sending address, that is, the address “17A16Q...” sends the majority of its balance to itself. This means the overall volume this transaction generates on the blockchain is large. However, the economically meaningful volume generated in the transaction (the real volume), which is the volume between different entities, is small.

The above situation where an address sends its balance to itself or to another address controlled by the same entity is very common. In part,

---

<sup>7</sup> This is the second transaction in block 600,000 and can be seen e. g., at <https://explorer.btc.com/btc/block/600000>.

it is a consequence of the design of the Bitcoin protocol. The outstanding balance of an address is not stored in the address but is imputed from the whole history of transactions involving this address by traversing back the Bitcoin ledger. For computational efficiency, the Bitcoin protocol allows one to send only the amounts that have been previously received by an address. For example, suppose an address previously received 5, 7, and 10 bitcoins, so the outstanding balance is 22 bitcoins.

To send 8 bitcoins from this address one can either send 10 bitcoins, or any of the following linear combinations:  $5 + 7$ ,  $5 + 10$ ,  $7 + 10$ ,  $5 + 7 + 10$ . Since in any case, the amount is larger than 8 bitcoins the sender needs to collect the difference using one of his addresses. This process creates a large amount of spurious volume that obscures the true volume of transactions on the blockchain.

Another common reason for spurious volume is the preference of blockchain participants for anonymity. Because the bitcoin blockchain is a public ledger all payment flows between addresses are perfectly observable. Many Bitcoin users, therefore, adopt strategies designed to impede the tracing of bitcoin flows.

Consider, for example, a situation where a hacker demands payment from a company to be sent to a Bitcoin address he controls. Since the ransom address is public information, if the hacker later sends bitcoins from this address to a third party, the party could easily flag funds as coming from illegal activity. To prevent this from happening, hackers often try to obfuscate tracing by creating multiple addresses and splitting the initial payment among them. This process is usually repeated many times resulting in the so-called “peeling chains”, where funds travel a long distance from one address to another leading to a large amount of fictitious volume on the ledger.

Peeling chains are also commonly used by many exchanges, such as Coinbase and Kraken, and many mining pools. These entities, every time they need to collect a change as in the transaction in Figure 1: Bitcoin transactions and spurious volume (view original figure here: [https://www.nber.org/system/files/working\\_papers/w29396/w29396.pdf](https://www.nber.org/system/files/working_papers/w29396/w29396.pdf)), generate a new address instead of re-using the old address. This new address is then used to send funds to another entity, and the change is col-

lected in another new address. This process is usually repeated many times until all initial balance is spent. The addresses used in peeling chains are usually used only to receive and immediately send bitcoins with a typical lifetime span of 10 hours.

There are two ways how one can account for peeling chain transactions. First, one could modify the clustering algorithm to add addresses in peeling chains to the corresponding clusters. The other approach, which we follow in this paper, is to backtrack volume in peeling chains to the original clusters and discard any intermediate addresses from further analysis. To backtrack this volume we develop an efficient recursive algorithm detailed in the Appendix.

Factoring out peeling chains reduces the computational burden and results in significant reduction of addresses and clusters. While the original database has 896 million addresses, after we remove addresses in peeling chains we end up with 640 million addresses. These addresses belong to 189 million clusters, of which 116 million clusters are single-address clusters.

Figure 2: Decomposition of volume: Internal, pass-through, and real volume (view original figure here: [https://www.nber.org/system/files/working\\_papers/w29396/w29396.pdf](https://www.nber.org/system/files/working_papers/w29396/w29396.pdf)) shows the decomposition of total Bitcoin blockchain volume into what we call internal, pass-through, and real volume. Internal volume is the within-cluster volume, that is, the volume that is generated when a cluster sends bitcoins to itself. The pass-through volume is the transitory volume associated with peeling chains. Finally, the real volume is the remaining volume, which represents transfers between clusters. This volume accounts only for 10% of the total Bitcoin volume on the blockchain, with 90% of the Bitcoin volume on the blockchain not tied to economically meaningful transactions.

## References

1. **Foley S., Karlsen J., Putniņš T.** Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies? // *The Review of Financial Studies*. 2019. № 32 (5). P. 1798–1853. <https://doi:10.1093/rfs/hhz015>

2. **Cong Y, Ulasli M, Schepers H. et al.** Nucleocapsid Protein Recruitment to Replication-Transcription Complexes Plays a Crucial Role in Coronaviral Life Cycle // *J Virol*. P. 169–176. 2020. № 94 (4). doi: 10.1128/JVI.01925-19
3. **Bitcoin** Core. URL: <https://bitcoin.org/en/bitcoin-core/> (accessed: 05.08.2022).
4. **BlockSci**. URL: <https://github.com/citp/BlockSci> (accessed: 05.08.2022).
5. **Ron D., Shamir A.** Quantitative Analysis of the Full Bitcoin Transaction Graph. URL: <https://eprint.iacr.org/2012/584.pdf> (accessed: 05.08.2022).
6. **Meiklejohn C., Holmbeck M., Siddiq M. et al.** An Incompatibility between a Mitochondrial tRNA and Its Nuclear-Encoded tRNA Synthetase Compromises Development and Fitness in *Drosophila* // *PLOS Genetics*. № 9 (1). doi: 10.1371/journal.pgen.1003238

---

### Information about the authors

**Igor Makarov** – London School of Economics Houghton Street London WC2A 2AE UK  
i.makarov@lse.ac

**Antoinette Schoar** – MIT Sloan School of Management 100 Main Street, E62-638  
Cambridge, MA 02142 and NBER  
aschoar@mit.edu

---