Essays on Computational and Empirical Methods in Financial Economics

by

Anil Donmez

M.A., Bogazici University, 2017
B.Sc., Bogazici University, 2014

Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of
Doctor of Philosophy

in the
Department of Economics
Faculty of Arts and Social Sciences

SIMON FRASER UNIVERSITY
Spring 2022

Declaration of Committee

Name: Anil Donmez

Degree: **Doctor of Philosophy**

Thesis title: **Essays on Computational and Empirical Methods in Financial Economics**

Committee:

Chair: Simon Woodcock
Associate Professor, Economics

Alexander K. Karaivanov
Supervisor
Professor, Economics

Jasmina Arifovic
Committee Member
Professor, Economics

Deniz Anginer
Committee Member
Assistant Professor, Finance

Christina Attanasova
Examiner
Professor, Finance

Dogan Tirtiroglu
External Examiner
Professor
Department of Business Management
Ryerson University

Ethics Statement

The author, whose name appears on the title page of this work, has obtained, for the research described in this work, either:

a. human research ethics approval from the Simon Fraser University Office of Research Ethics

or

b. advance approval of the animal care protocol from the University Animal Care Committee of Simon Fraser University

or has conducted the research

c. as a co-investigator, collaborator, or research assistant in a research project approved in advance.

A copy of the approval letter has been filed with the Theses Office of the University Library at the time of submission of this thesis or project.

The original application for approval and letter of approval are filed with the relevant offices. Inquiries may be directed to those authorities.

Abstract

This thesis is composed of three essays on the applications of computational and empirical methods in financial economics.

Chapter 1, Transaction Fee Economics in the Ethereum Blockchain, is co-authored with Alexander Karaivanov. It examines the economic determinants of transaction fees in the Ethereum blockchain. We estimate an empirical model based on queueing theory and analyze the factors determining the "gas price" (transaction cost per unit of service, "gas"). Using block- and transaction-level data from the Ethereum blockchain, we show that changes in service demand significantly affect the gas price—when there is high block utilization, per-unit fees increase on average, with a strong nonlinear effect above 90% utilization. The transaction type is another important factor—a larger fraction of regular transactions (direct transfers between users) is associated with higher gas price.

Chapter 2, Individual Evolutionary Learning and Zero-Intelligence in the Continuous Double Auction, jointly with Jasmina Arifovic and John Ledyard, studies behavior in a Continuous Double Auction, the most preferred exchange mechanism of financial markets around the world. Particularly, we report on two models, Zero-Intelligence and Individual Evolutionary Learning, which we tested against each other with a key emphasis on price formation and trade efficiency using two very different data sets: a large, uncontrolled set from classroom experiments using the MobLab interface and a small, controlled set from experiments at Simon Fraser University.

Chapter 3, Racial Differences in Senior Executives' Access to Information, is co-authored with Deniz Anginer, Nejat Seyhun, and Ray Zhang. Based on a hand-collected sample of race data for executives of S&P 1500 firms, this paper provides evidence of differences due to race in insider-trading behavior and in the profitability of senior corporate executives. We document that, although non-African-American executives make positive abnormal profits from insider trading, African-American executives, on average, earn zero abnormal profits. In contrast, the abnormal profits of Asian-American executives are similar to or even exceed those of Caucasian executives. However, these race differences are less profound in firms that emphasize diversity and employee equity. These results suggest that African-American executives are disadvantaged relative to other executives in access to insider information.

Keywords: blockchain, Ethereum, queueing theory, time series, transaction fees, allocative efficiency, continuous double auction, learning & adaptation, insider trading, empirical asset pricing, discrimination.

Dedication

This thesis is dedicated to the memory of Professor Jasmina Arifovic, who provided me with her full assistance both academically and financially until her untimely passing away, the loving memory of my grandfather, my dear family for their constant encouragement, and my beloved wife, Elif, for her unwavering and unconditional support.

Acknowledgements

I am immensely grateful to my supervisors, Alexander Karaivanov, Deniz Anginer, and Jasmina Arifovic, for their invaluable advice and continuous support throughout my PhD studies. I deeply appreciate everything I have learned from them and their assistance in writing this thesis.

I am particularly indebted to my co-authors, Alexander Karaivanov, Deniz Anginer, Jasmina Arifovic, John Ledyard, Nejat Seyhun, and Ray Zhang, whom I have benefited a great deal from their expertise during our collaboration.

I would like to thank my family for the love, support, and constant encouragement they gave me over the years even from thousands of miles away.

Last but not least, I would like to express my sincerest gratitude to my dear wife, Elif, for her endless love and support that kept me running during these years.

Table of Contents

List of Tables

List of Figures

Chapter 1

Transaction Fee Economics in the Ethereum Blockchain

1.1 Introduction

We study the economic determinants of transaction fees in the Ethereum blockchain, the second largest blockchain by market capitalization.[1] While Bitcoin, launched in 2009, is the oldest, best-known and most widely researched blockchain platform, Ethereum, started in 2015, has a different design, with broader applicability going beyond digital payments. Similar to other blockchains, Ethereum has an internal digital (crypto) currency with decentralized and scarce supply, called Ether (ETH), that can be used as store of value or transacted between users anywhere in the world. However, unlike other blockchain platforms, Ethereum is much more flexible and programmable, hence many developers have used the platform to create a wide range of decentralized applications ("DApps") and "smart contracts".[2]

A distinctive feature of Ethereum is its internal metering variable called *gas*. Each blockchain transaction has an algorithmically defined *gas requirement* – a pre-specified execution cost expressed in units of gas. The more complex a transaction or a smart contract is, the higher is its gas requirement, which must be paid by the user for the transaction to be recorded and executed on the blockchain. The simplest and most common transaction type, which we call "regular transaction", is a transfer of ETH between two blockchain

[1] On December 10, 2020 the market capitalization of Ethereum was $64 bln, second only to Bitcoin with capitalization $341 bln. Source: http://coinmarketcap.com

[2] At the heart of Ethereum's extensive functionality is the "Ethereum Virtual Machine" (EVM), a Turing-complete virtual computation engine, Buterin (2013). Thanks to the EVM design, Ethereum users are not only able to execute simple transactions such as sending digital currency from one address to another, but can also create and make "calls" (e.g., supply data or make automated trade requests) to "smart contracts". A smart contract is computer code that can be used to digitally facilitate, verify, or enforce the negotiation or execution of trades (e.g, see Cong and He, 2019). Examples include "DeFi" (decentralized finance) applications that allow users to borrow, lend and invest digital assets; exchanges for trading digital currencies; Ethereum-based (ERC-20) digital tokens; games; gambling platforms; voting platforms; and supply chain management systems.

addresses. Such transaction requires 21,000 gas to be executed, for any transferred ETH amount. Smart contract creations and smart contract calls have substantially higher gas requirements (10 to 20 times higher, see Figure 1.4).

The key component of transaction costs in Ethereum, and the main focus of our analysis, is the *gas price*. The gas price is a bid price (in ETH) per unit of gas, specified by the user when posting a blockchain transaction (transfer or contract creation/call). Together, the gas requirement and the gas price determine the total transaction cost which equals (gas used)$\times$(gas price). For example, an ETH transfer between two addresses requires exactly 21,000 gas and hence its transaction fee is 21,000$\times$(gas price). This implies that both the transaction complexity/type (its gas requirement) and the users' willingness to pay per unit of gas (the gas price) determine the cost of posting a transaction. The transaction type choice, including contract creation or calls, and its interplay with the gas price adds a dimension to our analysis which is mostly absent in the Bitcoin blockchain.

We characterize the economic determinants and dynamics of transaction fees in the Ethereum blockchain by estimating an empirical time-series model based on queueing theory. In contrast to opinions that blockchain activity is highly speculative and volatile,[3] we find that, by and large, gas price levels and dynamics comply with standard economic predictions from queueing theory and supply/demand theory. First, both the marginal Ethereum gas price (the minimum gas price bid within a block) and the median block gas price are higher on average when the blockchain experiences higher utilization or congestion, consistent with the theory. Second, we find that the effect of blockchain utilization on the gas price is strongly non-linear. When the block occupancy is below a threshold (90 percent), the effect of higher utilization on the marginal and median gas price is insignificant. However, beyond that threshold, the effect of blockchain utilization on the gas price is positive and increasing in a convex way. Third, higher marginal and median gas prices are positively and statistically significantly associated with a higher fraction of regular transactions per block and negatively associated with the fraction of contract creations and contract calls.

On Figure 1.1 we display the daily average value of our main variable of interest – the marginal gas price in ETH or USD, defined as the minimum gas price in each block, over the period of study November 1, 2017 to January 31, 2019. We also show the implied marginal transaction fee for a regular transaction (ETH transfer). The minimum gas price in a block is the bid price per unit of gas at which the last (marginal) transaction was recorded and corresponds to the gas price at which a transaction is just on the margin of being included vs. not being included in the current block (see Section 2.2 for more discussion). We see that the gas price exhibits significant variation over time, allowing us to analyze and identify its

[3]For example, on April 14, 2021 US Federal Reserve Chairman Jerome Powell said about cryptocurrencies "They're really vehicles for speculation" (CNBC). See Foley et al. (2019), Gandal et al. (2018), Griffin and Shams (2020), Li, Shin and Wang (2020) or Yermack (2014) for formal discussion and analysis.

Figure 1.1: Gas price

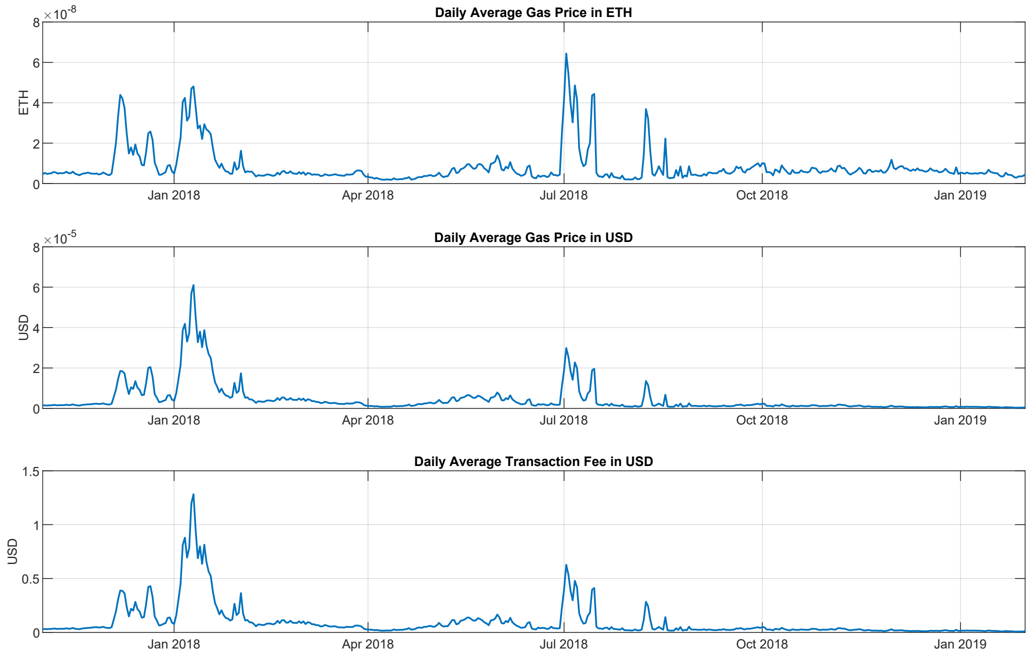

Notes: The top and middle panels of Figure 1.1 plot the daily average of the minimum gas price in each block ('marginal gas price') over time, measured in ETH or USD respectively. The bottom panel plots the implied minimum transaction fee in USD for a 'regular transaction' (ETH transfer between two accounts), computed as 21,000×(the daily average USD gas price).

determinants.[4] In addition, to manage spikes in the gas price, the Ethereum protocol can increase the service rate by raising the block gas limit to accommodate a larger number of transactions.[5] We control for these algorithmic supply changes in our empirical analysis.

Each Ethereum block has a block gas limit which can be thought of as the block's 'capacity' and which, together with the number of blocks created in a day, determines the gas supply.[6] The gas supply is kept stable by the blockchain algorithm automatically adjusting the cryptographic difficulty to ensure that block creation is spaced out evenly in time. Supply fluctuations are thus small and can be considered exogenous, arising mostly

[4]Although the median value of the marginal gas price is $0.000018, implying a $0.04 median fee for a regular transaction (see Table A.1), there are times, e.g., in January 2018 or July 2018, when the gas price was substantially higher, implying $1 or larger fee for a regular transaction. Relatively low transaction fees which are independent of the transfer amount are an important advantage of cryptocurrency platforms compared to conventional payment services providers. For example, Paypal, the largest online payment platform, charges $5 or more for international transfers. (`https://www.paypal.com/us/webapps/mpp/paypal-fees`).

[5]One such increase occurred on December 10, 2017, when the block gas limit was raised by 18%, see Figure 1.6.

[6]The block capacity (size) in Ethereum is measured in gas (not bytes), unlike Bitcoin which uses a fixed 2MB block size.

because of the randomness in the time to complete the cryptographic proof-of-work.[7] We do control for the few system-wide gas supply changes implemented via the blockchain code in the sample period (see Appendix A.2 for more details). Given the limited and stable gas supply, the gas prices bid by the users therefore determine the priority order of transactions included in each block. Too low gas price may cause a transaction's inclusion to be delayed, possibly indefinitely.

When posting any transaction on the Ethereum blockchain, in addition to the gas price, the user must also specify a transaction *gas limit*, that is, the maximum gas the transaction can use up. The transaction would be executed as long as its gas requirement does not exceed the transaction gas limit.[8] In addition, the sum of the gas requirements of all transactions included in a block cannot exceed the block gas limit. The need for a costly gas requirement and gas limit stem from the virtual machine basis of Ethereum and helps avoid problems such as infinite loops, coding errors, sabotaging the network, etc.[9]

The transaction fee (gas price) bid and paid by the users is the key mechanism providing economic incentives to operate the blockchain network and to verify and record transactions (a block creation reward also exists but is being reduced over time).[10] In Ethereum the blockchain transactions are verified and blocks are created by computing pools ("miners"), through a competitive cryptographic problem-solving consensus mechanism called Proof-of-Work. In order to maintain their operations, there must be sufficient financial incentive for the miners, covering their time, equipment and electricity costs.

Related literature

This paper relates to complementary theoretical and empirical research on cryptocurrencies in economics and finance. Huberman et al. (2019) show that Bitcoin's decentralized design

[7]The cryptographic problem is very costly to solve but its solution is easy to verify. Hence, the Proof-of-Work mechanism, together with the rest of the blockchain software code, ensures the security, consensus, and stability of the network.

[8]For example, a user would normally set the gas limit of a regular transaction to 21,000 gas, which would satisfy exactly the gas requirement for such transaction. In contrast, the gas limit for a smart contract creation or call must be set higher, to ensure that the limit covers the gas requirements of the resulting EVM instructions, depending on the contract's complexity.

[9]Ethereum is Turing complete – in theory a program of any complexity can be computed by the Ethereum Virtual Machine (EVM), see Buterin (2013). However, this flexibility can introduce security, stability and resource management problems. For example, Turing-complete systems can be set to run infinite loops. If executing arbitrary code did not face a resource constraint and if a transaction or smart contract causes such a loop, either by mistake or deliberately, this could destabilize or disable the Ethereum network. To prevent such problems, each allowed EVM instruction has a pre-defined cost in units of gas. The execution of any transaction or smart contract is automatically terminated if the gas consumed by running it exceeds the available gas (the gas limit) for the transaction. This 'resource constraint' feature of Ethereum therefore maintains Turing completeness while capping and putting a price on the system resources that any blockchain transaction can consume.

[10]See Catalini and Gans (2019) for further discussion and analysis of the verification and networking costs in blockchains.

protects users from monopoly pricing, derive closed-form expressions for fees and waiting times and compare Bitcoin payments to a traditional payment system. Consistent with our findings for Ethereum, they show that Bitcoin transaction fees increase in non-linear way with congestion, however, we find, in addition, that user choice over the transaction type is another key factor.

Easley et al. (2019) develop and test empirically a game-theoretic model of Bitcoin transaction fees and the strategic behavior of miners and users. They find that transaction fees are positively correlated with the average waiting time.[11] Since detailed waiting time data for the Ethereum platform is unavailable, we construct a variable measuring blockchain utilization directly from the blockchain records, which to our knowledge is new in the literature. Möser and Böhme (2015) study 45.7 mln Bitcoin blockchain transaction records and find that changes in the system protocol or the actions of big intermediaries can trigger regime shifts in the transaction fee level. Chiu and Koeppl (2019) examine theoretically the relationship between Bitcoin transaction fees, block size and user characteristics. Their simulation results suggest that users are willing to pay more when their transactions are more urgent and when the block size is smaller.

Our work complements these papers but differs in two important ways. First, we use detailed micro-level data (transaction-level and block-level) directly downloaded from the Ethereum blockchain. Second, unlike most authors who study Bitcoin, we instead analyze Ethereum's internal gas and gas price variables, which have clear economic interpretation as a scarce resource and its endogenous price.[12] We abstract from monetary policy issues related to the blockchain technology on which there exists a large separate literature.[13]

[11]In other related work, Kasahara and Kawahara (2019) model transaction execution as a non-preemptive priority queueing game and show that users' waiting time is determined not only by their own posted fee, but also by the arrival rate of users with higher transaction fees. Li et al. (2018) model transaction confirmation as a single-server queue with batch service and priority mechanism and show that the average waiting time is affected by the share of users from each priority class.

[12]See also Zochowski (2019) for a non-technical review of the network characteristics and transaction fee dynamics in several blockchain platforms including Ethereum.

[13]There is also a finance literature on cryptocurrency prices and returns determination, predictability and volatility, e.g., Ciaian et al. (2015), Kristoufek (2015), Corbet et al. (2018), and Athey et al. (2016) among others.

1.2 Data

1.2.1 Data sources

We use data obtained directly from the Ethereum blockchain.[14] By construction, the blockchain records a publicly accessible permanent copy of the complete transaction history since Ethereum's launch in 2015. We downloaded both block-level and transaction-level data, including transaction receipts data. The block-level data contain information about each block included in the ETH blockchain. The data contain the block's sequential number, difficulty level[15], the block *gas limit* which determines the maximum capacity of the block, the *gas used* which is the total gas consumed by all transactions included in the block, and a timestamp. The main block variables that we use in the empirical analysis are the *gas limit* and *gas used*, which determine the blockchain network utilization at the time of block creation.

The transaction-level blockchain data include: the block number in which the transaction is recorded, the transaction value in Wei,[16] the sender address, the receiver address, a "nonce" value which indicates the number of prior transactions posted by the sender address, the gas price in Wei set by the sender, the maximum gas amount that the transaction can use (transaction gas limit), and the actual gas quantity used by the transaction (obtained from the transaction receipt data). We merge the block-level and transaction-level data using the block number. The end result is a complete transaction-level blockchain dataset in which block specific information is preserved. For part of the analysis we also combine the blockchain data with additional data on the ETH price in USD or other currencies, obtained from the website min-api.cryptocompare.com at the daily level.

We analyze Ethereum gas prices in the period between November 1$^{\text{th}}$, 2017 and January 31$^{\text{th}}$, 2019. This period does not include significant changes in the core Ethereum blockchain protocol and code. This enables us to focus on the interplay between the blockchain platform and user demand. The first three months of the sample period, which we call the "peak period" and also analyze separately, capture several historic peak events as of the end of 2020 - the highest daily transaction count, the highest ETH price in USD, and highest total market capitalization of ETH.[17]

[14]We used the publicly shared Python scripts by E. Medvedev, github.com/blockchain-etl/ethereum-etl to extract and save the Ethereum blockchain data. The raw data files were then merged, processed and analyzed in Matlab and Stata using code written by the authors and available on request.

[15]Each block has a computational difficulty level corresponding to the algorithmic cryptographic problem solved by miners to verify and confirm the block transactions.

[16]Wei is the smallest denomination in Ethereum and equals 10^{-18} ETH.

[17]The Ethereum blockchain is experiencing another peak period in early 2021.

Our sample period includes 2,688,667 blocks and 309,760,480 transactions in total. The raw data have slightly fluctuating time frequencies, since block creation times differ.[18] To mitigate hour-of-the-day and other high frequency periodical variation in the data, we aggregate the raw transaction data at the daily level. This generates 457 daily observations that we use in the empirical analysis.

1.2.2 Variables and definitions

Our main variable of interest is the daily *marginal gas price*, defined as the daily average of the minimum observed gas price in each of the blocks recorded in a given day. The minimum gas price in a block is the user bid price per unit of gas at which the last (marginal) transaction was recorded. Hence, the marginal gas price variable captures most closely the price (fee per unit of gas) at which a transaction is just on the margin of being included vs. not being included in the current block. In robustness analysis we also use two alternative gas price variables: the *median gas price* (daily average of the median gas price in each block recorded on that day) and the *lowest 5-th percentile* gas price (daily average of the bottom 5th percentile gas price in each block recorded on that day).

Second, we construct a *blockchain utilization* variable to capture the usage level, or congestion rate, of the platform. Blockchain utilization is defined as the ratio between the sum of the gas requirements of all recorded transactions in a given day and the *gas supply* in the same day. We define the gas supply as the sum of the gas limits of all blocks recorded in a given day. The blockchain utilization variable therefore measures the fraction of total available gas supply that is used up, per unit of time. For example, daily blockchain utilization of 0.8 means that on average 80% of block capacity is used in a given day.

Third, to account for the fact that there are different types of transactions in Ethereum, we define the variable *regular transactions share* as the ratio of the number of regular transactions (simple ETH transfers between two addresses) to the number of all transactions recorded in a given day. This variable captures the possible effect of changes in the composition of posted transactions (for example, more vs. less urgent) on the gas price. In Section 4.4 we show that regular transactions are more likely to have higher gas prices and to be recorded near the top of their blocks; this is consistent with interpreting regular transactions as more urgent on average.

Table A.1 displays summary statistics of the variables defined above, and also the ETH price in USD or in terms of a basket of currencies.

[18] On average, an Ethereum block is created every 13 to 14 seconds.

1.2.3 Descriptive findings

Figures 1.2-1.6 illustrate the magnitudes and time dynamics of key Ethereum blockchain variables relevant for our empirical analysis, in our sample period. The “peak period" (Nov. 1, 2017 to Jan. 31, 2018) is shaded.

Figure 1.2 displays the daily average blockchain utilization (block occupancy rate). The Ethereum network operated close to its full capacity at times, particularly in the second half of the peak period and also around July and August 2018. For the remainder of the study period, the blockchain network utilization was well below 1, implying no significant congestion and available room for additional transactions in most blocks.

Figure 1.2: Blockchain utilization

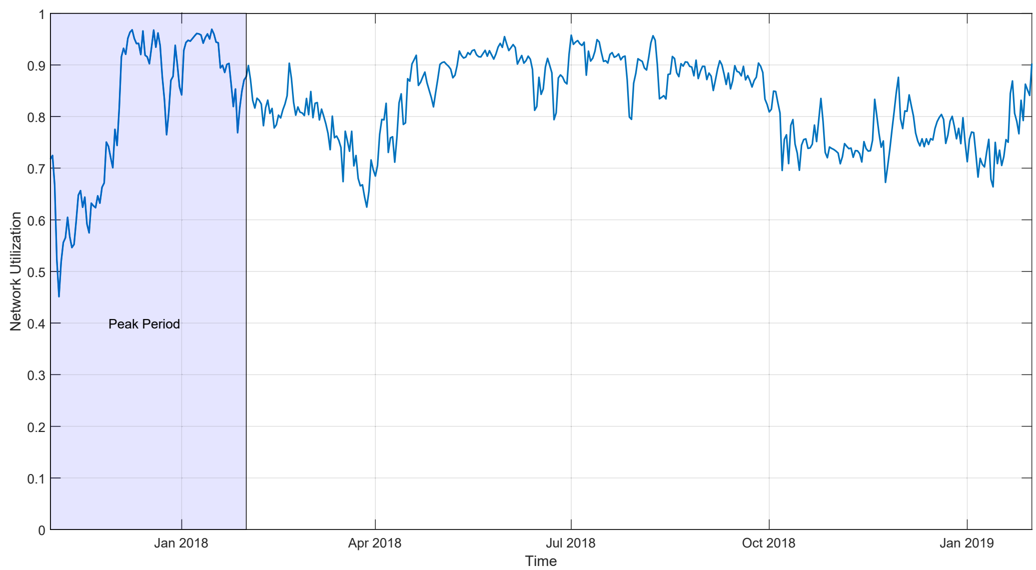

Notes: Regular transactions share. The regular transactions share is defined as the ratio between the number of regular transactions (direct ETH transfer between two addresses) and the number of all transactions recorded in a given day. The shaded region denotes the peak period between November 1, 2017 and January 31, 2018.

Figure 1.3 shows the daily regular transactions share. During the peak period in the beginning of the sample, the share of regular transactions rises steadily from about 0.45 to as high as 0.8. This suggests an increasing number of user transactions transferring ETH from one address to another, as opposed to smart contract creations or calls. In the remainder of the study period the regular transactions share goes back down and stays around 0.5.

Figure 1.4 displays the daily average gas used (gas requirement) by transaction type. Regular transactions (ETH transfer between two addresses) require and use 21,000 units of gas for execution, regardless of the transferred amount. In contrast, contract calls and contract creations have significantly higher gas requirement, depending on their complexity

Figure 1.3: Regular transaction share

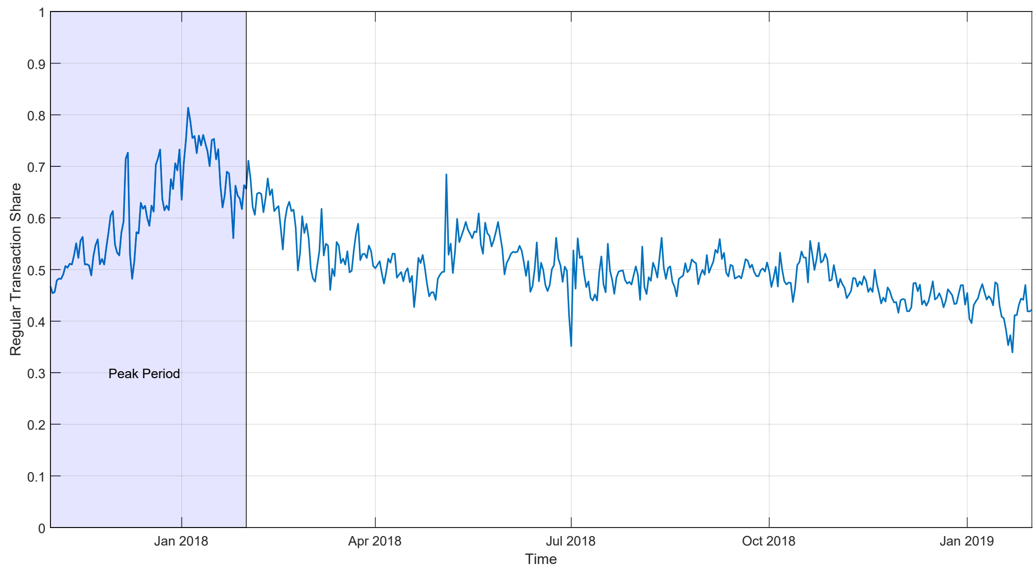

Notes: Required gas by transaction type. Daily average of the required gas amounts (gas requirements) by transaction type. The blue dashed (red straight) line denotes regular transactions (contract calls and contract creations).

Figure 1.4: Required gas by transaction type

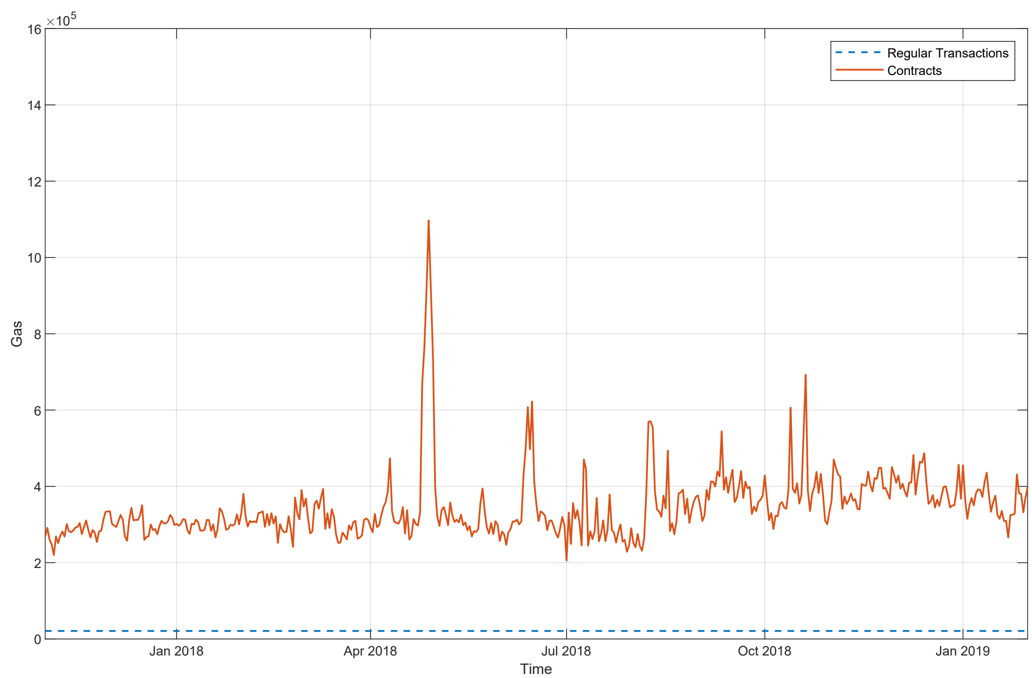

Notes: Required gas by transaction type. Daily average of the required gas amounts (gas requirements) by transaction type. The blue dashed (red straight) line denotes regular transactions (contract calls and contract creations).

– in our study period, the daily mean (median) gas requirement for contract calls and creations is 344,000 (324,000) gas.

Figure 1.5: Gas supply

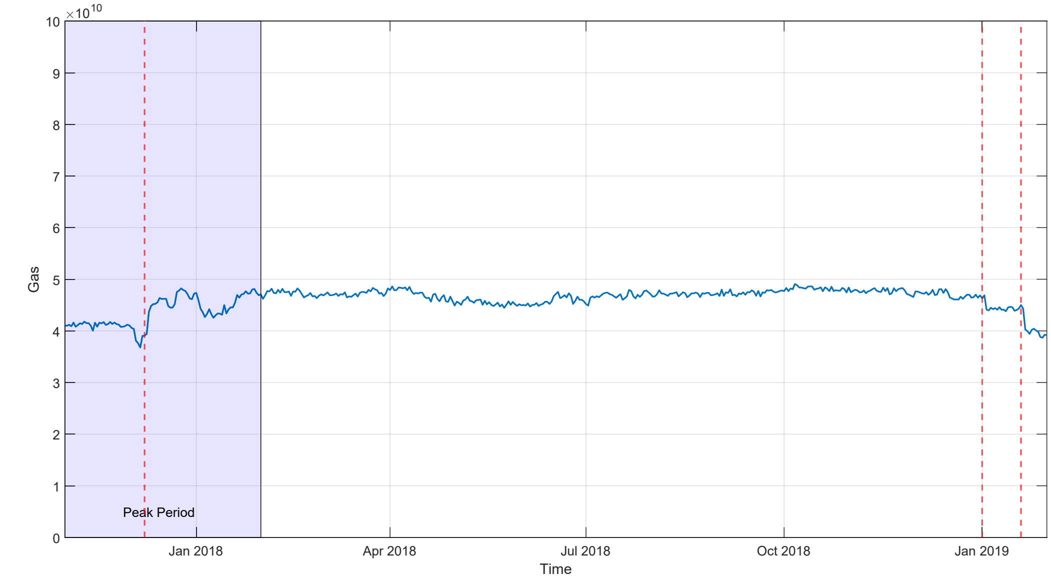

Notes: Gas supply. Gas supply is the sum of the gas limits of all blocks recorded on the blockchain in a given day. The shaded region denotes the peak period between November 1, 2017 and January 31, 2018. The vertical dashed lines denote exogenous system-wide changes in the gas supply.

Figure 1.6: Average block gas limit and daily block count

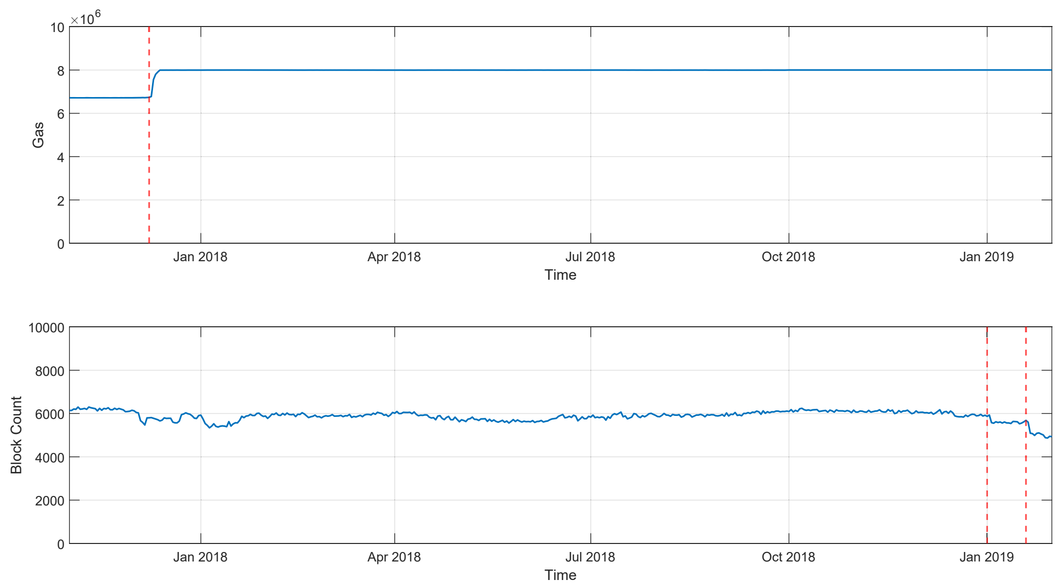

Notes: Average block gas limit and daily block count. The average block gas limit (upper panel) is the daily average of the gas limits of all blocks recorded on a given day. The vertical dashed line highlights December 10, 2017, when the Ethereum network initiated a system-wide increase of the block gas limit of approximately 18% that took three days to complete. The daily block count (lower panel) is the total number of blocks recorded on the blockchain in a given day. The vertical dashed lines denote January 3 and 21, 2019, the two dates on which the Ethereum protocol increased the cryptographic difficulty resulting in a decline in the number of blocks created per day.

Figure 1.5 plots the daily gas supply, defined as the sum of the gas limits of all blocks recorded in a given day. The gas supply variable proxies the total service rate of the

Ethereum blockchain, since block space is restricted by the block gas limit. Figure 1.5 shows that the gas supply is stable, with the exception of the beginning of the peak period (Dec. 2017) and the end of the sample (Jan. 2019). The Dec. 2017 increase was caused by a change in the Ethereum protocol which raised the block gas limit by approximately 18 percent as of December 10, 2017. The decrease in gas supply at the end of the sample was triggered by a system-wide difficulty increase. We control for these events in the empirical analysis (see also Figures 1.6 and 1.7 and Appendix A.2 for more details).

1.3 Model

1.3.1 Preliminaries

We model the demand for transactions using queueing theory. Transaction execution in the Ethereum blockchain is an example of a priority queueing system. Users who aim to obtain higher priority set higher per-unit transaction fee (gas price). This implies that if there are more users with higher waiting costs then, on average, the marginal gas price would be higher. Similarly, when the blockchain utilization (block occupancy rate) is higher, the cost of waiting would dominate the cost of transaction execution, which results in higher gas prices paid by the users. The predicted effect of the blockchain utilization on the gas price is hence positive.

The supply side of the blockchain consists of so-called "miners", decentralized providers of computing power who service user requests by verifying, executing and recording user-submitted transactions on the blockchain. To maximize their profits, miners sort the submitted transactions in descending gas price order, that is, the transactions with higher gas prices are included first (at the top) of the current block.[19] This means that a transaction's priority and position within its block is determined by the transaction's gas price. Moreover, miners' activity cannot be preempted by external "higher priority" job requests. Consequently, a natural way to model the transaction execution process in Ethereum is as a "non-preemptive priority queueing mechanism" (Shortle et al., 2018).

The Ethereum blockchain algorithm strives to minimize the impact of fluctuations in computer power supply (mining hash power) on transaction fees. Specifically, the Ethereum protocol is designed to keep the service rate (the average time between consecutive blocks) as stable as possible by automatically adjusting the Proof-of-Work cryptographic difficulty, so that higher availability of computational power (e.g., new miners or more computer power coming online) is quickly offset by higher cryptographic difficulty. This automatic adjustment, illustrated on Figure 1.7, generates a stable service rate (gas supply) in the blockchain

[19] 85% of all blocks in our data contain transactions which are perfectly sorted in descending gas price order and the rest of the blocks contain only minor exceptions from descending order. See Appendix A.2 for more details.

platform, as shown on Figure 1.5. Fluctuations in the gas supply are thus caused by exogenous factors (randomness from small variations in the block creation rate) or updates to the blockchain protocol. As an example of the latter, the observed increase in gas supply at the beginning of the study period occurred because of the implementation, via the blockchain protocol, of a system-wide increase in the block gas limit on December 10, 2017 (see Figure 1.6). We control for these system-wide changes in the empirical analysis.

Figure 1.7: Cryptographic difficulty

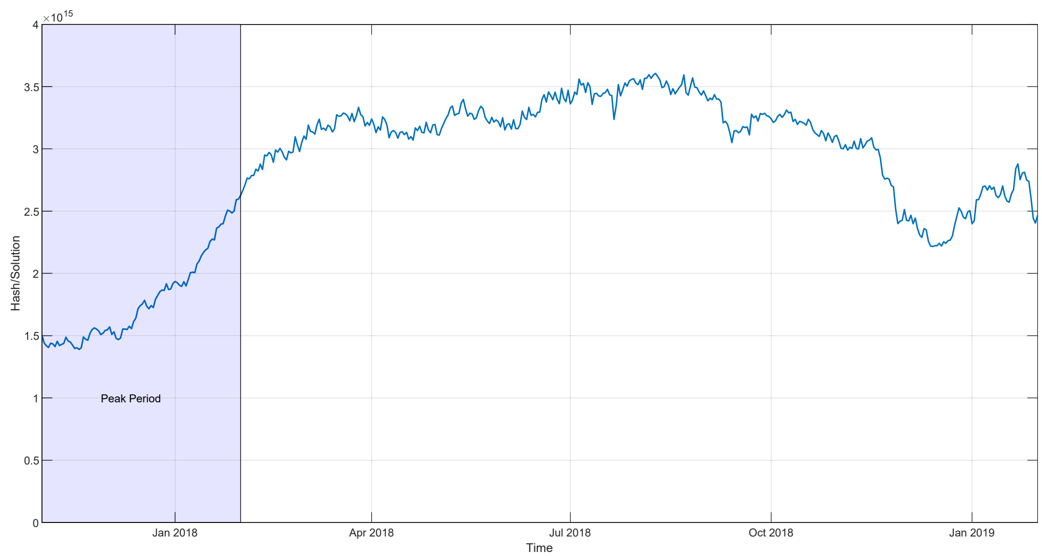

Notes: Cryptographic difficulty. The figure plots the daily average cryptographic difficulty measured as the number of hash function evaluations per block. The peak period is shaded.

1.3.2 Testable implications

Assume that each block is created at rate μ, called the service rate and that there are K priority transaction classes with corresponding arrival rates λ_i, $i = 1, ..., K$, where $i = 1$ denotes the highest priority class. In our setting, transaction(s) with the highest gas price constitute the highest priority class. The users' priority classes are thus not fixed in Ethereum – we can think of each user being assigned a type depending on the gas price they choose. The users' waiting time is then endogenously determined.

The total transaction arrival rate is $\lambda = \sum_{i=1}^{K} \lambda_i$. To simplify the notation, define the variable, $\rho = \frac{\lambda}{\mu}$, which measures how busy the platform is. To satisfy system stability we must have $\rho \leq 1$.[20] Applying Little's Law[21] yields

$$w_1^q = \frac{\lambda_1}{\mu(\mu - \lambda_1)} \tag{1.1}$$

where w_1^q is the waiting time of the highest priority class, $i = 1$. For priority classes $i = 2, ..., K$, the average waiting time is therefore

$$w_i^q = \frac{\sum_{j=1}^{i} \lambda_j}{(\mu - \sum_{k=1}^{i-1} \lambda_k)(\mu - \sum_{j=1}^{i} \lambda_j)} \tag{1.2}$$

There are two takeaways. First, the average waiting time of priority class i increases with the arrival rates of the higher priority classes $k \leq i$. Second, the average waiting time of each priority class $i = 1, ..., K$ decreases in the service rate μ.

The main trade-off in the queuing model is between transaction costs and waiting costs. Users of type i are assumed to have exogenous waiting costs per unit of time, c_i. The total cost incurred by user type i is thus

$$C_i = c_i w_i^q + p_i g_i$$

where p_i is the gas price bid by user i and g_i is the required gas for user i's transaction.

Specifically, if a user sets a low gas price, she would pay a low transaction fee if her transaction is picked by the block's miner. However, the transaction's priority among all other submitted transactions would be low, which would increase the waiting time and consequently the cost of waiting. Alternatively, if a user's unit waiting cost c_i (transaction urgency) is high, then the user would choose to bid a higher gas price p_i to shorten the waiting time which again affects the total cost.

The modeled relationship and trade-off between the gas price, transaction urgency and waiting time generates the following testable hypotheses which we evaluate empirically using the Ethereum blockchain data:

Hypothesis H1. When the blockchain experiences higher utilization or congestion (high λ), the marginal and median gas price are higher, holding all other variables constant.

Hypothesis H2. When there are more urgent transactions (high c_i's), the marginal and median gas price are higher, holding all other variables constant.

The model also implies that when the service rate is higher (high μ), the marginal and median gas price would be lower, holding demand and all other variables constant. As em-

[20]If ρ >1, i.e., the arrival rate is greater than the service rate, then the system queue will grow infinitely long.

[21]Little's Law, a theorem in Little (1961), asserts that the average number of customers in a queueing system is equal to the rate at which customers arrive and enter the system times the average sojourn time of a customer.

phasized earlier, the blockchain protocol by design automatically adjusts the cryptographic difficulty to offset changes in computing power supply and maintain a constant service rate and gas supply. Therefore, we do not include the gas supply variable in our main empirical specification and results (we show in a robustness check that including it does not affect our findings).

1.4 Empirical analysis

1.4.1 Estimation strategy

We estimate the following empirical specification:

$$p_t = \alpha + f(B_t) + \beta_2 R_t + \beta_3 X_t + \epsilon_t \tag{1.3}$$

The dependent variable p_t in equation (1.3) is the natural logarithm of the gas price (daily average) observed in the Ethereum blockchain, measured either in ETH or USD. We report results with the marginal, median and the lowest 5-th percentile gas price, as defined in Section 2.2. The variable B_t is the daily average blockchain utilization (block occupancy) which captures the usage or congestion level in the network. We estimate both a simple linear specification, $f(B_t) = \beta_1 B_t$, as well as a piece-wise linear specification (see Section 4.2) which allows for a threshold effect in blockchain utilization (a quadratic specification is also considered as robustness check). The variable R_t is the regular transactions share, included to control for changes in the composition of submitted transactions (higher vs. lower urgency). In Section 4.4, we show evidence suggesting that regular transactions have higher urgency on average compared to contract calls. Finally, X_t denotes control variables, e.g., the price of ETH in USD in Tables 1.1 and 1.2.

In the Ethereum blockchain, users make their transaction fee payment in ETH, the internal cryptocurrency of the platform. That is, the ETH gas price determines the transaction's priority among all other waiting transactions. However, it is possible that users are less concerned about the ETH transaction fee they must pay and more concerned about its real-value equivalent in US dollars or the national currency in which they receive income and pay bills. To capture the possible effect of the price (exchange rate) of ETH in terms of conventional currency on the gas price bidding choice of users, we also include the ETH price in USD in the specifications in which the gas price p_t is measured in ETH. An increase in the dollar price of ETH makes the ETH value of a given transaction fee relatively cheaper in USD terms, hence a user may be willing to to bid a higher ETH gas price. The expected correlation between the ETH gas price and the ETH price in USD is negative. In Section 4.3 we also consider a specification using the ETH price in terms of a basket of the three most common local currencies used by Ethereum nodes (USD, CNY and EUR).

Mapping blockchain utilization, B_t and the regular transactions share, R_t in the data to system congestion and the share of urgent transactions, respectively, in the queueing model, hypotheses H1 and H2 imply a positive association of the gas price p_t with both B_t and R_t. The daily gas supply is stable during the study period except for three system-wide events described in Appendix A.2, see Figure 1.6.[22]

1.4.2 Baseline results

Table 1.1 reports our baseline estimation results using equation (1.3). The dependent variable is log of the marginal gas price, as defined in Section 2, measured in ETH or USD. To correct for possible heteroskedasticity and autocorrelation in the error terms we use Newey-West standard errors with maximum lag 4.

Columns (1) and (3) in Table 1.1 use a linear specification for blockchain utilization, i.e., $f(B_t) = \beta_1 B_t$ in equation (1.3). A scatterplot of the marginal gas price and blockchain utilization, Figure 1.8, however, shows that their relationship is non-linear, with much larger gas prices observed when the blockchain utilization (block occupancy) is close to 1. We formally estimate this non-linear relationship in columns (2) and (4) of Table 1.1, where we allow for a threshold effect at 90% utilization. Specifically, in addition to $\beta_1 B_t$, we also include in $f(B_t)$ the binary variable D_t equal to 1 if $B_t > 0.9$ and zero otherwise and its interaction with the utilization variable, $TR_t = B_t * D_t$. We selected the 90% threshold using a formal structural break test (Andrews, 1993), which shows that the blockchain utilization series has a break at 0.9. In addition, at 90% utilization there may be insufficient block space to include certain types of transactions.[23] We also perform a robustness check using a quadratic form for $f(B_t)$, see Section 4.3.4 and Table A.4.

The results in Table 1.1, columns (1) and (3) show that the blockchain utilization variable, B_t measuring the network usage or congestion is positively and statistically significantly associated (at the 1% significance level) with the marginal gas price. Columns (2) and (4), which allow for a non-linear utilization effect, further clarify that the utilization impact is negligible below the threshold level of 0.9 (the B_t coefficient is weakly or not statistically significantly different from zero). However, in both the ETH and USD specifications (2) and (4), the coefficient estimate on the threshold interaction term TR_t is positive and statistically significantly different from zero at the 1% level, implying a strongly non-linear relationship between blockchain utilization and the gas price above 90 percent utilization. Quantitatively, the estimate 27.85 in column (4) means that 0.01 increase in blockchain uti-

[22] We control for these exogenous events affecting gas supply by including separate binary variables that take value of 1 on December 10-12, 2017, on January 3, 2019 and on January 21, 2019, respectively and equal 0 otherwise.

[23] Contract creations are the most complex transactions, with the largest gas requirement. A direct examination of the data shows that there are contract creation transactions using more than 10% of their block gas limit.

Figure 1.8: Network utilization and the marginal gas price

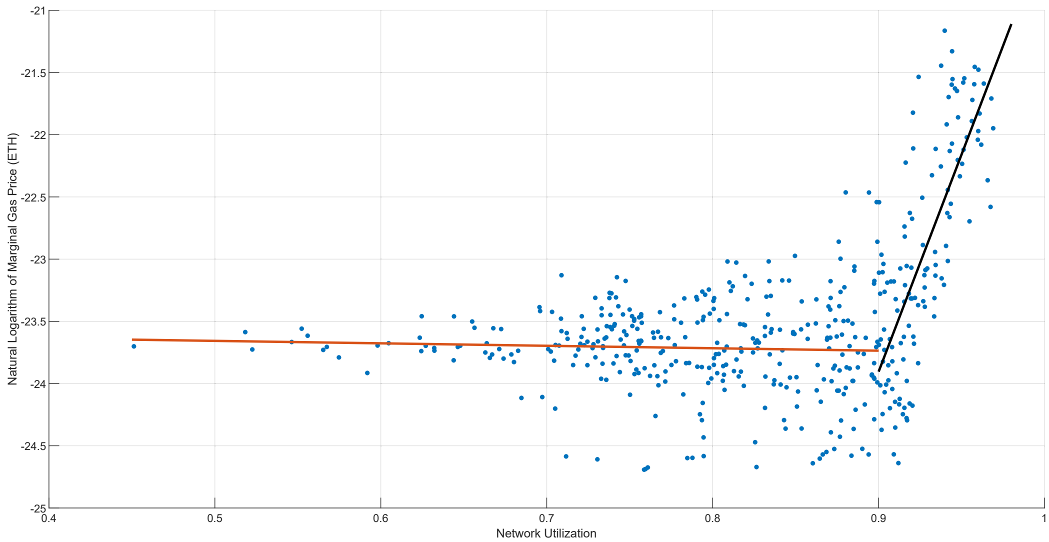

Notes: Network utilization and the marginal gas price. Scatterplot of blockchain utilization, B_t and log of the marginal gas price, p_t, with fitted lines

Table 1.1: Main results

	marginal gas price, ETH		marginal gas price, USD	
	(1)	(2)	(3)	(4)
blockchain utilization, B_t	2.263***	0.090	3.310***	0.885*
	(0.624)	(0.306)	(0.796)	(0.476)
utilization > 90%, TR_t	-	33.03***	-	27.85***
	-	(4.052)	-	(4.868)
regular transactions share, R_t	4.464***	2.539***	8.255***	6.659***
	(1.000)	(0.668)	(0.796)	(0.744)
ETH price in USD, X_t	-0.732***	-0.651***	-	-
	(0.208)	(0.160)	-	-
sample size	457	457	457	457
R-squared	0.293	0.595	0.647	0.760

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable, "marginal gas price" is the natural logarithm of the minimum observed gas price (in 10^{-8} ETH or 10^{-5} USD) over all transactions in a block, averaged over all blocks created on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization $> 90\%$" is the product of the blockchain utilization, B_t and a binary variable which equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price in USD" is the daily average ETH price in 10^3 USD. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

lization above the threshold level of 0.9 is associated with $100(e^{27.85(0.01)})$ or 32.1% increase in the marginal gas price in USD. Given that the mean USD marginal gas price in our sample is $\$0.45(10^{-5})$ (see Table A.1), a user executing a transaction with complexity 100,000 gas would pay $0.45(10^{-5})(10^5)(32.1\%) = 0.15$ USD (15 cents) more on average following a 0.01 increase in the blockchain utilization when the latter is above 90%.

The regular transactions share, R_t has positive and statistically significant estimates in all four specifications of Table 1.1. This implies that a larger share of regular transactions is associated on average with a higher gas price in ETH and USD. For example, the estimate 6.66 in column (4) suggests that a 0.01 increase in the regular transactions share is associated with $100(e^{6.66(0.01)})$ or 6.9 % increase in the marginal gas price in USD. This implies that a user executing a transaction with complexity 100,000 gas would pay $0.45(10^{-5})(10^5)(6.9\%)$ $= 0.03$ USD (3 cents) more, on average, following a 0.01 increase in the regular transactions share, holding all else equal.

We also include the price of ETH in USD in specifications (1) and (2) in Table 1.1, to measure the responsiveness of blockchain users to the conventional currency cost of a given ETH-denominated transaction fee. We obtain a negative estimate which implies that, when executing a transaction is more expensive in USD terms, the marginal user tends to bid a lower ETH gas price, all else equal. This result is consistent with the notion that users take the real value of transaction costs into consideration in their bidding decisions.

Our main results in Table 1.1 show strong evidence in support of Hypotheses H1 and H2 from Section 3. Higher blockchain utilization, measured by the average block occupancy rate, is associated with a higher marginal gas price (Hypothesis H1). Higher share of urgent transactions, proxied by the regular transactions share, corresponds to a higher marginal gas price (Hypothesis H2).

So far we interpreted the right-hand side variables in equation (1.3) as exogenous. However, a potential simultaneity problem may exist – a higher gas price makes all transactions costlier but the associated impact could be larger for more complex transactions such as contract calls and creations.[24] Hence, a higher gas price may raise the share of users choosing regular transactions because of the lower gas requirement.

To address the potential simultaneity between the gas price, p_t and the transaction type choice, we use the first lag of the regular transactions share, R_{t-1} and the daily number of new sender accounts as instruments. Since R_{t-1} was decided at time $t-1$, it does not have a direct causal impact on the dependent variable p_t, while R_t and R_{t-1} are highly correlated (Reed, 2015). Similarly, the number of new sender accounts is highly correlated with the regular transactions share (correlation 0.68) and does not have a direct causal effect on the minimum gas price.

[24]Contracts and regular transactions can be used interchangeably for some purposes, e.g., sending ETH between accounts could be done by both a contract call and a regular transaction.

Table 1.2: Instrumental variables (IV) results

	marginal gas price, ETH		marginal gas price, USD	
	(1)	(2)	(3)	(4)
blockchain utilization, B_t	2.266***	0.114	3.154***	0.830*
	(0.624)	(0.303)	(0.929)	(0.471)
utilization > 90%, TR_t	-	32.37***	-	26.79***
	-	(4.367)	-	(5.750)
regular transactions share, R_t	5.621***	2.990***	8.830***	7.130***
	(1.494)	(0.981)	(1.419)	(1.218)
ETH price in USD, X_t	-0.980***	-0.743***	-	-
	(0.287)	(0.191)	-	-
sample size	456	456	456	456
F-statistic (C-D weak instrument test)	267.1	239.7	327.7	282.0
χ^2-stat (DWH endogeneity test)	2.455	0.761	0.385	0.408
R-squared	0.285	0.594	0.646	0.759

Notes: Instrumental variables (IV) regressions with daily data including a constant. Newey-West standard errors reported in the parentheses. The first-stage regresses the regular transactions share on its first lag, the total number of new sender accounts and on the blockchain utilization and the ETH price in USD. The dependent variable "marginal gas price" is the natural logarithm of the minimum observed gas price (in 10^{-8} ETH or 10^{-5} USD) over all included transactions in a block, averaged over all blocks created on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price in USD" is the daily average ETH price in 10^3 USD. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

Table 1.2 reports the IV regression results, using the same specifications (linear or piece-wise linear in blockchain utilization) of equation 1.3 as in Table 1.1. The first-stage regresses the regular transactions share R_t on the instruments, blockchain utilization and the price of ETH in USD. The reported coefficient estimates in Table 1.2 are for the instrumented share. We also report two IV diagnostic statistics – the Cragg-Donald Wald F-statistic for weak instruments and the chi-squared Durbin-Wu-Hausman endogeneity test. The null hypothesis is that all regressors are exogenous which can be rejected if the test statistics are sufficiently large.

In all specifications in Table 1.2, our main results from Table 1.1 remain robust and the coefficient estimates change in only minor ways. In addition, the Cragg-Donald Wald F-statistics are much greater than 10, the standard threshold for weak instruments, which suggests that the chosen instruments are not weak (Staiger and Stock, 1997). The DWH endogeneity test value is smaller than the critical level, showing that we fail to reject the null that all regressors are exogenous. Overall, the Table 1.2 results suggest that our main conclusions are robust to the potential simultaneity problem in the regular transactions share.

1.4.3 Robustness and sensitivity analysis

Using a basket of currencies

In Tables 1.1 and 1.2 we control for the price of ETH in USD (in columns 1 and 2) or use the marginal gas price in USD as the dependent variable (in columns 3 and 4), to incorporate the possibility that the users' choice of gas price may be affected by the real cost of the transaction fee, in terms of the users' local conventional currency. In this section we extend this analysis and check the sensitivity of our results by considering the price of ETH in terms of a basket of currencies (BoC), as opposed to USD only. Specifically, we take the local currencies of the three geographic locations with the largest numbers of active Ethereum nodes (USA, China and the Euro-zone, see Kim et al., 2018) and use the share of active nodes from each location as weights for the respective currencies in the basket, USD, CNY and EUR (see Appendix A.2 for a detailed description).

Table 1.3: Marginal gas price results – Basket of currencies (BoC)

	marginal gas price, ETH		marginal gas price, BoC	
	(1)	(2)	(3)	(4)
blockchain utilization, B_t	2.263***	0.085	3.322***	0.898*
	(0.624)	(0.306)	(0.795)	(0.475)
utilization > 90%, TR_t	-	33.01***	-	27.83***
	-	(4.043)	-	(4.854)
regular transactions share, R_t	4.421***	2.489***	8.206***	6.611***
	(1.003)	(0.665)	(0.796)	(0.745)
ETH price (BoC), X_t	-0.570***	-0.504***	-	-
	(0.166)	(0.127)	-	-
sample size	457	457	457	457
R-squared	0.291	0.593	0.646	0.759

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable, "marginal gas price" is the natural logarithm of the minimum observed gas price (in 10^{-8} ETH or 10^{-5} BoC, weighted basket of currencies consisting of USD, EUR, and CNY) over all transactions in a block, averaged over all blocks created on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price (BoC)" is the daily average ETH price in terms of a basket of currencies consisting of US Dollars, Euro, and Chinese Yuan, weighted by the country/region specific number active nodes (see Appendix A.2). *, **, *** denote 10%, 5%, and 1% significance level, respectively.

Tables 1.3 and 1.4 repeat the analysis from Tables 1.1 and 1.2, using the price of ETH in terms of the basket of USD, CNY and EUR currencies (BoC). Our main results are essentially unchanged. The relationship between blockchain utilization and the marginal gas price is positive and strongly non-linear. The regular transactions share remains statistically significantly positively associated with the gas price in all specifications. The estimates on ETH price (BoC) in columns (1) and (2) show that the gas price is negatively associated with the ETH value in terms of the basket of currencies.

Table 1.4: Instrumental variables (IV) results – Basket of currencies (BoC)

	marginal gas price, ETH		marginal gas price, BoC	
	(1)	(2)	(3)	(4)
blockchain utilization, B_t	2.266***	0.112	3.168***	0.844*
	(0.624)	(0.302)	(0.927)	(0.470)
utilization > 90%, TR_t	-	32.52***	-	26.79***
	-	(4.348)	-	(5.726)
regular transactions share, R_t	5.541***	2.903***	8.772***	7.072***
	(1.493)	(0.968)	(1.417)	(1.212)
ETH price (BoC), X_t	-0.762***	-0.571***	-	-
	(0.228)	(0.151)	-	-
sample size	456	456	456	456
F-statistic (C-D weak instrument test)	267.5	239.8	327.7	282.0
χ^2-stat (DWH endogeneity test)	2.367	0.659	0.378	0.399
R-squared	0.288	0.592	0.645	0.758

Notes: Instrumental variables (IV) regressions with daily data including a constant. Newey-West standard errors reported in the parentheses. The first-stage regresses the regular transactions share on its first lag and the total number of new sender accounts and on the blockchain utilization and the price of ETH. The dependent variable "marginal gas price" is the natural logarithm of the minimum observed gas price (in 10^{-8} ETH or 10^{-5} BoC, weighted basket of currencies consisting of USD, EUR, and CNY) over all included transactions in a block, averaged over all blocks created on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price (BoC)" is the daily average ETH price in terms of a basket of currencies consisting of US Dollars, Euro, and Chinese Yuan, weighted by the country/region specific number active nodes (see Appendix A.2). *, **, *** denote 10%, 5%, and 1% significance level, respectively.

Alternative marginal gas price definition

The marginal gas price used in our baseline analysis in Table 1.1 was defined as the daily average of the *minimum* gas price observed in each block recorded on that day. A possible concern could be that the minimum gas price in a block may be unusually low, e.g., because of miner error or other reasons. This happens very rarely in our data, since executing transactions with very low gas price is not profitable for the miners. Still, to address potential concerns with using the minimum gas price in a block, we instead construct an alternative marginal gas price variable defined as the daily average of the *lowest 5-th percentile* gas price in each block recorded on that day. In this way any extreme or outlier gas price values are avoided, while we still focus on transactions that are at the margin or very close to the margin of being included vs. not included in a block.

Table A.1 shows descriptive statistics for the lowest 5-th percentile of the gas price (daily average). As expected, its mean is slightly higher than that of the baseline marginal gas price definition (the block minimum gas price) used in Table 1.1, but otherwise their distributions are very similar. Table 1.5 reports estimation results using this alternative marginal gas price definition. We do not observe any notable change in our main results when using the lowest 5-th percentile gas price as the dependent variable. Blockchain utilization has a positive and statistically significant estimate in columns (1) and (3) as before. The relationship between blockchain utilization and the lowest 5-th percentile gas price is strongly non-linear, reflected in the large positive and statistically significant threshold coefficients in columns (2) and (4). Likewise, the share of regular transactions has positive and significant coefficient estimates in all specifications. The negative and statistically significant estimates in columns (1) and (2) shows that the ETH gas price remains negatively associated with the price of ETH in USD, as in Tables 1.1 and 1.2. These results show that our main results are not sensitive to outliers in the marginal gas price.

Median gas price

Above we focused on the gas price for transactions on the margin of being included vs. not included in a block. We now look instead at the median transaction in terms of the gas price distribution within a block. We define the *median gas price* as the daily average of the median gas prices in each block recorded on a given day.[25]

Table A.1 reports summary statistics for the median block gas price. Its mean value over the studied period is more than twice larger than that of the marginal (block minimum) gas price but its distribution is right-skewed, similar to that of the marginal gas price. Table 1.6 reports regression results using log of the median gas price as the dependent variable. The results are broadly consistent with the theory hypotheses H1 and H2 and our baseline results for the marginal gas price, except for a few differences explained below. The estimate on blockchain utilization is positive and statistically significant and the non-linear 90% threshold effect is still present and large in magnitude. The regular transactions share is positively and statistically significantly associated with the median gas price.

[25]We do not study the mean gas price as it can be heavily influenced by outliers, e.g., abnormally high gas prices. See, for example, cryptonews.com/news/ethereum-transaction-fee-mystery-just-got-more-mysterious-6816.htm.

Table 1.5: Alternative marginal gas price definition, lowest 5-th percentile

	lowest 5-th percentile, ETH		lowest 5-th percentile, USD	
	(1)	(2)	(3)	(4)
blockchain utilization, B_t	2.239***	0.123	3.324***	0.957*
	(0.619)	(0.305)	(0.795)	(0.485)
utilization > 90%, TR_t	-	32.40***	-	27.18***
	-	(4.071)	-	(4.942)
regular transactions share, R_t	4.587***	2.703***	8.558***	7.000***
	(0.980)	(0.670)	(0.786)	(0.759)
ETH price in USD, X_t	-0.663***	-0.584***	-	-
	(0.204)	(0.160)	-	-
sample size	457	457	457	457
R-squared	0.314	0.598	0.658	0.761

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable "lowest 5-th percentile" gas price is the natural logarithm of the bottom 5-th percentile of gas prices (in 10^{-8} ETH or 10^{-5} USD) in each block, averaged across all blocks created on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price in USD" is the daily average ETH price in 10^3 USD. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

One difference from Table 1.1 when we use the median gas price as the dependent variable is that the price of ETH in USD does not have a statistically significant association with the median gas price. This can be interpreted as the median user being less concerned about the real cost of transaction execution, unlike the marginal user in the left tail of the gas price distribution. Possibly, the median user may be wealthier compared to the marginal user and less sensitive to the transaction fee, or the median user may be more risk-averse or have higher urgency hence bidding a higher gas price to avoid delays. Unfortunately, we do not have data to directly test these possible mechanisms.[26]

Additional robustness checks

We perform four additional robustness checks of our main results.

Quadratic model. First, in Table A.4 in the Appendix we consider a quadratic specification for the effect of blockchain utilization on the gas price, the term $f(B_t)$ in equation

[26]The reason for posting a marginal gas price vs. median gas price transaction could differ too - e.g., a non time-sensitive advertising or faucet payment from a website vs. time-sensitive transfer or purchase.

Table 1.6: Median gas price

	median gas price, ETH		median gas price, USD	
	(1)	(2)	(3)	(4)
blockchain utilization, B_t	0.790*	-0.581**	2.187***	0.572
	(0.490)	(0.268)	(0.701)	(0.523)
utilization > 90%, TR_t	-	24.38***	-	18.81***
	-	(3.819)	-	(5.047)
regular transactions share, R_t	3.954***	2.596***	9.042***	8.329***
	(0.748)	(0.588)	(0.735)	(0.835)
ETH price in USD, X_t	-0.094	-0.037	-	-
	(0.161)	(0.139)	-	-
sample size	457	457	457	457
R-squared	0.371	0.580	0.676	0.726

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable "median gas price" is the natural logarithm of the median gas price (in 10^{-8} ETH or 10^{-5} USD) in each block, averaged across all blocks created on date t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price in USD" is the daily average ETH price in 10^3 USD. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

(1.3). We demean the blockchain utilization variable to avoid multicollinearity issues and for easier interpretation of the coefficient on the squared term. Consistent with our main results in Tables 1.1, 1.5 and 1.6 we find that blockchain utilization is positively and statistically significantly associated with the gas price, with a strong non-linear effect reflected in the large coefficient on squared utilization. The regular transactions share estimate remains positive and statistically significantly different from zero at the 1% significance level. The price of ETH in USD has a negative estimate which is statistically significant, except for the median gas price, as in Table 1.6.

Other alternative specifications. In Appendix Table A.5 we re-estimate our main specifications from Table 1.1 by using as dependent variable the marginal gas price in levels instead of logs. Our main results remain robust – there is a significant positive non-linear relationship, with a strong threshold effect above 90% utilization, between the blockchain utilization and the marginal gas price (in ETH or USD). The results about the regular transactions share and the ETH price in USD (in columns 1 and 2) also remain robust.

In Appendix Table A.6 we re-estimate equation (1.3) by additionally including gas supply, defined as the sum of the gas limits of all blocks created in a given day. As in all tables, we control separately for the system-wide blockchain protocol events (jumps) affecting the block gas limit or difficulty via dummy variables. The gas supply estimate not statistically significantly distinguishable from zero which confirms that, once we account for the system-wide events, the blockchain gas supply is stable and any remaining minor fluctuations in it do not affect the gas price on average. Our main results, on the non-linear positive association of blockchain utilization with the gas price, the positive association of the regular transaction share with the gas price, and the negative relationship of the ETH price in USD with the gas price remain robust and very close in magnitude to the baseline estimates in Tables 1.1 and 1.2.

Finally, in Appendix Table A.7 we re-estimate the ETH specifications with a non-linear utilization effect (column (2) in Tables 1.1 and 1.2) by using the first lag of the ETH price in USD which allows for potential delays in the users' transaction posting behavior or bid gas prices in reaction to changes in the ETH/USD exchange rate.[27] Our results remain essentially unchanged.

1.4.4 Are regular transactions more urgent?

We examine further our results about the regular transactions share and its positive correlation with the observed gas prices. Specifically, we hypothesize that regular transactions are more urgent and associated with higher waiting costs. We construct a variable which measures the average position of regular transactions within a block, defined as follows. Suppose there are n regular transactions in a block containing N total transactions, where $N > n > 0$. Define the average regular transaction position (RTP) as:

$$\text{RTP} = 1 - \frac{\sum\limits_{i \in \text{regular}} pos_i/n - \frac{n+1}{2}}{N - n}$$

where $pos_i \in [1, N]$ denotes the position of regular transaction i in its block relative to the block middle. Remember from Section 3.1 that transactions are ordered in descending gas price order within each block; that is, more urgent transactions, with higher gas prices, are recorded nearer the top of the block. By construction, the RTP measure takes values between 0 and 1. A RTP closer to 1 (larger than 1/2) means that regular transactions are recorded on average nearer the top (beginning) of the block, i.e., their associated gas prices are higher.[28] The opposite is true if RTP is closer to 0 (smaller than 1/2). Figure

[27] For example, many new users joined Ethereum during the peak period when the USD price of ETH went up significantly (see Figure 1.10 and Figure A.1 in the Appendix).

[28] For example, suppose there are N=6 transactions in a block, n=3 of which are regular transactions. If the positions of the regular transactions are 1, 2 and 4 (i.e., toward the beginning/top of the block), then the RTP measure equals $1 - (7/3 - 2)/(6 - 3) = 0.89$ which is larger than 0.5. Similarly, if the positions of the regular transactions in the block were instead 3, 5 and 6 (i.e., toward the end/bottom of the block), then RTP equals $1 - (14/3 - 2)/(6 - 3) = 0.11$.

1.9 plots the daily average RTP value. The Figure confirms that regular transactions are indeed located, on average, in the upper (higher gas price) half of the Ethereum blocks.[29] This supports our hypothesis that regular transactions are more urgent on average and likely to be associated with higher waiting costs and justifies using the regular transactions share, R_t as proxy for transaction urgency in the estimation equation (1.3).

Figure 1.9: Position of regular transactions within blocks

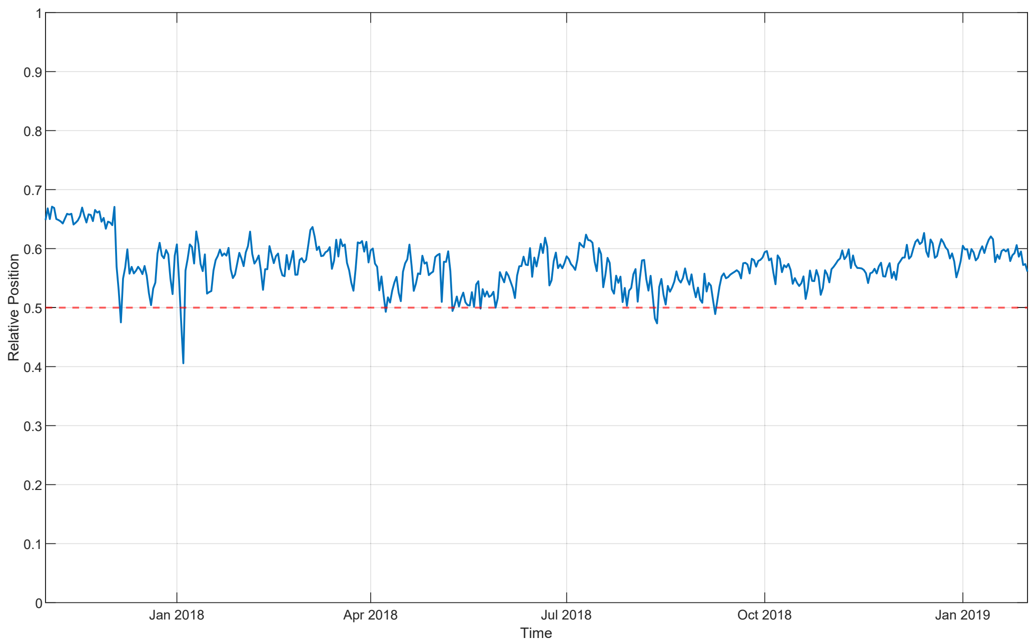

Notes: Position of regular transactions within blocks. Regular transaction position (RTP) within blocks, daily average. RTP close to 1 (0) indicates that regular transactions are located close to the top (bottom) of the block and their associated gas prices are higher (lower). The dashed line indicates the mid-block position, RTP = 1/2.

We also checked whether there is any bunching in the gas prices of regular transactions, since blockchain users often use special software (“wallet") to automate transaction execution and set the transaction fee. If the software systematically picks relatively high gas price values, and if many regular transaction are posted by such wallet users, then this could be an alternative explanation for the observed upper-half relative position of regular transactions in their blocks depicted on Figure 1.9. Direct inspection of the data confirms that there is no noticeable bunching in the gas prices for regular transactions.

1.4.5 Peak period

In this section we analyze a specific three-month sub-period of our sample, which we call the “peak period”, between November 1, 2017 to January 31, 2018. This is a time in which

[29]The mean and median RTP in our data are both equal to 0.57.

the USD price of ETH increased sharply and the blockchain utilization was very high. Figure 1.10 shows that the ETH price in USD went up by 400% during this period. In addition, Figure 1.2 shows a substantial increase in blockchain utilization. This period is also associated with a large and growing number of new Ethereum accounts/addresses (more than 10,000 per day), part of which could be because of new users joining and transacting at least once on the Ethereum platform, see Figure A.2 in the Appendix. Furthermore, the peak period exhibits sharp oscillations in the blockchain conditions: from an individual user's perspective the average per unit transaction cost could vary by more than 200% from one day to the next, see Figure 1.1. Table A.1 in the Appendix reports summary statistics for the main blockchain variables during the peak period. There is a noticeable difference in the gas price levels compared to the full sample – both the marginal and median gas prices set by users are much higher during the peak period. In addition, the share of regular transactions is considerably larger during the peak period compared to the full sample.

Figure 1.10: Average ETH price in USD

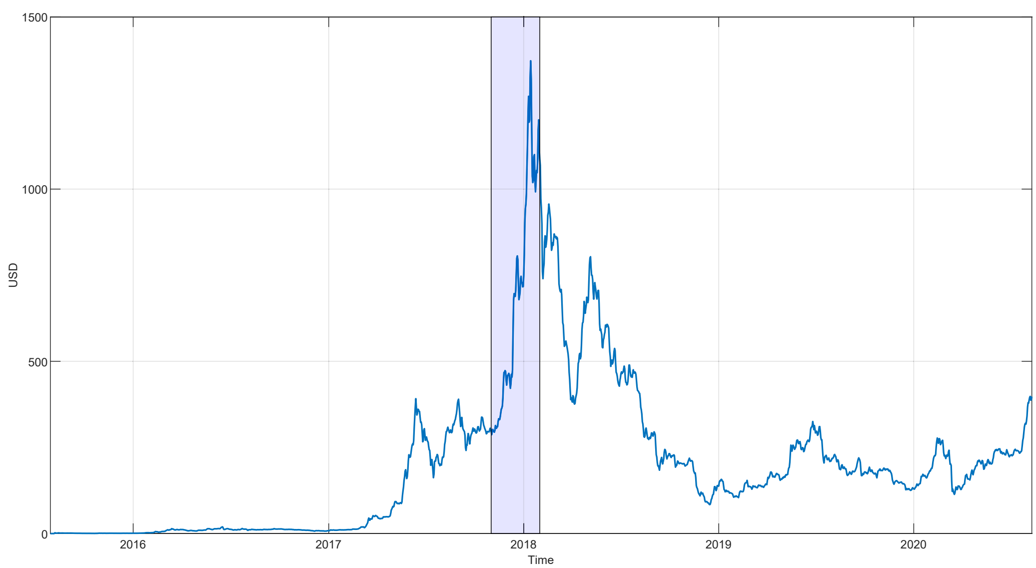

Notes: Average Ether (ETH) price in United States dollar (USD). The figure plots the daily average ETH price in USD. The shaded region highlights the peak period between November 1, 2017 and January 31, 2018.

Because the peak period looks different from the full sample in several dimensions, we test the robustness of our main results using the peak period data only. This helps us evaluate whether there are any significant omitted factors determining gas prices in the peak period or whether the observed high gas prices and utilization rate were mainly driven by demand and the other factors we account for, as in the full sample. We run regressions using the specifications in Table 1.1 for each of the three gas price measures (marginal, lowest 5-th percentile and median). Table 1.7 and Appendix Table A.8 summarize the results. Columns (1) and (2) report results using the daily average of the marginal gas price as the dependent

Table 1.7: Peak period (November 2017 - January 2018)

	gas price in ETH					
	marginal		lowest 5-th percentile		median	
	(1)	(2)	(3)	(4)	(5)	(6)
blockchain utilization, B_t	2.787**	0.534	2.921***	0.858	1.519**	0.626
	(1.122)	(0.741)	(1.039)	(0.695)	(0.681)	(0.579)
utilization > 90%, TR_t	-	17.43***	-	16.19***	-	9.554***
	-	(3.097)	-	(2.998)	-	(2.702)
regular transactions share, R_t	3.419**	1.185	3.243**	1.194	1.972*	1.035
	(1.416)	(0.933)	(1.333)	(0.863)	(1.014)	(0.769)
ETH price in USD, X_t	-0.336	0.057	-0.286	0.072	0.346	0.480**
	(0.483)	(0.324)	(0.451)	(0.299)	(0.296)	(0.235)
sample size	92	92	92	92	92	92
R-squared	0.642	0.854	0.684	0.867	0.733	0.819

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable "marginal gas price" is the natural logarithm of the minimum observed gas price (in 10^{-8} ETH) over all included transactions in a block, averaged across all blocks created on day t. The dependent variable "lowest 5-th percentile" gas price is the natural logarithm of the bottom 5-th percentile of gas prices (in 10^{-8} ETH) in each block, averaged across all blocks created on day t. The dependent variable "median gas price" is the natural logarithm of the median gas price (in 10^{-8} ETH) in each block, averaged across all blocks recorded on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price in USD" is the daily average ETH price in 10^3 USD. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

variable while columns (3)-(4) and (5)-(6) report results using the lowest 5-th percentile gas price and the median gas price, respectively.

The results in Tables 1.7 and A.8, using gas prices in ETH and USD respectively, show that the peak-period results remain broadly in line with our full-sample findings, although the peak period estimates are noisier. Blockchain utilization still has a positive and strongly non-linear association with the gas price overall, however the standard errors are larger than in the full sample results. A larger regular transactions share is positively associated with higher gas prices in columns (1), (3) and (5), as in the baseline results, but the estimates are not statistically significantly different from zero in the specifications with a non-linear utilization threshold, columns (2), (4) and (6). The regular transactions share estimates are statistically significantly positive in all columns of Appendix Table A.8, when the gas price is measured in USD. The estimate on the ETH price in USD in Table 1.7 is not significantly different from zero, except in column (6), which could be explained by the possibility that the sharply increasing ETH price in USD in the peak period may have been seen as a (speculative) investment or trading opportunity by many users, despite the high associated transaction costs.

1.5 Conclusion

Blockchain cryptocurrency platforms like Ethereum are unique examples of financial markets that, because of their decentralized and anonymous nature, remain largely unregulated by official authorities. Furthermore, these digital markets lack external formal contract enforcement, beyond the algorithmic protocol and computer code of the platform itself (see Karaivanov, 2021 and Townsend, 2020 for further discussion). There has also been a lot of debate on the extreme volatility, speculation activity and illicit transactions on these platforms (for example, Foley et al., 2019; Griffin and Shams, 2020; Li et al., 2020). It is therefore a valid empirical question whether blockchain platforms like Ethereum operate as financial markets for payments and related services, subject to standard demand and supply and other economic factors.

We find that the answer is by and large affirmative. We analyze transaction fees (gas prices) in the Ethereum blockchain and find that demand factors, measured by block utilization rates, and the choice of transaction type are the primary economic determinants of observed fees. Blockchain utilization has a positive association with the marginal and median gas price which is strongly non-linear above the 90% block utilization threshold. Additionally, we find that blockchain users endogenously vary the mix of regular transfers vs. smart contract transactions, with a larger fraction of regular transactions observed when the gas price is high.

We abstracted from studying the determinants of the ETH price in USD, instead treating it as a potential factor affecting gas prices. Further research on the joint determinants of blockchain transaction fees and the ETH/USD exchange rate could be beneficial. In addition, more research on the network structure of Ethereum blockchain addresses and on the heterogeneity of recorded transactions (by size, sender address, or recipient address) can provide additional useful insights on the economics of blockchain platforms.

Chapter 2

Individual Evolutionary Learning and Zero-Intelligence in the Continuous Double Auction

2.1 Introduction

The Continuous Double Auction (CDA) is a centralized market in which buyers and sellers are free at any time to make bids and asks and to accept the bids and asks of others. The open book CDA is the preferred exchange mechanism of financial markets around the world, and in particular, of stock exchanges (NYSE, Euronext, LSE, NASDAQ, etc.).

Trader behavior and market performance in a controlled double auction setting were first examined by Smith (1962). Following his pioneering work, in thousands of experiments many other economists have documented features of the CDA some of which are summarized in Holt (1995) and Plott (2008). For example, it is well known that, in later periods, prices and allocations converge to their competitive equilibrium values. But, in the first periods of these experiments, competitive equilibrium is not reached immediately. See, for example, Plott (2008, p.16) where he exposits the “sawtooth" property of transaction price equilibrium.

What happens in the first period of these experiments is particularly important if one is to better understand how a CDA generates price discovery. One can think of the first period as the immediate time after a change in demand-supply conditions creates a disequilibrium situation. In the long-run, the market may settle down on a single competitive equilibrium price. But in the short-run, the prices and allocations are rarely efficient. It is important to understand how price discovery occurs in this information sparse environment.

There is no generally accepted theory about the dynamics of price formation *within the first period* that has survived experimental testing. Cason and Friedman (1996) tested three theories of within period dynamics: the sophisticated Bayesian equilibrium of Wilson (1987), the simpler Bayesian Game Against Nature (BGAN) of Friedman (1991), and the Zero-Intelligence traders (ZI) of Gode and Sunder (1993). They found that “none of these three models adequately explains price formation in double auction markets.” (p. 1333) But they gave more support to the ZI model. “ZI now seems the only natural source for a null hypothesis in assessing the performance of any more complicated mechanism.” (p. 1333)

Lin et al. (2020) also tested the three theories using a much larger data set than that of Cason and Friedman (1996). The data were obtained from more than 9,000 market periods in classroom experiments conducted in various countries using the MobLab platform.[1] Their conclusions supported the Cason and Friedman (1996) findings. They rejected the Wilson model and BGAN, and accepted ZI with reservations. “We find much stronger support for ZI theory compared to Cason-Friedman (1993). ...It appears ZI explains dynamics within periods. ... However, a non-negligible portion of our data falls outside of the simulated 95% confidence region."(p.921)

In this paper, we test two theories of the CDA: Zero-Intelligence (ZI) and Individual Evolutionary Learning (IEL). The literature on ZI traders in the double auction is extensive. For instance, see Gode and Sunder (1993, 1997, 2004), and Cliff and Bruten (1997). The IEL agent has been shown to behave like human subjects in a variety of repeated, discrete, synchronous games, such as call markets (Arifovic and Ledyard 2007), voluntary contribution mechanism, VCM, (Arifovic and Ledyard 2012), and repeated battle of the sexes game (Arifovic and Ledyard 2018). But, the CDA is a different type of game: it is continuous and asymmetric. In this paper, we adapt IEL for the CDA.

We use two very different data sets to test the theories. One is a subset of the MobLab data created and used by Lin et al. (2020).[2] These data afford researchers an unusual opportunity to test theories of behavior in CDAs against a very large number of observations. We will use the 2090 observations of the first period of a CDA experiment with the same demand-supply configuration. There are, however, some potential problems with these data as they are generated in classroom experiments without the usual controls for a standard economic experiment. 25% of the periods involve 5 or more trades that lose money. We, therefore, ran some experiments of our own as a comparison. This, the second data set, which we will call the SFU data set, contains 25 observations of the first period of a CDA experiment with the same demand-supply configuration as the MobLab data set.

In testing the two models, we focus on the two main measures of performance: efficiency and price. We first test ZI and IEL against the MobLab data. We find that IEL generates efficiencies closer to the MobLab efficiencies than does ZI. And, ZI generates prices closer to MobLab prices than does IEL. Neither does very well as a full theory. We then considered a mixed pool of agents: part IEL and part no intelligence (NI). NI traders generate bids and asks randomly from a fixed set. They are related to, but not the same as, Gode-Sunder’s ZI-U (ZI-unconstrained) agents. We find that a mixture of 70% IEL and 30% NI provides

[1]MobLab is an educational software company that provides a platform for carrying out in-class experiments in economics courses. Full disclosure: Ledyard is on the Board of Directors of MobLab.

[2]The data are publicly available at https://osf.io/9mfws/.

a close fit to the MobLab data in both efficiency and price. That mix outperforms ZI. We also find that these results are robust to variations in the two basic parameters of IEL.

We then test ZI and IEL against the SFU data. We find that IEL generates efficiencies closer to the SFU efficiencies than does ZI. And, ZI generates prices closer to SFU prices than does IEL. Although neither theory is exactly correct, if we weigh efficiency and price equally, IEL out performs ZI. We also find that adding NI into the mix with IEL does not improve the performance. And, the results are robust to variations in the two basic parameters of IEL.

2.2 Experimental design and measurements

All of the experiments and simulations in this paper are run within the context of the same environment and market design.

2.2.1 The environment

The general structure of the environment used in this paper is standard. There is one commodity which is available in integer amounts. There are N traders. Each trader is either a buyer or a seller but not both. K is the maximum number of units a buyer may buy or a seller may sell. A buyer's/seller's payoff at the end of the market is based on a vector of (marginal) values/costs. For buyer i, the values are $V^i = (V_1^i, ..., V_K^i)$ where $V_1^i > V_2^i > ... > V_K^i$. For seller i, the costs are $C^i = (C_1^i, ..., C_K^i)$ where $C_1^i < ... < C_K^i$.

If buyer i buys r units of the good during the market period and pays prices $(p_1, ..., p_r)$ for them, the payoff will be $U(r; V^i, p) = \sum_{k=1}^{r}[V_k^i - p_k]$. If seller i sells r units of the good during the market period and receives prices $(p_1, ..., p_r)$ for them, the payoff will be $U(r; C^i, p) = \sum_{k=1}^{r}[p_k - C_k^i]$.

The specific environment used in this paper is straight-forward. There are 5 buyers and 5 sellers. Each buyer can buy up to 3 units. Each seller can sell up to 3 units. There is no resale and no buy back. The values for each buyer and costs for each seller yield the demand-supply configuration found in Figure 2.1. There are multiple possible competitive equilibria. The competitive equilibrium set of prices is $[110, 114]$. The competitive equilibrium quantity is 13. At any equilibrium, there are 13 infra-marginal and 2 extra-marginal units.

One important feature of this design is the asymmetry between the buyers' surplus and the sellers' surplus in equilibrium. The sellers' surplus is larger. We will see later in this paper that this has an effect on the dynamics of transactions prices.

2.2.2 The CDA market design

All of the experiments and simulations we consider in this paper use a standard CDA design. There is a public book - a queue of the bids and asks that have been made but not accepted.

Figure 2.1: Demand-Supply configuration

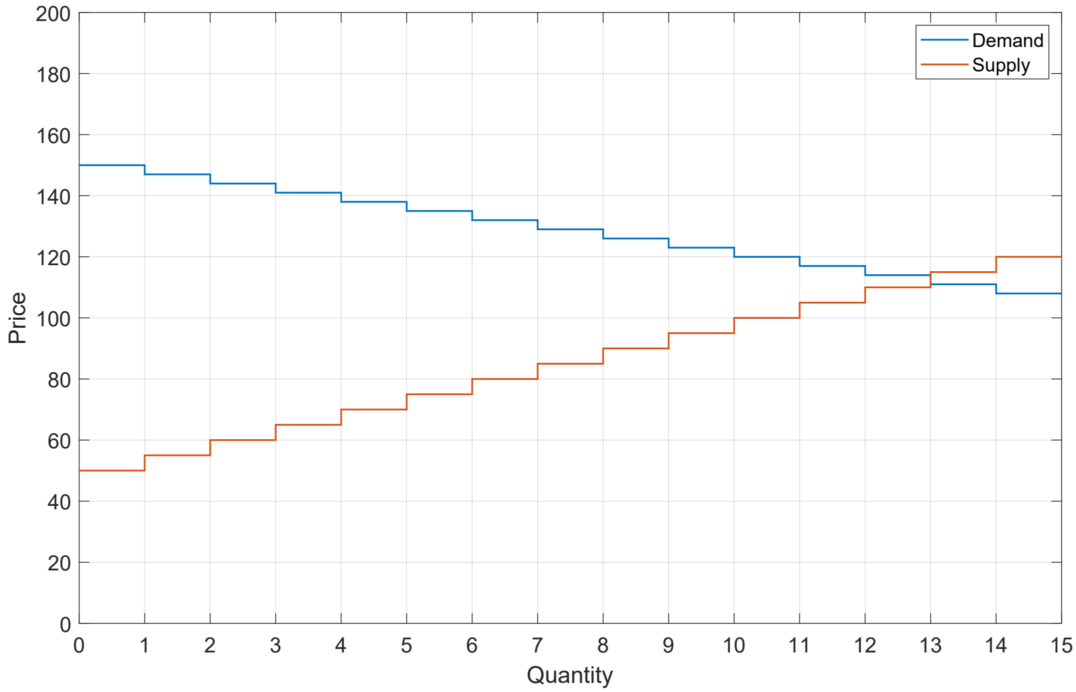

Traders, at anytime, can add orders to the book or cancel any of their orders in the book.[3] There is no resale. A trader is either a buyer or seller. A buyer can always accept the lowest ask in the book. A seller can always accept the highest bid in the book.

Orders are restricted to be for only 1 unit at a time. Bids and asks are restricted to the interval $[s_l, s_u]$. For this paper, $[s_l, s_u]$ always is equal to $[0, 300]$. There is no requirement that bids be less than a buyer's value for the item or higher than a seller's cost for the item. There is also no spread reduction rule; that is, a new bid does not have to be lower than the highest bid in the book and a new ask does not have to be higher than the lowest ask in the book.

2.2.3 Key observables

The two key performance measurements for CDAs are *Efficiency* and *Price*. Efficiency measures how much of the surplus is captured by the traders through the final allocation. Price determines how that surplus is distributed among the traders. Any good theory must explain both efficiency and price.

We use four measures of efficiency related performance.

[3]We try to be consistent with our language in this paper. Buyers and sellers submit *orders*. A *bid* is an order submitted by a buyer. An *ask* is an order submitted by a seller.

1. *Efficiency* is the surplus attained in a period divided by the maximum possible surplus in that period. Efficiency is

$$E = \frac{\sum_i \pi_i}{\sum_i \pi_i^C}$$

   where π_i is trader i's profit in the experiment and π_i^C is trader i's profit in the competitive equilibrium. This is standard.

2. *Quantity* is the number of completed trades in a period. This is standard.

3. Inefficiency in the CDA has two potential causes: (i) infra marginal units that are not traded and (ii) extra marginal units that are traded. The *Inefficiency Ratio* is a measurement designed to identify the sources of inefficiency in a CDA. The measurement was developed by Cason and Friedman (1996) and used by Lin et al. (2020). It measures the percentage of the surplus lost due to traded extra marginals. If it is low, then most of the loss in efficiency is due to untraded infra marginals.

4. The *Buyers' Correlation coefficient* measures the correlation between the rank of the value (from high to low) of an item and the order in which it traded. The *Seller' Correlation coefficient* measures the correlation between the rank of the cost (from low to high) of an item and the order in which it traded. We use Spearman's ρ. Higher correlations can lead to higher efficiency since there is then a higher probability that infra-marginal units will trade first and a lower probability that extra marginal units will trade at all.

We use three measures of price related performance.

1. *Price* is the average price of all trades.

2. The *Buyers' Profit Split (BPS)* is the percent of the profits achieved by the buy side.

3. We measure *price volatility* using α from Smith (1962).

$$\alpha = \sqrt{\frac{1}{J} \sum_j \left[\frac{(P_j - P^C)}{P^C} \right]^2}$$

   where J is the number of trades, P_j is the price of the jth trade, and P^C is the competitive equilibrium price. If there is a range of competitive equilibrium prices, we use the mid-point of the range for P^C.

2.3 The MobLab data

In the MobLab data set, there are 2090 observations for the first period in the environment described in Figure 2.1 above. We display the averages of the key observables for these data in Table 2.1.

Table 2.1: MobLab experimental data: averaged over 2090 observations

Efficiency	85.9
Quantity	11
Inefficiency Ratio	0.49
Buyers' Correlation	0.56
Sellers' Correlation	0.57
Average Price	113.2
Buyers' Profit Split	0.36
Price Volatility	0.26
Orders per period	55

There is a relatively high average efficiency of almost 86%. The inefficiency ratio of 0.49 indicates the 14% loss in efficiency is due almost equally to infra marginal units going untraded and extra marginal units being traded.

The relatively high correlation coefficients indicate that high values and low costs are usually trading early which helps efficiency. One should, however, temper this observation with the fact that, for example, each buyer must trade their 2 more valuable units before being able to trade their third and least valuable. So there is perfect correlation within a trader. This could lead to high correlation coefficients across traders even if their bidding order is entirely random.

The average price of 113.2 is a competitive equilibrium price. In fact, we can say something stronger. In Table 2.2 we display the average prices of each trade in the the order of that trade; that is, the average price of the first trade, the average price of the second trade and so forth. The average price of each of those trades is an equilibrium price. The MobLab buyers' profit split is almost exactly equal to the buyers' profit split at the competitive equilibrium price of 113.2.

Table 2.2: MobLab data price dynamics: average price of the first 10 trades

1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
111.6	113.2	113.4	113.7	113.7	114.6	112.8	112.3	113.6	112.1

That trade is occurring on average at a competitive equilibrium price, no matter the order of trade, is somewhat surprising. A feature of the Demand-Supply configuration we are using is the asymmetry between the buy side and the sell side. According to Holt (1995), summarizing observations from many well controlled experiments, "price tends to converge to the competitive equilibrium from below if producer surplus exceeds consumer surplus at the competitive equilibrium price, and from above in the reverse situation."(p. 372). We will explain the reasons for the difference between Holt and these data later in this paper.

Ordering activity is not very high. 55 orders are used to generate 11 trades. Each trade requires, a minimum of at least 2 orders: a buy and an ask. There are 2 1/2 times the number of orders than is minimally necessary. The split of the orders between bids and asks is almost equal. 47.7% of the orders are bids.

While the number of observations of the MobLab data set is truly impressive and the data seem consistent with conventional wisdom about experimental CDAs, there is a potential problem. MobLab experiments are usually run in classrooms and not necessarily subjected to the usual rigorous control of the experimental economics laboratory. Subjects might be paid in course points or some other way but usually are not compensated with money in the standard way. The particular compensation scheme for each observation is not identified in the data set. Further this may be the first, and only, market experiment that these students participate in. So the subjects, and perhaps the instructor running the experiment, are mostly inexperienced.

In the 2090 observations we are using, there are many violations of an individual rationality constraint; i.e., a buyer pays more than their value or a seller accepts less than their cost. In the 2090 periods we are using, 90% have at least one trade that involves a violation. 25% have at least 5 such trades,[4] almost half of the average number of trades made. The most telling statistic is that 4.3% of the 2090 periods have an *average* price greater 150 – the most that any buyer should be willing to pay.

Lin et al. (2020) are clearly aware of this issue and devote some space to discussing it. They conclude that "a limited analysis across different types of incentives offered by 10% of the instructors suggests different levels of incentives do not create substantial differences in behavior."(p. 918) At this point in our analysis we accept their assumption and proceed to examine whether the IEL behavioral model explains the data.

2.4 Theories

We describe the theories behind Individual Evolutionary Learning (IEL) and Zero Intelligence (ZI) in the standard Demand-Supply environment and CDA design.

Any trader in a CDA has to decide (a) when to order, (b) what to order, (c) when to cancel an order, and (d) when to accept an order in the book. In many CDA market designs, such as the ones in this paper, the way to accept the best ask in the book is to submit a bid higher than that ask. The market then crosses such a bid at the price equal to the ask. Therefore, we ignore (d) as being covered by (a) and (b).

We begin with the ZI theory, since it is easier to understand.

[4]But, many such trades still increase aggregate surplus. The loss on one side of the trade is covered by the gain on the other. The "irrational" behavior does not seem to affect efficiency.

2.4.1 ZI in the CDA

ZI was introduced in Gode and Sunder (1993) for environments where buyers each only wanted 1 unit and sellers each only had 1 unit to sell. We modify their model slightly to fit the environment with 3 units per trader.

When to order Ordering occurs in iterations. At each iteration, a trader is drawn at random, with replacement, and bids if a buyer or asks if a seller.

What to order The selected trader bids uniformly from $[s_l, V_{k+1}]$ if a buyer or asks uniformly from $[C_{k+1}, s_u]$ if a seller.

When to cancel an order The ZI trader cancels their order in the book, if they have one, whenever they are about to submit a new order.

Each trader is permitted three transactions only and is not allowed to re-trade any item bought or sold.

The ZI process continues for a pre-specified number of computer draws.

2.4.2 IEL in the CDA

In our previous work, IEL agents participated in a synchronous, repeated game. Each IEL agent was defined primarily by a set of possible strategies and a foregone utility function. With these, an IEL agent generated an action in each round of the game. For example, in a repeated call market, the agent decided what bid to make in each round. In a CDA, which is an asynchronous continuous game, we need to specify a little bit more.

An IEL agent in a CDA is similar to a ZI agent. "When to order" and "when to cancel an order" are the same as in ZI. "What to order" is different. The IEL agent selects what to order in the same way it has selected strategies in the other Arifovic-Ledyard papers.[5] We describe the process for a buyer. A seller is symmetric.

- Each agent maintains a finite set of remembered potential bids, A.

  The remembered set is a collection of J numbers from $[s_l, s_u]$, the set of admissible orders.[6]

  At any instant in time t, an IEL buyer will have purchased k(t) units of the commodity. The value to them of the next unit they buy is $V_{k(t)+1}$. Since bids greater than $V_{k(t)+1}$

[5] We assume the reader has some familiarity with our previous work with IEL agents. For those who do not, a more detailed description of IEL and its applications is in Appendix B.3.

[6] At the beginning of a period, we initialize the set A by randomly selecting J items from $[s_l, V_1^i]$ for buyer i and from $[C_1^j, s_u]$ for seller j.

are certainly going to lose money, we restrict the remembered set of a buyer to be a subset of $[s_l, V_{k(t)+1}]$. For a seller, we restrict the remembered set to $[C_{k(t)+1}, s_u]$.

When an IEL buyer accepts an ask or has a bid accepted, the acceptable set $[s_l, V_{k(t)+1}]$ becomes $[s_l, V_{k(t)+2}]$. At this time there may be bids b in the remembered set for which $V_{k(t)+2} < b < V_{k(t)+1}$. Such bids are sure losers. We remove these from the remembered set. Similarly, when an IEL seller accepts a bid or has an ask accepted, the acceptable set $[C_{k(t)+1}, s_u]$ becomes $[C_{k(t)+2}, s_u]$ and there may be asks o such that $C_{k(t)+1} < o < C_{k(t)+2}$. We remove these.

- Each action in the set A is evaluated by a foregone utility function, u.

  The book at time t is $\{\mathcal{B}_t, \mathcal{O}_t\}$ where $\mathcal{B}_t$ is a set of bids $\{b_t^1, ..., b_t^n\}$ ranked from high to low, and $\mathcal{O}_t$ is a set of asks $\{o_t^1, o_t^2, ..., o_t^m\}$ ranked from low to high.[7] $b_t^1 = B_t < o_t^1 = O_t$. B_t is the best bid in the book at time t. O_t is the best ask in the book at time t.

  If IEL sends a buy bid of $b < O$, there will be no transaction and their payoff will be 0. If they bid $b > O$, they are accepting the offer and their payoff will be $V_{k+1} - O$. Therefore, reacting to the book, a buyer's foregone utility for the potential bid b, when they have completed k buys, is:

$$u(b|k, B, O) = \begin{cases} 0 \text{ if } b \leq O \\ V_{k+1} - O \text{ if } b > O \end{cases} \tag{2.1}$$

  Since the foregone utility function is only used to evaluate entries in the remembered sets, IEL only needs to consider $b < V_{k+1}$.

  Similarly, a seller's foregone utility for a potential ask, when they have completed k sales, is:

$$u(o|k, B, O) = \begin{cases} 0 \text{ if } o \geq B \\ B - C_{k+1} \text{ if } o < B \end{cases} \tag{2.2}$$

  It should be noted that this particular foregone utility function ignores the possibility that, if $b < O$ and is in the book, b might eventually be accepted, yielding a higher profit than $V_{k+1} - O$. Our IEL agent is myopic and non-strategic, and ignores this option value.

- When it is a trader's turn to bid, a bid is randomly selected from A in proportion to its foregone utility, u.

[7] The bid queue can be of different length than the ask queue. $n \neq m$ is possible.

- Before the next bid is chosen, A is updated as follows:[8]
  - Experimentation: with probability μ, we randomly replace an action a in A with one randomly drawn from $[s_l, V_{k+1}]$ using a truncated normal with mean a.
  - Replication: a pairwise fitness tournament eliminates alternatives that would not have done well per the foregone utility, u.

Although the IEL process is the same as we have used in the past, it may seem to be a complicated way to pick a bid. A simple way to understand what the IEL agent is doing is to compare it to a Zero Intelligence (ZI) agent.

2.4.3 IEL vs. ZI

In the implementations in this paper, IEL and ZI are identical in “when to order" and “when to cancel an order". They differ only in “what to order". When it comes time to bid, a ZI buyer generates a bid uniformly distributed on $[s_l, V_{k+1}]$. If $V_{k+1} < O$, the IEL agent does exactly the same.[9]

But if $V_{k+1} \geq O$, there is a significant difference in the behavior of the two types of agent. A ZI agent continues to select a bid uniformly distributed from $[s_l, V_{k+1}]$. *Even though O is a perfectly acceptable ask*, ZI only accepts the best ask with probability $\lambda = \frac{V_{k+1}-O}{V_{k+1}-s_l}$. IEL, on the other hand, updates its remembered set in a way that increases the probability it will accept the best ask. This happens because of the replication process. Consider an example where $V_{k+1} = 100, s_l = 0$, and $O = 50$. Then $\lambda = 0.5$. The foregone utility is $V_{k+1} - O$ if $b > O$ and is 0 if $b \leq O$. In replication, a bid below O will remain in the remembered set with probability $(1-\lambda)^2$. Thus after replication the probability that IEL will accept the offer O is $1-(1-\lambda)^2$. For example, if $V = 100, s_l = 0$, and $O = 50$, then ZI accepts with probability 0.5 and IEL accepts with probability 0.75. If IEL happens to bid below 50, then the next time they are selected to bid they will do replication again, and this time will accept with probability $1-(1-\lambda)^4$ or 0.94. ZI will continue to accept with probability 0.5.

IEL is less patient than ZI.[10] The probability that an IEL agent accepts an ask is higher than that of a ZI agent. And over time the probability that IEL accepts a profitable ask approaches 1. Trading will proceed faster with IEL agents than with ZI agents.

[8]Details on these can be found in Appendix B.3.

[9]For IEL this is a three step process. First the remembered set consists of uniformly randomly chosen elements. Second, since the foregone utility is the same for all bids in the remembered set, it remains essentially composed of uniformly randomly chosen elements, even after experimentation and replication. Third, again since the foregone utility is the same for all elements in the remembered set, the selection of the bid is done uniformly randomly.

[10]There is another type of impatience in a CDA: impatience in offering options to others by adding a bid to the queue. By using the Gode-Sunder process to determine when to bid, this version of IEL does not address this type of impatience.

2.5 Simulations

To discover the performance of IEL and ZI agents in a CDA, we ran a series of simulations. Each run simulated the initial period of one experimental CDA market. There were 10,000 runs for each simulation. This seemed to be more than enough to stabilize the performance measures to within 2 significant digits. The upper bound on bids and asks was set at 300, the same as is used in the MobLab experiments.

Because we use the Gode-Sunder procedure to determine when agents should bid, a key parameter for our simulations is the number of draws; i.e, how many orders are submitted in one period. This parameter is new to our research with IEL and plays an important role in determining the performance of both IEL and ZI in the CDA. In our simulations we used a range of values for the number of draws, in order to explore the impact that it has on performance.

2.5.1 IEL simulations

We used the same parameters for IEL that we have used before; the size of the remembered set is $J = 100$ and the probability of experimentation for any one member of the remembered set is $\mu = 0.033$. We also used other values of J and μ to check for robustness. The details of those additional simulations can be found in Appendix B.2. We will comment as necessary on the effect of alternative parameters, but, as will be seen, there is no substantive change in the key results of this paper across those different parameters.

The values of the performance measurements for our IEL simulations are in Table 2.3. We will focus on the two key measures: efficiency and price. In the column labeled E, we list the average efficiency attained for each number of draws. In the column labeled P, we list the average price attained for each number of draws.

Table 2.3: IEL performance measures

draws	E	Q	ED	ρ_B	ρ_S	I	P	BPS	α
20	48.5	4.5	16.3	0.37	0.34	0.006	103.9	0.49	0.21
25	57.1	5.5	20.8	0.41	0.39	0.012	104	0.49	0.2
30	63.8	6.4	25.3	0.44	0.42	0.02	104.4	0.48	0.2
35	69.4	7.1	29.7	0.46	0.45	0.031	104.5	0.49	0.2
40	74.1	7.8	34	0.48	0.47	0.046	104.8	0.48	0.19
45	77.5	8.3	38.1	0.49	0.49	0.06	105	0.48	0.19
50	80.3	8.8	42	0.5	0.51	0.074	105.1	0.48	0.19
55	82.7	9.2	45.7	0.51	0.53	0.092	105.3	0.48	0.19
60	84.5	9.5	49.4	0.52	0.53	0.108	105.4	0.48	0.19
65	85.7	9.7	52.6	0.52	0.54	0.125	105.5	0.48	0.18
70	87	10	55.9	0.53	0.55	0.14	105.7	0.48	0.18
75	87.9	10.1	58.8	0.54	0.56	0.153	105.8	0.48	0.18
80	88.6	10.3	61.9	0.53	0.57	0.169	105.9	0.48	0.18
85	89.1	10.4	64.3	0.54	0.57	0.183	106	0.48	0.18
90	89.7	10.5	66.6	0.54	0.58	0.193	105.8	0.48	0.18
95	90	10.6	68.9	0.54	0.58	0.2	106	0.48	0.18
100	90.4	10.7	70.5	0.54	0.59	0.205	106.1	0.48	0.18
105	90.5	10.7	72.4	0.54	0.59	0.216	106	0.48	0.18
110	90.9	10.8	73.9	0.54	0.59	0.228	106	0.48	0.18
115	91	10.8	75.6	0.54	0.6	0.225	106	0.48	0.18
120	91.2	10.8	76.8	0.55	0.59	0.231	106.1	0.48	0.18
125	91.2	10.9	78	0.54	0.59	0.235	106.2	0.48	0.18
130	91.4	10.9	79	0.55	0.59	0.236	106.1	0.48	0.18

Notes: E = efficiency, Q = quantity, ED = draw of last trade, ρ_B = buyers' correlation coefficient, ρ_S = sellers' correlation coefficient, I = inefficiency ratio, P = price, BPS = buyers' profit split, α= price volatility.

Observation 1.
The number of draws matters. But, both efficiency and price level off after about 100 draws.

The explanation for Observation 1 is obvious. Remember, each draw represents one order submission to the book. As the number of draws increases, there are more orders and thus more trading. At some point, the number of draws will be large enough that there are few or no mutually beneficial trades left. To determine at what level the artificial cap becomes non-binding, we tracked for which draw the last trade in each run occurred. The results are in the column labeled ED in Table 2.3. For a small number of draws, a significant number are used up in trading. When the number of draws is over 100, only 75 to 90 of those are used, on average, to complete trading. By 100 draws, the cap seems not to be important. All of the measurements have leveled off by then.[11]

Observation 2.

The average efficiency is increasing at a decreasing rate in the number of draws.

The explanation for Observation 2 rests on two facts: (1) trading is Marshallian on average and (2) each trade increases the surplus attained. By Marshallian trading we mean that higher value units generally trade earlier than lower valued units and lower cost units are more likely to trade earlier than higher cost units. While this is certainly true for a single individual with multiple units,[12] it is not necessarily true across individuals. But IEL bids/asks are probabilistically proportional to values/costs. Higher bids and lower asks are probably going to trade earlier. Thus, higher value and lower cost items are probably going to trade earlier. This means that on average, higher surplus trades occur earlier in the period. In the columns labeled ρ_B and ρ_S in Table 2.3, we list the correlation coefficients for buyers and sellers that measure this Marshallian effect. These have the right sign, and are mostly larger than 0.5 confirming that the trading by IEL agents is Marshallian.[13]

With Marshallian trading, since higher surplus trades occur earlier, efficiency increases at a decreasing rate. Marshallian trading also implies that the source of the inefficiency is mostly due to untraded infra marginal units. There are more of these for lower numbers of draws. This is confirmed by the low inefficiency ratio in the column labeled I in Table 2.3.

Observation 3.

(a) The average price is increasing very slightly in the number of draws.
(b) Prices converge toward the competitive equilibrium prices of $[110, 114]$ *from below for each value of draws.*
(c) Buyers and sellers split the profits on average almost equally.

The explanation for Observation 3 rests on three facts: (1) trading is Marshallian on average, (2) buyers and sellers split the surplus on average, and (3) sellers' surplus is larger than buyers' surplus in competitive equilibrium, as can be seen in Figure 2.1.[14] Consider the extreme case of Marshallian trading where the lowest cost and highest value trade first.

[11]In Table B.1, we also display results for 135 to 175 draws in increments of 5. The leveling off is even more obvious. For all numbers of draws greater than 100, the average efficiency is about 91%, the average quantity is about 11, and the average price is about 106.

[12]If all traders traded their first units and only those, efficiency would be 62%. If they all traded only their first 2 units, efficiency would be 94%.

[13]The sellers' correlation is larger than the buyers' correlation. We believe this is related to the asymmetry between values and costs, but do not have a proof of that.

[14]Observation (b) was also noted by Cliff and Bruten (1997) for ZI traders when there is asymmetry in the surplus.

Further suppose they trade at a price that splits the surplus between them. Then the second highest value and second lowest cost trade next and trade at a price that splits the difference. Because of the asymmetry in surplus, the second price will be higher than the first.[15] In the simulations, trading is more probabilistic than this simplistic example, but over 10,000 runs things average out. This can be seen most clearly in Table 2.4. In that table we list the average price of each trade in the order they are traded. For example, the average price of the first trade, when there are 50 draws is 104.4. The average price for the sixth trade is 105.4. It is easy to see that, for all draws, the average price is generally increasing as trade progresses.

Table 2.4: IEL price dynamics: average price of the first 10 trades

	Order of Trade									
draws	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
25	103.4	103.4	104.2	104.5	104.3	104.9	104.6	104.1	106.5	98
50	104.4	103.7	103.7	104.3	104.9	105.4	106.3	106.7	107.7	108.2
75	103.8	103.8	104	104.7	105.1	105.6	106.2	107.3	108	108.8
100	103.8	103.9	104.4	104.7	105.2	105.8	106.4	107.3	108.4	109
125	103.9	103.5	104	104.5	105.2	105.5	106.4	107.4	108.7	109.1
150	103.7	103.6	103.8	104.1	105.1	105.7	106.4	107.2	108.5	109.4
175	103.9	103.8	103.8	104.4	105	105.5	106.6	107.3	108.3	109.3

2.5.2 ZI simulations

Our ZI simulations ran exactly the same as our IEL simulations with only the bid selection rule changed. Each run simulated the initial period of one experimental CDA market. There were 10,000 runs for each simulation. The results can be found in Table 2.5.

Qualitatively, ZI and IEL are similar. The number of draws matters, efficiency increases at a decreasing rate in the number of draws, and price increases slightly in the number of draws. The reason for these observations is the same as that for the IEL results. But quantitatively there are significant differences between IEL and ZI.

Observation 4.
The number of draws matters. But, efficiency does not level off after 100 draws.

Observation 5.
The average efficiency is increasing at a decreasing rate in the number of draws. But it takes more draws for ZI to attain a given level of efficiency than it does for IEL.

[15] In the environment in this paper, these prices would be 100, 101, 102, ... , 112. With 11 trades, the average price would be 105.

Observation 6.
(a) The average price is increasing very slightly in the number of draws.
(b) Average prices are higher for ZI than for IEL.

Table 2.5: ZI performance measures

draws	E	Q	ED	ρ_B	ρ_S	I	P	BPS	α
20	21.9	1.9	13	0.19	0.19	0	108	0.44	0.19
25	27.2	2.3	17	0.24	0.22	0	108.3	0.44	0.18
30	32.3	2.8	22	0.26	0.24	0	108.8	0.43	0.18
35	36.3	3.2	26	0.29	0.27	0	109	0.43	0.18
40	40.5	3.6	30	0.31	0.3	0	109.1	0.42	0.18
45	44.6	4	35	0.33	0.33	0	109.4	0.42	0.18
50	47.7	4.3	39	0.34	0.34	0	109.4	0.42	0.17
55	50.8	4.7	44	0.37	0.36	0	109.4	0.42	0.17
60	53.7	5	48	0.39	0.38	0	109.5	0.42	0.17
65	56.7	5.3	53	0.4	0.4	0.01	109.3	0.42	0.17
70	58.9	5.6	57	0.41	0.4	0.01	109.3	0.42	0.17
75	61.4	5.9	61	0.42	0.42	0.01	109.3	0.42	0.17
80	63.3	6.1	65	0.43	0.43	0.01	109.5	0.42	0.17
85	65.4	6.3	70	0.44	0.44	0.01	109.5	0.42	0.17
90	67.3	6.6	74	0.45	0.45	0.01	109.6	0.42	0.17
95	68.9	6.8	78	0.46	0.46	0.01	109.5	0.42	0.17
100	70.6	7	82	0.47	0.47	0.02	109.6	0.42	0.16
105	71.9	7.2	87	0.48	0.48	0.02	109.7	0.42	0.16
110	73.3	7.4	91	0.48	0.49	0.02	109.6	0.42	0.16
115	74.9	7.6	95	0.49	0.49	0.02	109.6	0.42	0.16
120	76.1	7.8	99	0.5	0.5	0.02	109.6	0.42	0.16
125	76.9	7.9	103	0.5	0.51	0.03	109.6	0.42	0.16
130	78.1	8.1	107	0.51	0.51	0.03	109.7	0.42	0.16

Notes: E = efficiency, Q = quantity, ED = draw of last trade, ρ_B = buyers' correlation coefficient, ρ_S = sellers' correlation coefficient, I = inefficiency measurement, P = price, BPS = buyers' profit split, α = volatility.

For a given number of draws, the efficiency attained by ZI agents is significantly less than that attained by IEL agents. IEL agents attain an efficiency of 72% in 35 draws while it takes ZI traders 105. IEL attains an efficiency of 90% in 70 draws while it takes ZI 250. It takes ZI approximately 3 times the number of draws that it takes IEL to reach the same level of efficiency. In addition, the increases in efficiency and other measurements for IEL level off after 100 draws while those increases do not level off under ZI even after 500 or more draws.[16] The main reason for the difference is that IEL is less patient than ZI as we explained in Section 2.4.2. Both IEL and ZI traders satisfy the Marshallian trading principle - the buyer and seller correlation coefficients are positive - but because of its impatience, IEL does more so. Given the same number of draws, more infra marginal units are being traded by IEL than by ZI. This is quantified in the Inefficiency Ratios of the two simulations.

ZI average prices are about 4% higher than those of IEL but still slightly under the competitive equilibrium prices of $[110, 114]$. ZI average prices start out higher in the first trade and remain higher throughout. See Table 2.6. The buyers' share of profits for IEL traders is 48%. Because of the higher average prices, the buyers' share of profits with ZI traders is only 42%. That is about what the BPS would be at the competitive equilibrium price of 110.

Table 2.6: ZI price dynamics: average price of the first 10 trades

	Order of Trade									
draws	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
25	107.7	108	107.3	105.8	104.8	102.6	92.9			
50	108	109.4	109.3	109.1	108.7	108.5	108.7	105.8		
75	107.6	109.1	109.5	109.6	109.5	109.1	109.3	109.6	108.1	110.9
100	108.6	109.5	109.5	109.6	110	109.7	109.5	109.4	109.6	108.3
125	108.1	109.3	109.5	110.1	109.8	109.6	109.8	109.7	109.8	110
150	107.3	109.4	109.8	109.7	109.6	109.8	110	109.8	109.5	109.9
175	108	109.2	109.7	109.9	109.9	110.1	109.7	109.9	110	109.8
200	108.1	109.5	109.5	109.8	109.8	109.8	109.9	109.9	110.1	110

2.6 Comparing IEL and ZI to MobLab subjects

Cason-Friedman (1996) concluded that "ZI now seems the only natural source for a null hypothesis in assessing the performance of any more complicated mechanism." Lin et al. (2020) went further and said "We find much stronger support for ZI theory compared to Cason-Friedman ... It appears ZI explains dynamics within periods." In this section we explore whether IEL improves on ZI in explaining the first period dynamics of a CDA in the MobLab experiments.

One of the key findings of both the IEL and ZI simulations is that the number of draws matters. This means any comparison of IEL or ZI to humans is not straight-forward. What value should we choose for the number of draws in order to compare to the human data?

[16]In Table B.4, we provide the simulation results for draws of 135 to 750 in various increments.

To see the problem in detail, consider the data in Table 2.3. IEL with 65 draws produces an average efficiency of 85.7 which is close to the human number of 85.9%. But, for 65 draws, the quantity traded by IEL is 9.7 which is less than the 11 of humans. To get IEL to trade 11 units requires 160 draws. To get the sellers' correlation of 0.57 we need 80 draws. What is the right choice? 65, 80, or 160?

Rather than arbitrarily fitting the number of draws to the data *ex post*, we decided to tie our hands and use a number of draws equal to the average number of orders in the human data. For the 2090 observations of the MobLab data, the average number of orders is 55. Using that for our number of draws, we get Table 2.7 comparing IEL and ZI statistics to those of the MobLab human subjects.

Table 2.7: IEL vs. ZI vs MobLab

	MobLab	IEL (55 draws)	ZI (55 draws)
Efficiency	85.9	82.7	50.8
Quantity	11	9.2	4.7
Inefficiency Ratio	0.49	0.09	0
Buyers' Correlation	0.56	0.51	0.37
Sellers' Correlation	0.57	0.53	0.36
Average Price	113.2	105.3	109.4
Buyers' Profit Split	0.36	0.48	0.43
Price Volatility	0.26	0.19	0.17
Orders	55	55	55

Observation 7.
IEL is better than ZI at explaining the MobLab efficiency data.

All of the efficiency statistics of IEL are closer to the MobLab measures than the ZI statistics. This is also apparent from the plots of the cumulative distributions of the sample efficiencies in Figure 2.2.

Observation 8.
ZI is better than IEL at explaining the MobLab price data, although neither explains all of it.

Consider Figure 2.2. It is obvious that the ZI distribution is closer to the MobLab distribution than IEL is. But, neither ZI nor IEL can explain all of the MobLab prices. About 25% of the MobLab average prices are higher than the highest IEL average price of 125 or the highest average ZI price of 144. MobLab average prices are more dispersed than either ZI or IEL average prices at both the upper and lower ends of the distributions. MobLab prices are also more volatile.

Figure 2.2: IEL vs ZI vs MobLab: average price and efficiency CDFs

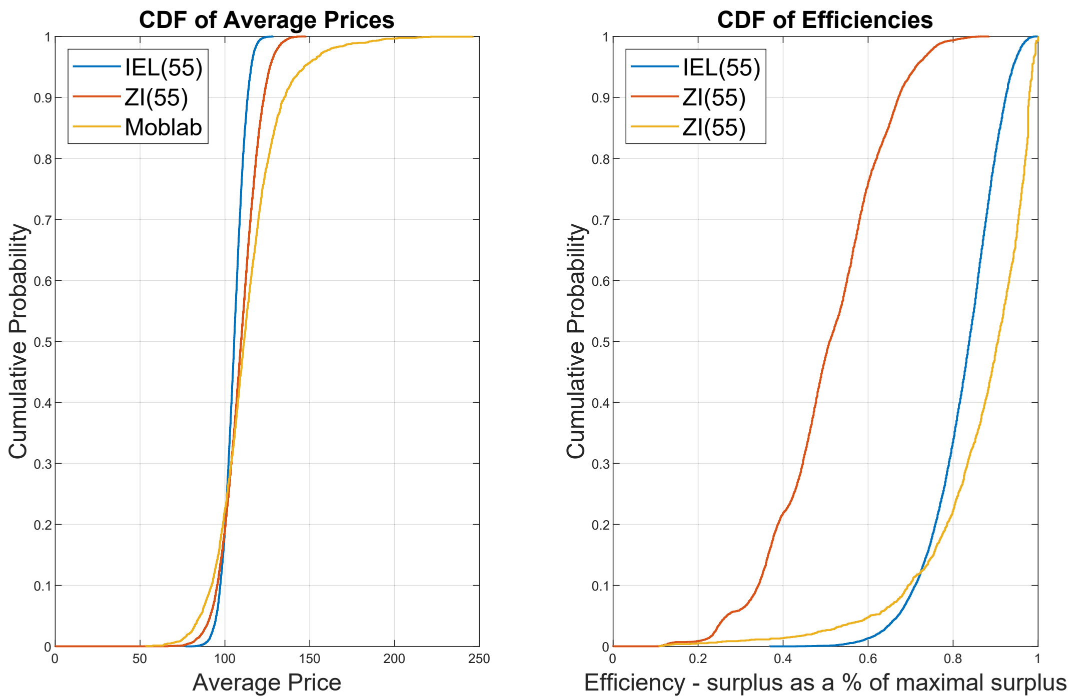

2.6.1 Full trading

It would be easy to conclude at this point that neither IEL nor ZI is a very convincing model of the MobLab subjects. But, one might argue that we have inappropriately handicapped ZI traders by arbitrarily choosing 55 draws. Perhaps, with enough draws, the ZI efficiency statistics would be closer to the MobLab data than IEL and we would have to reject IEL in favor of ZI. To test this out, we let both IEL and ZI have enough draws "to allow sufficient time for trading" as was done in Gode-Sunder (1993, p.122). For IEL that would mean 100 or more draws. We pick 115. For ZI it would mean 600 or more draws. We pick 750. We then get Table 2.8 and Figure 2.3.

Table 2.8: MobLab subjects vs IEL vs ZI: full trading

	E	Q	ρ_B	ρ_S	I	P	BPS	α
MobLab(55 orders)	85.9	11	0.52	0.57	0.49	113.2	0.36	0.26
IEL(115 draws)	91	10.9	0.55	0.59	0.23	106	0.48	0.18
ZI(750 draws)	97.6	12.2	0.66	0.70	0.42	109.9	0.42	0.14

Notes: E = efficiency, Q = quantity, ρ_B = buyers' correlation coefficient, ρ_S = sellers' correlation coefficient, I = inefficiency measurement, P = price, BPS = buyers' profit split, α = volatility.

Observation 9.
With full trading

(a) IEL is better at explaining efficiency.
(b) ZI is better at explaining prices.
(c) But, neither IEL nor ZI are very good models for the MobLab subjects.

The statistics in Table 2.8 tell the same story as before: for efficiency, IEL is closer to MobLab while for price, ZI is closer. But a glance at Figure 2.3 reveals that neither IEL nor ZI with full trading are good models for the MobLab subjects. Efficiencies are too high. Prices are low and not dispersed enough. Full trading is not the answer.

Figure 2.3: IEL vs ZI vs MobLab subjects: full trading

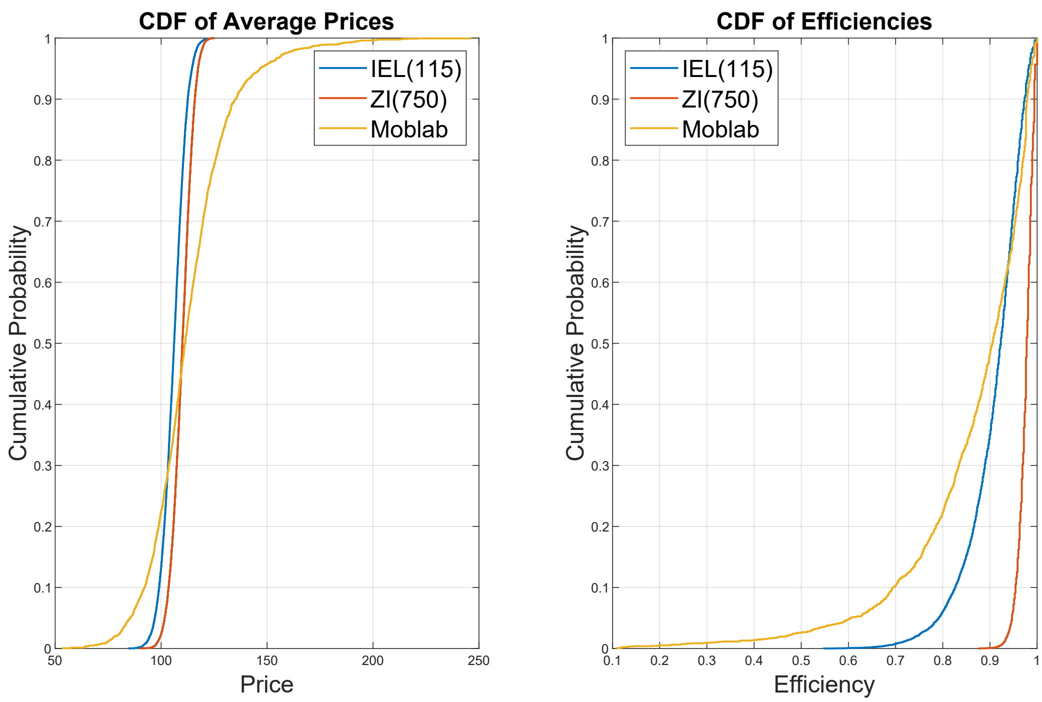

2.7 Improving the theory

If any model is to explain or predict human behavior in the MobLab experiments *by itself*, then it will have to make mistakes. That is, IEL buyers would have to be willing to bid or accept asks above their values and IEL sellers would have to be willing to ask or accept bids below their costs. But, those are dominated strategies and not something economists usually put into their models.[17] We take a different route. We introduce heterogeneity into the agent pool.

[17] As Gjerstad and Shacht (2021) point out in their very nice paper, the standard assumption of "individual rationality" implies buyers will bid no higher than their value and sellers will not ask below their costs. Clearly the MobLab data reject that assumption.

Suppose that the pool of subjects that each experimental session draws from is composed of IEL agents plus agents who either do not take the rules seriously or are confused. Suppose these other agents submit bids and asks randomly, violating the individual rationality constraint, but in a somewhat constrained manner. In particular, they randomly select their bids and asks from the range $[s_l, s_N]$. We call these agents No Intelligence (NI) agents. In the special case, when $s_N = s_u$, these are the ZI-U (ZI unconstrained) agents in Gode-Sunder (1993).

In our simulations, we assume that η percent of the pool of potential agents are NI traders. In each simulation, at the start of a run, we randomly assign each trader to be NI with probability η and IEL with probability $1 - \eta$. Other than that, the simulation remains the same as in Section 2.4.1. We do this 10,000 times and generate efficiencies, prices, etc. We do this for a number of possible values for both s_N and η. η ranges between 0.05 and 0.4 while s_N ranges between 200 and 300. We restrict attention to s_N less than or equal to 300 because in the MobLab experiments bids and asks are restricted to $[0, 300]$.

To determine appropriate values for η and s_N, we did a grid search. We minimized a standard sum of squares distance measure on efficiency and average price The results are in Table 2.9. We found that $\eta = 0.3$ and $s_N = 250$ yielded results closest to the MobLab efficiency and price at a distance of 0.08.

Table 2.9: IEL+NI search

	S$_N$=200			S$_N$=250			S$_N$=300		
%NI	e	p	d	e	p	d	e	p	d
0.05	83.2	105.4	57.36	83.0	106.7	44.37	83.0	107.7	35.00
0.10	83.8	105.6	51.05	83.8	108.0	27.08	83.5	110.0	15.80
0.15	84.4	105.7	46.95	84.4	109.1	16.17	84.1	112.4	4.89
0.20	85.0	105.7	44.99	85.0	110.6	6.37	84.5	115.1	5.47
0.25	85.5	105.7	44.11	85.5	111.8	1.75	85.0	117.0	12.37
0.30	86.1	105.4	47.53	86.0	112.9	0.08	85.6	119.1	27.29
0.35	86.7	105.6	45.94	86.5	114.0	0.99	85.9	121.0	47.48
0.40	87.1	105.3	50.66	87.1	115.1	4.77	86.5	123.9	89.83
iel only	82.7	105.2	63.82						
zi only	50.8	109.4	1680.93						

Notes: e=efficiency, p=price, d= distance = $(e - e^*)^2 + (p - p^*)^2$ where (e^*, p^*) are the MobLab values.

Table 2.10: IEL + NI vs ZI vs MobLab

	E	Q	ρ_B	ρ_S	I	P	BPS	α
MobLab (55 orders)	85.9	11	0.52	0.57	0.49	113.2	0.36	0.26
IEL(55)& NI(0.3,250)	86	10.5	0.55	0.56	0.30	112.9	0.42	0.35
IEL(55)	82.7	9.2	0.51	0.53	0.09	105.3	0.48	0.19
ZI(55)	50.8	4.7	0.37	0.36	0.00	109.4	0.43	0.17

Notes: E = efficiency, Q = quantity, ρ_B = buyers' correlation coefficient, ρ_S = sellers' correlation coefficient, I = inefficiency measurement, P = price, BPS = buyers' profit split, α = volatility.

Observation 10.
The model of 70% IEL agents and 30% NI agents, with $S_N = 250$, outperforms ZI in explaining the MobLab data.

Figure 2.4: IEL + NI vs ZI vs MobLab: average price and efficiency

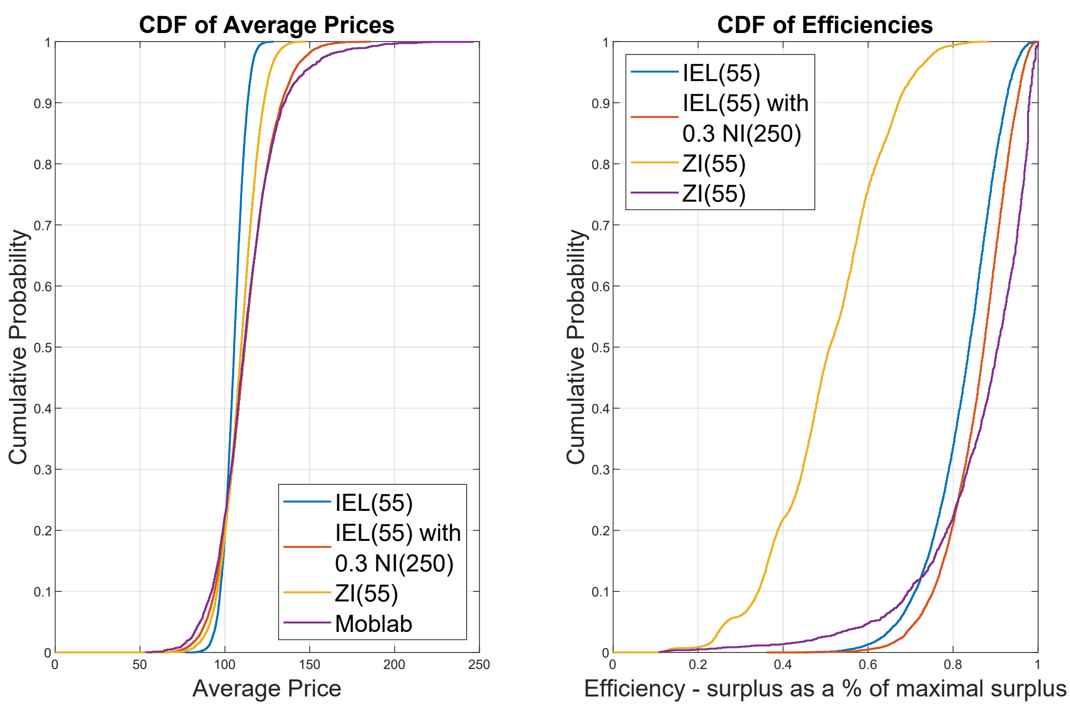

Looking at the statistics in Table 2.10 and the Figure 2.4, it is easy to see that mixing in ZI agents with the IEL agents gets the efficiency statistics a little closer to the MobLab statistics. But the effect on the price statistics is dramatic. In Figure 2.4 the CDF of average prices generated by the mixture of IEL and NI agents is almost identical to the CDF of the MobLab average prices. The mixture of IEL and NI is clearly a better explanation of the MobLab CDA price and efficiency data than is ZI.

2.7.1 Robustness

We end this section with a brief comment on the robustness of these results to variations in the IEL parameters J, the size of the remembered set, and μ, the probability of experimentation. We ran the same simulations with IEL as before for various values of J and μ. The results are in Table 2.11.

Varying μ has very little effect. This is because for IEL in the CDA, it only matters how many of the remembered set are above the best ask and how many are below. Increasing J increases efficiency and decreases price up to $J = 300$ and then they both level off. This is because for IEL in the CDA, there are only 300 possible bids in $[0, 300]$ and so adding extra adds very little.

Table 2.11: Robustness: efficiency and price

	$\mu = 0.0033$	$\mu = 0.033$	$\mu = 0.25$
$J = 50$	(77.2, 105.8)	(77.4, 105.7)	(77.2, 105.8)
$J = 100$	(82.5, 105.5)	(82.7, 105.2)	(82.6, 105.3)
$J = 200$	(84.9, 104.8)	(85.0, 105.0)	(85.2, 105.0)
$J = 300$	(85.5, 104.8)	(85.6, 104.8)	(85.4, 104.7)
$J = 400$	(85.9, 104.8)	(85.9, 104.7)	(85.9, 104.7)
$J = 500$	(86.2, 104.6)	(86.1, 104.6)	(86.2, 104.5)

But the key finding of our robustness check is that it does not change our conclusion that a mixture of 70% IEL and 30%NI with a 250 upper bound is the best fit to the MobLab data. The results of a grid search, when $J = 300$, across the percent of NI and the upper bound are in Table B.8. The closest result, in the sum of squares distance between the simulation result and the MobLab efficiency and price, would have 15%NI with an upper bound of 300. 30%NI with an upper bound of 250 came in a very close second. Now consider Figure 2.5 where the CDFs for efficiency and price are displayed. It is very difficult to distinguish the 70-30 mix with an upper bound of 250 (our model when $J = 100$) from the 85-15 mix with an upper bound of 300 when $J = 300$.

2.8 Better data

In the previous section, we showed that a pool with a mix of 70% IEL agents and 30% NI agents leads to statistics that closely approximate those of the MobLab data. But the question still remains, does IEL best model those agents who are not confused or fooling around. One way to answer that is to look at some data generated under stricter controls than MobLab in the hope that any human NI-like agents will be eliminated.

We report on three new CDA market sessions we conducted online through the Simon Fraser University (SFU) on 4/20/2021 and 4/21/2021. In each period, the environment was

Figure 2.5: IEL+NI robustness: average price and efficiency

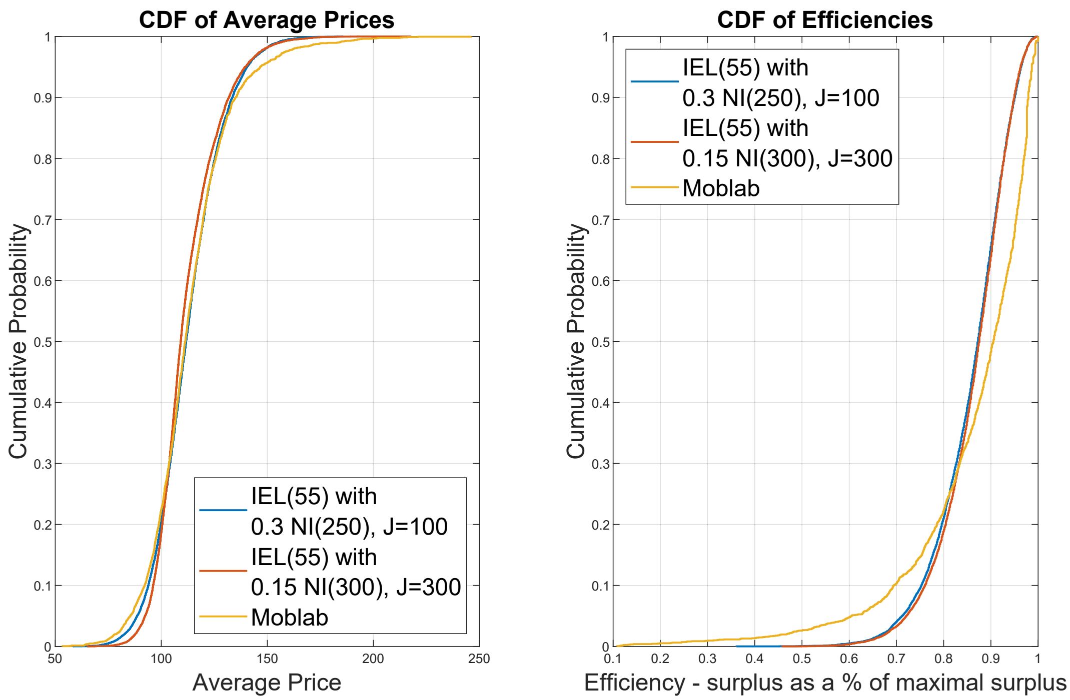

the same as the environment in Section 2.2.1. There were 5 buyers, each of whom could buy up to 3 units, and 5 sellers, each of whom could sell up to 3 units each. In each session, participants traded for 14 periods. There were 2 practice periods, then 10 regular market periods, and then 2 periods where IEL and humans trade together. For this paper, we only use the data from the 10 regular market periods.

Experiments were conducted via an online software called Flex-E-Markets, which provides software as a service platform to organize and manage multiple, simultaneous online marketplaces.[18] Because of Covid restrictions, these experiments were run online. Screen shots are provided in Section B.4.1. Subjects earnings consisted of a show-up fee ($7) and the profits made during the non-practice trading rounds. Subjects were paid for all non-practice sessions. Total earnings ranged from $7 to $25 with an average of $18. Each session lasted 70 minutes.

In one of the sessions, a subject, acting as a seller in the second 5 regular market periods, consistently sold more than 3 units in 3 in spite of instructions to the contrary. Flex-E-Markets does not have a way to prevent this. We set aside the data for those second 5 periods. Thus, there are a total of 25 observations from the SFU experiments.

For these experiments, we wanted each group to participate in 10 market periods during a session. We also wanted each period appear to involve a new demand-supply configuration. One way to do that would be to randomly draw new values and costs each period as in Cason-

[18]See http://flexemarkets.com.

Friedman (1993). But we also wanted to have as much data as possible from the configuration in Section 2.2.1. We adopted a compromise. Each period we randomly reallocated buyers along the demand curve and sellers along the supply curve (that is, we shuffled values and costs) and we also drew a random amount (positive or negative) that we then added to each value and cost. This had, we believe, the effect of providing an environment each period that subjects thought was new, from their point of view, but which were all equivalent, under quasi-linearity. Subjects were either a buyer for 5 rounds and then a seller for 5 rounds, or vice versa. Table B.10 contains the values and costs resulting from our randomization process.

At the beginning of each session, all traders were informed that all values and costs would be drawn independently from [0, 200].[19] Before the start of each trading period, each trader saw their own values or costs for that period and did not see the values and costs of others.

In Table 2.12, we display the performance measurements for the SFU sessions along with those from the Moblab sessions for comparison.

Table 2.12: SFU vs Moblab vs IEL vs ZI

	SFU	Moblab	IEL(40)	ZI(40)
Efficiency	71	85.9	74.1	40.5
Quantity	7.6	11	7.8	3.6
Inefficiency Ratio	0.06	0.49	0.05	0.00
Buyers' Correlation	0.25	0.56	0.49	0.31
Sellers' Correlation	0.32	0.57	0.47	0.30
Average Price	105.7	113.2	104.8	109.1
Buyers' Profit Split	0.47	0.36	0.48	0.42
Price Volatility	0.17	0.26	0.19	0.18
Orders	40.4	55	40	40

The incentives in the SFU sessions were clearly stronger than those in the Moblab sessions. In the SFU data 13.8% of the trades have one loser and this does not seem to go away in later periods. In the Moblab data that is 90%. Only one period in the SFU data, which was the first period of one session, had at least 5 trades with a loser. This is 4% of the periods. For Moblab that is 25%. The highest trade price over all periods in the SFU data is 133 compared to an average price greater than 150 in 4.3% of the periods in the Moblab data.

[19]Each period, all values and costs are initially in [50, 150]. For each session, we added a constant randomly drawn from U $\sim$ [-50, 50]. This implies that values and costs are coming from [0, 200]) since the lowest cost is 50 and the highest value is 150 in the original S/D schedule

Comparing SFU to Moblab, we see lower efficiency due to lower trading and lower prices with a more equal split between buyers and sellers. As we will see, this is can be explained by the fact that the SFU subjects do not seem to be NI-like.

How do IEL and ZI do in explaining these data? We used the same rationale for the number of draws for IEL and ZI that we did earlier. Since the average number of orders in the SFU experiments was 40,[20] we used 40 draws in all our simulations. We provide a comparison of the statistics in Table 2.12 and Figure 2.6.

Based on the statistics in Table 2.12, IEL is much closer to behaving as the SFU subjects than ZI is. With the exception of the correlation coefficients and Price Volatility, the IEL statistics are closer to the SFU data than are the ZI statistics. Figure 2.6 clearly supports that conclusion for the efficiency.

Figure 2.6: IEL vs SFU vs ZI: efficiency

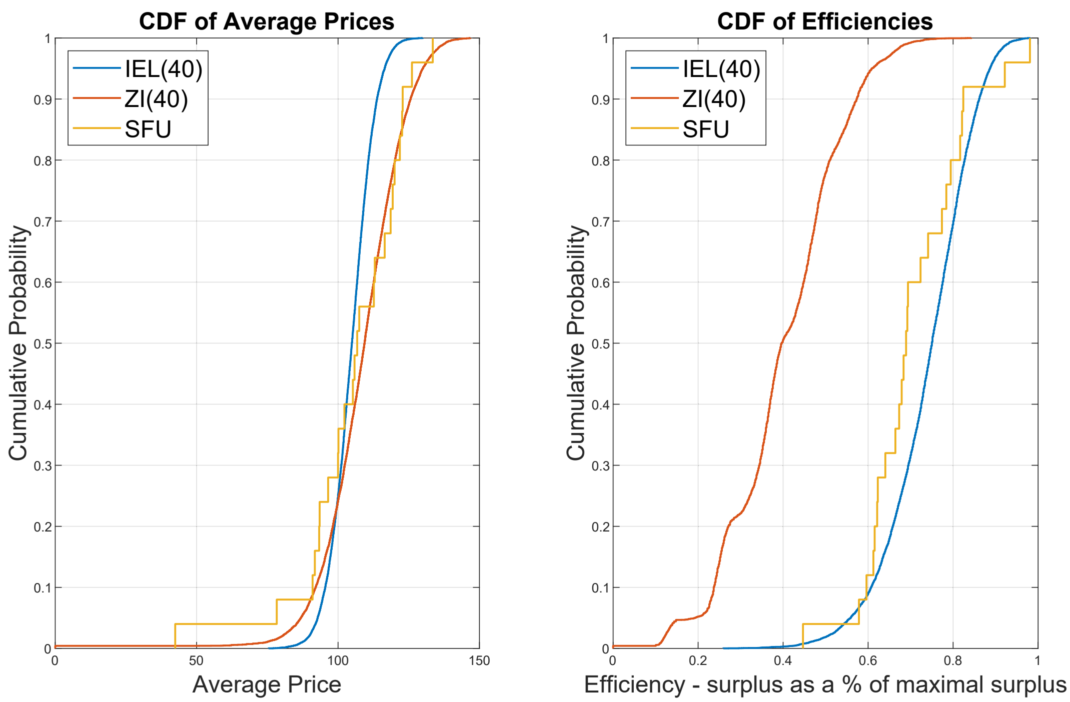

But the left panel of the same figure reflects a different story for prices. Except at the very top, the distribution of the average prices for ZI agents seems to be a better fit to the distribution of average prices for the SFU subjects than is the IEL distribution. IEL prices are less dispersed than either ZI or SFU prices.

We explored whether a mixture of NI and IEL might improve on the IEL fit for SFU as it did for MobLab. Surprisingly, at least to us, the answer was no. The results across the grid, varying the % of NI and the upper bound on NI orders, are in Table 2.13. A robustness

[20]51% of the orders were bids from buyers.

check, redoing the simulations with $J = 300$, also indicated that IEL alone was the best alternative. Those results are in Table B.9.

Table 2.13: IEL+NI search

	S_N=200			S_N=250			S_N=300		
%NI	e	p	d	e	p	d	e	p	d
0.01	74.0	104.8	18.58	74.0	105.0	18.29	74.0	105.4	17.93
0.02	74.2	105.3	20.46	74.1	105.3	19.21	74.3	105.8	21.61
0.03	74.4	105.1	23.25	74.1	105.5	19.10	74.2	106.5	21.14
0.04	74.4	104.8	23.66	74.6	105.9	25.74	74.2	106.8	21.40
0.05	74.6	105.1	26.03	74.5	106.1	24.44	74.6	107.3	28.00
0.10	75.2	105.1	35.32	75.1	107.7	36.93	74.8	109.9	44.43
0.15	75.9	105.1	47.95	75.8	109.1	56.05	75.4	112.4	78.58
0.20	76.5	105.0	60.45	76.6	110.3	81.15	76.0	114.8	123.71
iel only	73.9	104.9	17.26						
zi only	40.5	109.1	1855.71						

Notes: e=efficiency, p=price, d= distance = $(e - e^*)^2 + (p - p^*)^2$ where (e^*, p^*) are the MobLab values.

Observation 11.
(a) IEL clearly is better than ZI as an explanation for the efficiencies from the SFU experiment.
(b) ZI appears to be a little better than IEL as an explanation for the average prices from the SFU experiment.

2.9 Summary and discussion

2.9.1 Summary

The goal of our research is to develop a model of individual behavior similar to that of human subjects in continuous double auction experiments. In this paper we report on two models, ZI and IEL, which we tested against each other using two very different data sets: a large, uncontrolled set from classroom experiments using the MobLab interface and a small, controlled set from experiments at SFU.

We found that drawing subjects from a pool composed of 70% IEL agent and 30% NI agents, who randomly order from $[0, 250]$, generated results that were a very good fit to the 2090 observations from MobLab. That mixture outperformed a pool composed only of ZI agents. Given the large number of violations of individual rationality in the MobLab data, it is not surprising that some collection of irrational agents, like the NI agents, is needed to explain that data.

We found that a pool composed of only IEL agents was a good fit to the 25 observations from SFU. With respect to the distribution of efficiency, that IEL model dominated the ZI model. But, with respect to the distribution of average prices, the ZI model does a little better than IEL. Weighing the efficiency and price fits equally, IEL seems to be a better overall fit to the SFU data than ZI. We further found that mixing NI agents into the IEL pool does not improve the fit. The tighter controls and more salient payoffs of the SFU experiments seem to have eliminated any NI type behavior among the subjects.

2.9.2 Robustness

There are several free parameters in our simulations. For IEL, there are the size of the remembered set, J, and the rate of experimentation, ρ. We used $J = 100$ and $\mu = 0.033$, the same values we have used in many other papers which cover many different economic situations. To be sure that this apparent arbitrariness was not crucial, we tried variations of the two parameters and found our results were robust to those variations.

For the simulations, the free parameters are the range of allowable orders and the number of draws to use in the Gode-Sunder bid timing process. For the allowable orders, we used $[0, 300]$, the exact set used in both the MobLab and SFU experiments. We found that the number of draws significantly affects the simulations of both IEL and ZI agents. Rather than choosing the number of draws to generate the best results, we used the average number of orders for each data set; 55 for MobLab and 40 for SFU. The results confirm that this was not an inappropriate choice.

2.9.3 Discussion

IEL seems to be doing very well as a model of the behavior of humans in the first period of a CDA. But there are a couple of improvements that should be made in our future work.

Although the choice to make the number of draws in the simulations equal to the average number of orders in the experiments has some justification, it is to some extent arbitrary. To be a complete theory of behavior in a CDA, IEL needs to be revised to endogenize the agent's decision as to when to bid. One possibility would be to use the model of Gjerstad-Dickhaut (2003), but it is complicated. Another would be the approach of van de Leur-Anufriev (2018) that also involves multi-period learning.

IEL outperformed ZI in explaining the efficiencies and prices generated by humans with one exception – the distribution of average prices in the SFU experiments. The SFU distribution has a larger range than that generated byI which, in turn, is larger than that generated by IEL. Our future work will try to understand the reasons for these differences and to modify IEL to account for it.

It is important to note that IEL is myopic and learns. IEL is not strategic. As long as an IEL agent is interacting with others of the same type, this doesn't seem to be a problem. The data generated are similar to experimental data. We conjecture, however, that a more strategic agent might create problems for IEL agents. For example, if humans and IEL agents interact in an experiment, will they behave similarly? If not then the IEL agent will need further modification to achieve our goal of a model of individual behavior similar to that of human subjects in continuous double auction experiments.

Chapter 3

Racial Differences in Senior Executives' Access to Information

3.1 Introduction

Race inequality is well-documented in the literature. African-Americans do not enjoy the same access to opportunities as their non-African-American counterparts (see, e.g., Pope and Sydnor, 2011; Ravina, 2019; Bartlett, Morse, Stanton and Wallace 2021; Begley and Purnanandam 2021). Prior studies document that African-Americans are less likely to have economic opportunities because of statistical discrimination (Arrow, 1972; Phelps, 1972) and taste-based discrimination (Becker, 1957).[1] However, less is known about race inequality in corporate leadership. In this work, we investigate race differences in the insider-trading behavior of corporate leaders of S&P 1500 companies to determine if African-American corporate executives have equal access to the networks that generate valuable insider information.

Prior studies show that informal work-related and social interactions are important in building human and organizational capital within the workplace (O'Leary and Ickovics, 1992; Korenman and Turner, 1996; Bolte, Immorlica, and Jackson, 2020). For instance, managers usually require the cooperation and support of many people within their social and professional networks but outside their formal reporting hierarchy. Homophily—the notion that individuals are more likely to interact and form relations with others similar to them—can lead to personal and professional groups being segregated and homogenous with respect to individual characteristics, including race.[2] This can lead to African-American

[1]Statistical discrimination refers to an approach that relies on a discriminative factor as a statistical tool to predict behavior when economic agents have imperfect information about individuals they interact with, whereas taste-based discrimination refers to discrimination because of prejudice. See Lahey (2008) for a review of this literature.

[2]Homophily in race and ethnicity are the strongest divides in social networks. After race and ethnicity, the greatest social divides are created by age, religion, education, occupation, and gender in approximately that order (McPherson, Smith-Lovin, and Cook, 2001).

executives having fewer network ties than their Caucasian counterparts (Ibarra, 1993).[3] If social and professional networks are indeed racially homophilous, then African-American executives may have limited access to information to make decisions because pervasive homophily can lead to material information that flows through social networks in a workplace to be highly localized (Golub and Jackson, 2012).[4]

In this paper, we use insider trading data from 2010 to 2018 to analyze race differences in insider trading profitability and formally test the hypothesis that African-American executives have limited access to networks that generate valuable inside information due to race-based discrimination. Prior research has shown that insiders have access to material nonpublic information and that they use this informational advantage when they engage in open market sales and purchases of their firm's stock.[5] In particular, insiders tend to buy stock before good news and sell stock before bad news. If African-American executives are excluded because of their race from some social and professional networks where nonpublic information is discussed, then the availability and access to such private information should be reflected in the profitability of their insider trades.

Following prior studies, we focus on insider purchases because insiders often sell shares of the companies where they work for of liquidity reasons (Inci, Narayanan and Seyhun, 2017). Purchases, therefore, contain greater information content. To obtain information on the race of executives, we manually downloaded photos of executives from the internet. We then evaluated each photo to identify the race of an insider.[6] In our database, African-Americans make up 1.06% of executives and 1.10% of the purchases in terms of the number of transactions.[7]

Consistent with prior studies, we show that insider trades are highly profitable (Jaffe, 1974; Seyhun, 1986; Lakonishok and Lee, 2001). The 90-day equal-weighted average of market-adjusted returns after insider purchases is 5.5%, which is economically significant. More importantly, we document that there is a significant race difference in insider profitability: non-African-American executive purchases are 4.6% (3-month cumulative market-

[3]Lalanne and Seabright (2011), Mengel (2015), and Lindenlaub and Prummer (2016) provide evidence for gender differences in networking outcomes including at top executive jobs.

[4]Another possible explanation for segregation within a workplace would be "home bias" effect, which refers to the phenomenon wherein agents (businesses, funds, etc.) are more likely to conduct transactions with parties who are geographically closer to them, either in the same country or same state, rather than those outside. For more details of home bias effect, see French and Poterba (1991), Coval and Moskowitz (1999), Ahearne, Griever and Warnock (2004), and Tesar and Werner (1995).

[5]See Seyhun (1998) for a review. Corporate insiders are officers, directors, and large shareholders of public corporations.

[6]The pictures were cross-checked by two separate individuals to ensure accuracy. The data collection and verification procedures are explained in more detail in Section 3.

[7]This number is consistent with the observation that less than one percent of all companies throughout the US have a Black leader at the helm (source: https://afrotech.com/richest-black-ceos)

adjusted return) more profitable than the African-American executive purchases.[8] Finally, we document that abnormal returns of African-American executives do not attain statistical significance at any time during the year following their transactions, indicating that African-American trades do not contain private information.

We also examine the insider-trading profitability of Asian-American executives to see if the experience of African-American executives is unique with respect to other minority groups. The results show that the abnormal profits of Asian-American executives are similar to or even exceed those of Caucasian executives. These results are consistent with a number of papers that have shown that Asian-Americans have greater access to economic opportunities and experience higher economic outcomes compared with African-Americans (see, e.g., Chetty, Friedman, Saez, Turner and Yagan, 2020), given the unique experience and history of African-Americans in the United States (Acharya, Blackwell and Sen, 2016).

To rule out other alternative explanations of the results, we use empirical specifications to control for confounding factors. Specifically, the results could be driven by omitted industry characteristics, such as insider trading opportunities within an industry, that may be simultaneously correlated with the number of African-American executives working within a specific industry. We control for industry-fixed effects to take into account that unobserved industry-specific factors may explain the findings. In addition, the differences between non-African-American executives and African-American executives can be attributed to the fact that, on average, non-African-American executives hold higher-ranking positions (e.g., non-African-American executives are more likely to be CEOs and CFOs), which give them greater access to private material information. To rule out this alternative explanation, we use position-fixed effects. We further account in our empirical model for other personal characteristics such as insider holdings, compensation, and education level of executives. Our evidence suggests the differences between African-American executives and non-African-American executives cannot be explained by personal characteristics such as position, age, educational level or compensation.

The differences between African-American and non-African-American executives we find may also be attributed to disposition factors such as greater risk-aversion on the part of African-American executives. By conducting three tests, we rule out the possibility that race differences in trading profitability are driven by dispositional factors. First, we investigate the timing and intensity of trading patterns of African-American executives. Specifically, if African-Americans are less optimistic, lack confidence, or are more risk-averse than their counterparts, they would be less likely to conduct insider trades. On the contrary, if we document that no differences exist in trading intensity, then the better access to informa-

[8]Using size-momentum-value adjustments, non-African-American executives earn 3.6% whereas African-American executives earn 0.6%, a difference of 3.0%. Once again, none of the abnormal returns for African-American executives attain statistical significance at any horizon.

tion of non-African-American executives likely plays a role in their greater trading profits. The evidence suggests the latter—almost no differences in trading frequency exist between African-American executives and non-African-American executives.

Second, we examine favorable (profitable) insider trading 30 and 45 days before an earnings announcement. Profitable trades before an earnings announcement are more likely to draw attention from public and regulatory authorities and are more likely to be investigated by the Securities and Exchange Commission (Huddart, Ke and Shi, 2007). If African-American executives are more risk-averse, we would expect a lower intensity of favorable insider trading before earnings announcements. However, between African-American and non-African-American executives, no significant differences appear in favorable trading intensity before an earnings announcement.

Third, we examine differences in the backdating of options granted to African-American and non-African-American executives. Backdating refers to the fraudulent practice of choosing a past date as the option-award date to maximize the value of the option award. In particular, previous studies have shown that corporate insiders tend to receive options awards at or near minimum stock prices just before an increase in the company's share prices to maximize their compensation at the expense of shareholders (Narayanan, Schipani, and Seyhun, 2007; Heron and Lie (2007); Narayanan and Seyhun, 2008; Bizjak, Lemmon and Whitby, 2009; Anginer, Narayanan, Schipani and Seyhun, 2011). Backdating options is considered an unethical practice and is subject to legal and regulatory enforcement. If African-American executives are more risk-averse, we would expect them to engage less in backdating options than non-African-American executives. Empirically, the results show no difference in option award patterns between the two groups, again suggesting that differences in the information content of insider trades are unlikely to be driven by differences in risk aversion.

If race-exclusive informal networks are the channels through which non-African-American executives gain an informational advantage, then we should expect the advantages to such executives to be smaller in firms that discourage discriminative behavior and promote equity and diversity. We thus investigate whether the informational advantage of non-African-American executives is attenuated in workplaces that promote equity and diversity. Following prior literature (see, e.g., Flammer and Luo, 2017), we use employee-related corporate social responsibility (CSR) scores from the Kinder, Lydenberg, and Domini (KLD) database as a proxy for equity at the workplace. We test whether race differences in trading and profits are less profound in firms that emphasize equity and diversity. The results suggest that race differences in information access are mitigated by a strong corporate emphasis on equity and diversity. However, the advantages of equity and diversity do not extend to Asian-American executives who already trade on superior information.

In addition, to control for hidden power differences in titles, and as a further and out-of-sample check, we obtained race data for corporate directors from a different data vendor, Institutional Shareholder Services (ISS), and examined whether race differences exist in information access for African-American directors. We document similar results for corporate directors as those of our main analysis of corporate executives.

These findings make several important contributions to the literature. To the best of our knowledge, this is the first examination of racial differences in access to material information by top-level insiders. We document strong evidence supporting the existence of these racial differences, showing that African-American executives are at a disadvantage relative to their Caucasian and Asian-American counterparts. Second, the findings also complement those reported in the insider-trading literature by revealing new channels, namely homophilous networks, through which executives gain an informational advantage and increase their insider-trading profits. This paper also offers insights into how to tackle taste-based discrimination in the workplace. In particular, the results suggest that corporate investment in CSR, especially the equity and employee aspects, may mitigate discrimination and promote equity and diversity.

The remainder of the paper is organized as follows. Section 3.2 presents the literature review and develops the hypotheses. Section 3.3 describes the data and empirical methodology. Section 3.4 explains the results, and Section 3.5 concludes.

3.2 Related literature and hypotheses development

3.2.1 Literature review

The importance of race characteristics is acknowledged by many studies in economics and finance, which show that an individual's race profile often has a strong impact on a wide range of social and financial outcomes (Hellerstein and Neumark, 2008; Pope and Sydnor, 2011; Chetty, Friedman, Saez, Turner and Yagan, 2020). The race differences are partly driven by statistical discrimination, which is based on statistical inference of economic behavior (Lang and Spitzer, 2020). For example, Ayres and Siegelman (1995) show that the disparate treatment of African-Americans by car dealerships may be explained by dealers' statistical inferences about consumer reservation prices. In addition, a fraction of the race differences is due to taste-based discrimination, which reflects prejudices or preferences (Lang and Spitzer, 2020). For instance, studies show that African-American borrowers face higher borrowing costs even after controlling for the borrower's financial attributes, personal characteristics, and the ex-post probability of default (Pope and Sydnor, 2011; Ravina, 2019).

Our paper advances this literature by showing that, even among top-level executives and directors, African-American executives face information disadvantages in the labor market, as reflected in their lack of insider-trading profits. This result is consistent with the findings of Bertrand and Mullainathan (2004), who document that discrimination makes it harder for African-Americans not only to find a job but also to improve their employability. In addition, the results show that race differences are attenuated in firms that place greater emphasis on equity and diversity. Collectively, by ruling out potential alternative explanations, our evidence suggests that the documented race differences in the access to information of senior executives are likely due to taste-based discrimination.

This paper also advances the literature that examines information networks in insider trading. Cohen, Frazzini and Malloy (2008) show that portfolio managers take on larger positions in companies they are connected to through their network, and earn significantly higher returns on these positions. Jagolinzer, Larcker, Ormazabal and Taylor (2020) study political connections and opportunistic trading by corporate insiders and show that political connections facilitate information flow and lead to greater profitability by connected insiders. Cao et. al (2015) examine the trades of independent directors socially connected to their firms' senior executives. They show that connected directors earn significantly higher returns than unconnected independent directors. Ahern (2017) examines social relationships in illegal insider trading networks. He shows evidence of private information flow through friends and family of corporate insiders. In a closely related paper, Inci, Narayanan and Seyhun (2017) examine gender differences in executives' access to information using insider trading. This paper contributes to this strand of the literature by showing the importance of homophilous networks through which some executives gain an informational advantage and increase their insider-trading profits.

This paper also contributes to a growing literature on the determinant role of physical appearance in economics and finance. Although the literature provides ample evidence that individuals with desirable physical attributes receive more favorable judgment in various social and economic settings, it is not until recently that physical appearance has been recognized to profoundly impact the job performance of sophisticated financial market participants such as financial analysts (Cao, Guan, Li and Yang, 2020; Peng, Teoh, Wang and Yan, 2020) and CFOs (Hsieh, Kim, Wang and Wang, 2020; Hrazdil, Li, Lobo and Zhang, 2021). Our evidence complements this growing literature by examining how race affects the profitability of insider trading.

Finally, this paper also relates to the literature on the use of CSR policies to incentivize employees (see, e.g., Bode, Singh, and Rogan, 2015; Flammer, 2015; Flammer and Luo, 2017; Cassar and Meier, 2018). These studies suggest that CSR policies constitute an effective non-monetary form of compensation to attract, engage, and retain employees. The present findings support this view by further suggesting that companies that invest in CSR programs improve employee social engagement and mitigate adverse behavior at the workplace. Taste-based discrimination in the workplace is often costly to the employer (e.g., Pope and Sydnor, 2011). We show herein that emphasis on equity and diversity may mitigate such discrimination.

3.2.2 Hypothesis development

To make decisions, executives typically acquire essential information from prescribed and emergent networks (Ibarra, 1995). The former is formally structured and consists of a set of specified relationships between superiors and subordinates and among representatives of different committees and departments. Individuals in prescribed networks are obligated to interact to accomplish a defined organizational task such as generating sales. In contrast, emergent networks involve more discretionary patterns of interactions based on social networks and often offer potential access to information and resources. Emergent networks are important to the functionality of corporations and to executives because of the complex nature of business operations (Kanter, 1983).

Prior studies report that, compared with their Caucasian counterparts, minority managers have fewer intimate social network relationships within the organization (Ibarra, 1995; Korenman and Turner, 1996). This finding is not surprising because considerable evidence in the social psychology literature suggests people tend to interact with those from the same race (Blau, 1977; Feld, 1981). This phenomenon is called the homophily effect. Given that senior executive positions are held disproportionately by Caucasian executives,[9] interactions between the dominant group of Caucasian executives could be much more frequent and rewarding than similar interactions within the minority group. This situation could have undesirable consequences for members of racial minorities who do not have the advantages of being within a larger or more strongly connected social network.

In addition to intra-firm differences, racial differences in access to useful insider information arise from social networks outside the firm. Social networks enhance trust between individuals and facilitate a more timely, transparent, and efficient flow of information between two individuals at two different firms. For example, executives may obtain information from their peers at their customer and supplier firms (Chen, Levy, Martin and Shalev, 2021). Cai and Szeidl (2018) also show a positive effect of business networks on firm performance. Similar to the within-firm difficulty in accessing information, African-American executives may also have less access compared with their counterparts to information from executives outside their firm because of race differences in social networks. Alldredge and Cicero (2015) provide evidence of insiders trading their own company's stock profitably based on public information about their principal customers.

The insider trading literature documents the ability of senior executives to exploit their privileged inside information obtained through formal or informal means to earn greater returns on their trades (see, e.g., Jaffe, 1974; Seyhun, 1986, 1992, 1998; Lakonishok and Lee, 2001). We expect that access to unique information by corporate insiders is reflected in their insider trading patterns. We conjecture that African-American executives have less

[9] African-Americans account for more than 10% of the US population, but only around 1% of executives in our database.

access to information and, therefore, we expect African-American executives to earn lower returns from insider trades. Our first hypothesis is thus stated as follows:

H1: Insider trades by African-American executives are less informative and have lower abnormal returns compared with insider trades by non-African-American executives.

It is also possible that African-American executives may have to work harder and may require greater talent to overcome workplace discrimination and other social hurdles to become an executive in the first place (Bertrand and Mullainathan, 2004). It then becomes possible that they have advantages in access to information. This positive selection bias can potentially overcome the network effect and result in African-American executives being more informed than their counterparts.[10] Therefore, whether African-American executives earn higher insider trading profits is an empirical question.

This work also explores whether differences based on race in access to material information are mitigated if firms take initiatives to improve diversity and reduce adverse discriminative behavior. Prior studies show that CSR activities are a form of social capital that introduces social structure into the firm and encourages desirable social norms and behaviors (see, e.g., Mirvis, 2012). As such, CSR programs are often used as a form of compensation that is hard to imitate by other peer firms and can align employee interests with employer interests (Flammer and Luo, 2017). Shen and Benson (2016) document how socially responsible human resource management (SRHRM), defined as CSR directed at employees, affects employee work behavior. Employees of firms that adopt organization-level SRHRM have better task performance and are more likely to help coworkers with job-related problems. Thus, we conjecture that a firm's CSR investments in employees, equity, and diversity help align the incentives of employees and the firm and facilitate inter-racial communication and cooperation. Based on the above discussion, our second hypothesis is as follows:

H2: The differences in the information content and abnormal returns of insider trades by African-American executives and insider trades by non-African-American executives are lower in firms that invest in socially responsible human resource management programs.

There may be a relationship between the CSR strategy pursued by a firm and the hiring of African-American executives. In particular, firms that place greater emphasis on CSR may be more likely to hire African-American executives to meet an imposed quota of minorities in the upper management of the company. Executives hired under the pressure of

[10]To further differentiate between race and talent issues, we also control for age, education level, and professional achievements of insiders.

such quotas could be "token representatives" and may be placed in less important roles. In such situations, African-American executives may not be effective contributors to important decisions and therefore have less access to key information. We would thus expect them to be less informed and earn lower profits in firms that invest in CSR equity and diversity programs. Ultimately, whether African-American executives have greater access to information in high-CSR firms is an empirical question.

3.3 Data

We use ExecuComp to identify executives of public companies for the period from January 2010 to December 2018. To capture the ethnicity of executives, we obtained their names from the ExecuComp database and searched for a clear frontal photograph of each executive using Google Images. We followed several steps to ensure a high quality of ethnicity information from our image search: First, we used the executives' names and affiliations to search Google Images for photographs for each executive and downloaded the photographs from the search results. We then evaluated each photo to identify the race of each insider. In particular, five undergraduate research assistants evaluated the photographs to determine the racial profile of each executive.

To ensure high-quality data, one of the authors then examined all the images again to check the accuracy of the data. After manually categorizing insiders based on their photographs, 15,598 distinct insiders were identified, 149 of whom were African-American.[11] As a further data-validation procedure, we obtained race information from the ISS and cross-checked it with the race information obtained by using the first procedure. ISS covers the information of corporate directors, while ExecuComp covers executives. There are 3203 overlaps between the two databases, and they matched 99.8% of the time. We further reconciled the differences by again manually searching the information. To examine whether the African-American executive experience is unique and to compare their insider-trading behavior with that of another minority group, we repeated the same process to identify Asian-American executives.[12] We identified 546 (3.5%) insiders as being Asian-American.

Following prior literature, we obtained insider-trading data from the Thomson Reuters Insider Filing Data Feed. The data set contains all open-market insider trading of shares by officers, directors, and beneficial owners of publicly traded firms. Shares acquired through the exercise of options, stock awards, and trades with corporations were excluded. We obtained stock price information from the Center for Research in Security Prices (CRSP).

[11] After removing observations with missing variables, the share of African-American executives in the final sample is 1.1%.

[12] We follow the U.S. Census Bureau for definitions of Asians as "people having origins in any of the original peoples of the Far East, Southeast Asia, or the Indian subcontinent (for example, Bangladesh, Cambodia, China, India, Indonesia, Japan, Korea, Malaysia, Pakistan, the Philippine Islands, Thailand, and Vietnam)."

The primary objective of this study is to examine the profitability of insider trades by ethnicity. To accomplish this, we first aggregated insider trades by executives each day by calculating the net number of shares traded. If the net number of shares traded for insider i at day t is positive (negative), then this insider is considered a net buyer (net seller) on that day.[13] We removed any insider-day pair where the insider is considered neutral or a net seller.[14] That is, we removed all insider-day pairs with zero or a net negative number of shares traded. The reason for separating insider trades by purchases and sales is that purchases are most likely motivated by information, whereas sales are likely to be driven by the diversification and liquidity needs of the insider (Seyhun, 1998).[15] For the backdating analyses, we obtained the option awards of insiders from the Thomson Reuters Insider Filing Data Feed (obtained from forms 4 and 5 submitted to the SEC) with a transaction code of "A."

Because Thomson Reuters Insider Filing and ExecuComp have no common identifiers for corporate executives, we matched the names of insiders by using a fuzzy matching algorithm: 46% of executives with a photo are successfully matched with insider trading data.[16] Among successful matches, 64 of insiders are categorized as African-American. We then used the BoardEx data set to obtain executive characteristics such as gender, age, individual network size (number of overlaps through employment, other activities, and education), an indicator for an insider's past achievements (honors, educational titles, etc.) and education level. In addition, we obtained the financial information of firms from the COMPUSTAT industrial files.

In the analyses, we control for several firm characteristics. In particular, we control for book-to-market ratio (*bm*), calculated as the book value of equity divided by the market value of equity; firm size (*log(mktcap)*), the natural log of firm market value; and leverage (*leverage*), the total liabilities divided by shareholders' equity. All financial variables are lagged by at least six months. To eliminate the influence of extreme values, leverage, book-to-market, and market value variables are winsorized by 0.5% at the top and the bottom. We also control for some insider characteristics. Those include the age (*age*) and the gender (*male*) of an insider, the number of years that an insider has spent in his current firm (*tenure*), an individual's network size (*log(networksize)*), compensation

[13]We followed the standard way of insider trade aggregation following Seyhun (1988) and Lakonishok and Lee (2001).

[14]In the untabulated results, we also compare the race differences of insider sales and we generally do not find any statistical differences.

[15]For the trading before earning announcement analyses, we include insider sells before the announcement of bad news as they are likely to be opportunistic.

[16]Our matching rate depends on whether an insider with a photograph made an insider trade during the study period; if an insider did not make any insider trade, there is no match.

(*log(compensation)*) and equity holding (*shares owned*), an indicator for an executive's past achievements (*achievement*), and two dummies accounting for an executive's education level (*phd* and *cpa*). To eliminate the influence of extreme values, all continuous variables were winsorized at their 0.5 and 99.5 percentile values.

After removing the observations with missing key variables, our sample consists of 3,201 observations. Table 3.1 reports the summary statistics of our key variables. *African-American* is an indicator variable taking the value of unity if an insider is African-American. In the final sample, African-American insiders account for 1.1% of all insider purchases. In the sample, the average age is 56.4 years and they had worked, on average, 6.55 years in the same firm. Additionally, 94.4% of all insider purchases were made by male executives.

From the KLD database, we obtained data on environmental, social, and governance (ESG) criteria for company operations. This database is organized by year and contains a summary count of all strengths (indicated by +1) and concerns (indicated by 1) the company received in a general category in the given year. In our analysis, we focused on three ESG ratings: *humanity* (e.g. respect for the fundamental convention on human rights), *employee* (e.g. employee involvement, job satisfaction, health and safety of the workplace), and *diversity* (e.g. promotion of women and minorities, maintenance of diversity and equal opportunities). Our ESG ratings were calculated as the sum of all strengths and concerns associated with a given category. That is, higher values imply a better performance for a given category. For our sample, the average humanity and employee scores are positive (0.04 and 0.32) whilst the average diversity score is negative (-0.13).[17]

3.4 Empirical results

3.4.1 Main results

We begin our analyses by examining univariate changes in insider trading profitability across ethnicity groups. The profitability of insider trades is measured by the cumulative market adjusted and size-momentum-value adjusted cumulative returns following an insider trade. For each insider i in day t, we compute 3, 6, 9, and 12 months' cumulative market adjusted and size-momentum-value adjusted returns. The size-momentum-adjusted returns were calculated as per Daniel, Grinblatt, Titman and Wermers (1997). Tables 3.2 and 3.3 report the summary statistics of insider trading profitability and nonparametric t-test results by trade type and insider ethnicity, respectively. Collectively, these results indicate that non-African-American executives make positive, statistically significant abnormal profits from their purchases, whereas African-American executives make negative yet statistically

[17] Also, firms in our sample with at least one African American executive have higher employee (0.49 vs 0.31) and diversity (0.01 vs -0.15) scores whilst the differences of humanity score are negligible (0.03 vs 0.04).

insignificant abnormal profits. The profits of African-American executives do not attain positive significance at any holding period.

If informational asymmetry exists across different ethnic groups with African-American executives being less informed, then one would expect their insider trades to be less profitable. Tables 3.4 and 3.5 formally reports the mean and median differences across insider ethnicities. In Table 3.4, we report t-test results for the mean difference. The results show that insider purchases by African-Americans are significantly less profitable than those of the rest of the executives. Table 3.5 reports the Wilcoxon rank-sum test results for the median difference. The result shows that the median profitability of insider purchases by African-Americans is significantly less than that of insider purchases of all other ethnicities. Figures 3.1 and 3.2 show a graphical representation of this difference by plotting the equal-weighted average of market-adjusted returns around a 250-trading-day window of purchases made by African-American and by non-African-American executives. On average, non-African-American executives purchase when the stock price is at its lowest point within a 100-day window. This V-shaped pattern in abnormal returns is indicative of informed trading, and the same pattern does not appear for African-American insiders. Abnormal returns are negative after their purchases. In contrast, the profitability of purchases by Asian-American executives is similar to that of Caucasian executives and directors.

These differences in the univariate analysis may be driven by firm characteristics such as size and leverage or by insider characteristics such as age, gender, and tenure. To control for these differences and to test more formally the profitability of insider trades across ethnicity differences, we perform the following regression:

$$ret_{i,(t+1,t+\tau)} = \alpha + \beta_1 \textit{African-American}_{i,t} + \beta_2 \textit{controls} + \epsilon_{i,t} \tag{3.1}$$

Above, *ret* is the T-day cumulative market adjusted or size-momentum-value adjusted return following an insider i's trade at day t.[18] In our empirical specifications, we include 3, 6, 9, and 12 months' cumulative returns as our dependent variable.[19] i denotes each insider and t denotes each trading day; *African-American* is an indicator variable taking the value of 1 if an insider i's ethnicity is reported as African American; controls include logarithm of the market value (*log(mktcap)*), book-to-market ratio (*bm*), and leverage (*leverage*) of a firm that insider i is currently employed at, as well as the age (*age*), gender (*male*), and tenure (*tenure*) of an insider i. We also control for the logarithm of an individual's network size (*log(networksize)*), compensation (*log(compensation)*) and equity holding (*shares owned*), an indicator of an executive's past achievements (*achievement*), and two dummies accounting

[18]T value can change depending on how far we are looking into the future and when the trade is made. For three-month cumulative returns, for example, T can take different values depending on the insider trade day t. If an insider makes a trade on April 1^{st}, T will be 91; if the trade is made on May 1^{st}, T will be 92.

[19]It should be noted that cumulative returns are bounded below as returns cannot be lower than -100%.

for an executive's education level (*phd* and *cpa*). To take into account that unobserved role-specific[20] and time-specific factors may explain our findings, we also control for year and insider role (CEO, CFO etc.) fixed effects.[21]

The results of this regression are reported in Tables 3.6 and 3.7 where we report the results for market adjusted returns and size-momentum-value adjusted returns, respectively. The four columns report the abnormal return result for insider purchases for 3, 6, 9 and 12 months. . In all specifications, except for 3 months, we find that insider trades by African-American executives are significantly less profitable than those of other executives. These results are consistent with the first hypothesis that insiders of African-American ethnicity are less informed than other executives.

Next, we examine whether the race differences are due to access to information or to dispositional factors, such as risk aversion, overconfidence, or lack of confidence (Inci, Narayanan and Seyhun, 2017). We carried out three tests that show that race differences in trading profitability are not driven by dispositional factors. First, we examine trading intensity. If the documented race differences are driven primarily by risk aversion, non-African-American executives would trade more frequently. To examine differences in insider trading intensity between African-American and non-African-American executives, we run the regression specified in equation 3.1 but replace the dependent variable with insider-purchasing intensity. Following Inci, Narayanan and Seyhun (2017), we use the logarithm of the number of shares traded by insiders. Table 3.8 reports the trading intensity regression results. The results indicate that no significant differences exist in insider trading intensity between African-American and non-African-American insiders.

Second, we examine differences in favorable trading intensity before earnings announcements. We focus on earnings announcements because they contain important information for the public market. Profitable insider trades are more likely to be investigated by the Securities and Exchange Commission if made before an earnings announcement (Huddart, Ke and Shi, 2007). We run a logit and a probit regression where the dependent variable is an indicator variable that takes on a value of unity if an insider made a purchase (sale) in the 30 days preceding a positive-surprise (negative-surprise) earnings announcement. As an alternative, we use as a dependent variable a dummy that takes on a value of unity if an insider made a favorable trade in the 45 days preceding an earnings announcement. We use the same set of controls in the regression reported in Table 3.8. Table 3.9 reports the

[20]For example, our results could be driven by the fact that African-American executives may be less likely to be CEOs of the company who have the most insider information.

[21]We cannot conduct firm-fixed effects regressions since some of our variables are at the firm level. Furthermore, we also cannot simply compare the profitability and intensity of trading between African-American and non-African-Americans for the same firms since we would not be able to control for the position of the insiders. In a given firm, if the CFO is African-American, there would be no non-African-American CFOs to compare with.

regression results. Columns 1 and 2 report probit regression results, and columns 3 and 4 report logit regression results. In all columns, the coefficients of African-Americans do not differ statistically from zero. Overall, the results are similar to those reported in Table 3.8: no systematic differences exist between African-American and non-African-American executives. These findings are inconsistent with risk-aversion or lack-of-confidence explanations for the differences in profitability of insider trades by African-American executives.

Finally, we examine the timing of option awards to executives to see if African-American executives are less likely to engage in the fraudulent practice of backdating. Previous literature reports that insider option awards exhibit telltale signs of fraudulent backdating behavior (Heron and Lie, 2007, Narayanan, Schipani, and Seyhun, 2007, Narayanan and Seyhun, 2008). The evidence suggests that, when receiving an option award, top-level executives and the board of directors typically engage in collusive, fraudulent behavior and falsely designate a date in the past when the stock was at or near a minimum level as the option grant date to maximize the value of their option awards. This practice is tantamount to secretly (and illegally) awarding themselves in-the-money options at the expense of their shareholders and taxpayers. This line of reasoning again suggests that insiders option awards tend to occur at or near minimum or V-shaped prices.

The evidence of the potential backdating of option awards appears in Table 3.10. Following Anginer, Narayanan, Schipani and Seyhun (2011), we compute for each firm-insider pair in our sample the number of times options were granted at one of the three lowest stock prices over a 51-day window centered on reported grant dates. We then divide this number by the firm-insider pair's total number of option grant dates during the sample period. If the resulting ratio exceeds 10%, we classify the insider as having engaged in backdating. We run a logit and a probit regression where the dependent variable is an indicator variable that takes the value of unity if the insider is classified as having engaged in backdating. We designate African-American executives with an indicator variable and use the same set of controls as we did for the regression reported in Table 3.8. The probit (logit) regression results are reported in column 1 (2). The results in Table 3.10 show that no differences exist in the likelihood of option backdating between African-American and non-African-American executives.

Overall, the results suggest that dispositional factors do not explain why non-African-American executives earn greater profits from insider information than African-American executives.

3.4.2 Impact of Corporate Social Responsibility

In this section, we examine, while controlling for various firm characteristics, whether insider trades of African-Americans at firms with higher CSR scores are more profitable than those of African-Americans working at firms with lower CSR scores. Cross-sectional differences based on CSR scores could reveal potential remedies to discrimination problems within firms.

In particular, we examine how the insider race versus profitability relationship varies with scores for (1) the social record in human-rights protection, (2) strong employee relations, and (3) an emphasis on diversity. We begin by constructing from the KLD database three SRHRM measures: *humanity*, *employee*, and *diversity*.[22] The humanity score measures a company's effectiveness in respecting fundamental human rights such as labor rights at work. The employee and diversity scores measure a company's effectiveness concerning job satisfaction, health and safety in the workplace, and diversity and equal opportunities. Employee scores include items such as union relations, retirement benefits, and health and safety measures. Diversity scores include items such as policies for work-life balance and rules safeguarding the rights of female, LGBTQ, and minority employees from potential discrimination. The KLD database reports the number of strengths and concerns captured from global media sources for each category. Following prior literature (see, e.g., Marano and Kostova, 2016), we create our SRHRM measures by calculating the difference between the reported number of strengths and concerns,[23] where higher values indicate better performance for a given category. We hypothesize that higher scores in these categories would facilitate greater employee engagement, reduce workplace discrimination, promote equity and inclusion, and reduce social and professional networking barriers for African-American workers to improve information flow within a company. In turn, these would make insider trades by African-American executives more informative than insider trades by African-American executives working in companies with lower SRHRM scores.

To test our hypothesis, we construct the following specification:

$$\begin{aligned} ret_{i,(t+1,t+\tau)} =& \alpha + \beta_1 \textit{African-American}_{i,t} + \beta_2(\textit{African-American}_{i,t} \times \textit{SRHRM}_t) \\ & + \beta_3 \textit{controls} + \epsilon_{i,t} \end{aligned} \tag{3.2}$$

where, *SRHRM* is either humanity, employee, or diversity score of a company at time t.

The results of the regressions with humanity, employee, and diversity measures are reported in Tables 3.11 & 3.12, 3.13 & 3.14, and 3.15 & 3.16, respectively. Tables 3.11, 3.13, and 3.15 include the result for market-adjusted returns; the rest include the result for size-momentum-value-adjusted returns. For insider purchases, consistent with pre-

[22]KLD is an independent social choice investment advisory firm that compiles ratings of companies' effectiveness in addressing the needs of their stakeholders. The KLD database contains yearly social ratings of companies along several dimensions such as employee relations, diversity, the environment, and human rights.

[23]The KLD dataset presents a binary summary of positive (strength) and negative (concern) ESG ratings. In each case, positive and negative ratings are indicated by 1 and 1, respectively. If the company did not have a strength or concern in regard to an issue, it is indicated by zero. Our ESG ratings are calculated as the sum of all strengths and concerns associated with a given category.

vious results obtained in this work, insiders from other ethnic groups still outperform African-American executives, which is suggested by the negative coefficients of the African-American variables. More importantly. we document positive significant coefficients for the ($\textit{African-American}_{i,t} \times \textit{SRHRM}_t$) interaction variable across most specifications, suggesting that insider trades by African-American executives working at firms with higher SRHRM scores are more profitable than those of their African-American peers working at other firms. The economic magnitude of the coefficients is also significant. For instance, we find that, for an African-American executive, a one-standard-deviation increase in the employee score of a firm is associated with a 6.1% higher 6-month cumulative market-adjusted return following their insider purchases. These findings confirm our second hypothesis that African-American executives employed in firms that emphasize and invest in fairness and diversity have better access to information than their African-American counterparts in other companies, which is reflected in more profitable purchases.

An alternative explanation of our ESG results is that the title of the executive does not capture the power of the office. One stereotype is that African-Americans are disproportionately promoted into HR-type roles and that these roles are low-power roles. Lacking power, they lack information, which means that people are less inclined to network with them. In firms that value ESG, the HR department gains power, giving the HR executive more power and more access to networking. The alternative explanation suggests that the greater parity at ESG firms just reflects the greater power of the HR executives in ESG firms. However, the evidence from racial differences in director trading (shown in Tables 3.20 and 3.21) is inconsistent with this alternative explanation.[24]

Furthermore, some assert that ESG initiatives represent a zero-sum game at the expense of Caucasian and Asian-American executives. We also tested this alternative explanation. Although not shown, we replicated the analysis in Tables 3.11-3.16 for Asian-American executives. The results show that the profitability for Asian-American executives does not decline for ESG firms, which is inconsistent with the zero-sum explanation.

The results of the cross-section analysis have important implications. First, these results suggest that our initial findings are unlikely to be driven by statistical discrimination. For instance, African-American executives may have less access to information because of poor networking skills. Others may be less willing to communicate with them because the marginal benefit of networking with them is low. The cross-sectional results suggest that this is an unlikely explanation. On the contrary, the evidence suggests that taste-based discrimination likely drives race differences and could be mitigated by SRHRM programs that encourage equity and diversity. The evidence presented also contributes to the current

[24]It might be the case that minority groups seek for employment opportunities at firms that promote equity and diversity. This might result in employees from minority groups to get more populous in such firms, which, in turn, creates a better within company information flow and improves work performance of minority groups. Unfortunately, we are not able to test this potential selection bias due to lack of data.

policy debate regarding the regulation of mandatory equity and diversity programs for listed firms. For example, the Security Exchange Commission recently approved Nasdaq's proposal—the first of its kind—to require minimum race and gender representations and disclosures in its listing rules.[25] The present findings support such initiatives by indicating that taste-based racial discrimination exists in top corporate executive settings and that it may be worthwhile to establish stronger anti-racial discrimination initiatives.

3.4.3 Asian-American Insiders

Next, we examine the profitability of insider trading by Asian-American executives to see if the experience of African-American executives is unique. In particular, we examine whether Asian-American insiders face similar disadvantages in information access. Since the 1990s, sociologists have documented a new racial structure in America. As a result of Asians' efforts to integrate into the "mainstream" society, the economic and cultural differences between Caucasians and Asians have faded in the post-Civil Rights era (Lee and Bean, 2007). Conversely, African-Americans have yet to overcome these barriers and remain at the bottom of the social hierarchy (Gans, 2005).

To test whether differences exist for insider trading by Asian-American executives, we repeat the analyses reported in Tables 3.6 and 3.7 but use an indicator variable for Asian-American executives instead of an African-American indicator. The results are reported in Tables 3.17 and 3.18. After controlling for other determinants of insider trading profitability, Asian-American insiders have similar or better returns, suggesting that their access to information is similar to that of their non-Asian counterparts. These results are consistent with those reported in the social sciences literature that documents favorable socioeconomic outcomes for Asian-Americans despite widespread discrimination (especially in labor markets) encountered by Asian-Americans in the early part of the 20th century (Gans 2005).

Overall, our results suggest that the experience of African-American executives is unique, perhaps reflecting the unique historical experience of African-Americans with vestiges of slavery and state-enforced racial segregation, directly and indirectly, contributing to economic outcomes to this day (Acharya et al. 2016). Another potential explanation is that a significant portion of Asian-American executives tend to be foreign-born. As a consequence, these Asian-American executives may have already learned to overcome social, cultural and linguistic barriers in networking within organizations. This potential selection bias could also explain the differences we observe.[26]

[25] Nasdaq companies are required to meet gender and diversity requirements or explain in writing why they have have failed to do so. See https://www.cnbc.com/2021/08/06/sec-approves-nasdaqs-plan-to-boost-diversity-on-corporate-boards.html.

[26] We do not have enough observations of foreign-born executives with African ethnicity to test this potential explanation.

3.4.4 Robustness

We conduct two robustness checks of our main results. First, we show that our results are robust against nonlinear effects due to control variables by using a nearest-neighbor matching approach. In addition, we perform an out-of-sample check by using race data for corporate directors collected from a different data vendor: Institutional Shareholder Services (ISS).

The results may be attributed to the fact that African-American executives are not directly comparable to other executives. For example, African-American executives generally work at larger companies that yield, on average, lower absolute stock returns (and, consequently, lower trading profitability) for their insider trades. To account for this potential problem, we follow Dehejia and Wahba (2002) and conduct nearest neighbor matching where each observation from one ethnic group is matched with the "closest" or the "most similar" observation from another ethnic group. The similarity between subjects is based on a weighted function of the covariates for each observation. For this analysis, we include market value (*log(mktcap)*), book-to-market ratio (*bm*), *leverage*, gender (*male*), *age, tenure, humanity, employee, diversity*, network size (*log(networksize)*), compensation (*log(compensation)*), equity holding (*shares owned*) and insider role and year dummies as covariates. In particular, for each observation from one ethnicity group, we pick only one, i.e., the most similar, observation from the other group when performing matching.

The results of the nearest neighbor matching are reported in Table 3.19. The reported values are the returns of non-African American executives minus the returns of African-American executives following their insider trades. We find that the profitability of African American insider purchases are less profitable than that of the others, as implied by the positive and significant values for three- and six-month cumulative returns following insider trades. Nine- and twelve-month cumulative returns for purchases are also positive (but not statistically significant). These results suggest that the race differences in insider trading profits are robust against the nonlinear impact of observable covariates.

Next, we examine insider trades of corporate directors by ethnicity. We do not include the insider trades of directors in our main sample because we only have photographs of top executives. We use the ISS Compensation Data Set to obtain the ethnicity information of corporate directors.[27] The ISS dataset provides the opportunity to examine race differences in insider trading by using an external source for ethnicity information as an "out-of-sample" robustness check. The ISS dataset includes insider-specific variables that we used in our previous analysis with executives such as age, gender, tenure, network size, achievement, and education dummies. Because of the lack of common insider identifiers, we match the names of directors by using a fuzzy matching algorithm when merging the ISS data with

[27] To the best of our knowledge, ISS is one of the only databases that provide race information. Other common databases, such as BoardEx and ExecuComp, do not have this information.

Thomson Reuters Insider Filing data. To avoid the possibility that our results may be due to the executives who are also directors, we further restrict our sample to directors by eliminating observations that overlap with our ExecuComp database.[28]

After including the stock return and financial statement data, removing the observations with missing key variables, and winsorizing the market value, book-to-market ratio, and leverage variables by 0.5% at the top and bottom, our final sample contains 4,302 daily insider purchases, 138 of which belong to African-American executives. The sample period is from January 2010 to December 2018.

Tables 3.20 & 3.21 report the results where we summarize the result for market-adjusted and size-momentum-value-adjusted returns, respectively. Consistent with our previous findings, results indicate that the insider trades of corporate directors with African-American backgrounds are significantly less profitable than those of corporate directors of other ethnicities. In addition, the magnitude of the coefficients is like those reported in Tables 3.6 & 3.7 for executives.

The results for directors indicate that subtle differences in the executive power of officers cannot account for the prior results we report for executives. Directors tend to be more homogeneous across firms than executives. Consequently, this evidence from directors provides additional support for the discrimination hypothesis.

3.5 Conclusion

This paper investigates racial differences among corporate executives in S&P 1500 firms. The results show that the profitability of insider trading by African-American executives is significantly less than that of their Caucasian and Asian-American counterparts. These results cannot be attributed to the firm type or to the industries for which the executives work. We also rule out unequal risk-aversion, overconfidence, and ethics as the basis for these differences. These results also cannot be explained by insider attributes such as gender, education, achievement, or compensation. In addition, these race differences are smaller in firms that place greater emphasis on equity and diversity.

Overall, these results are consistent with the hypothesis that African-American executives are excluded from some social and professional networks where nonpublic information is discussed, including the top echelons of corporate hierarchy. Given the importance of social networks and professional networks in labor markets, these results suggest that a lack of equal access to information can place qualified African-American executives at a disadvantage when competing with non-African-American counterparts for career promotions and higher compensation. The evidence presented herein also supports new initiatives that have been introduced recently to promote racial equity and diversity in public corporations.

[28] The results are identical if we remove this restriction.

Figure 3.1: Returns from insider purchases - executives

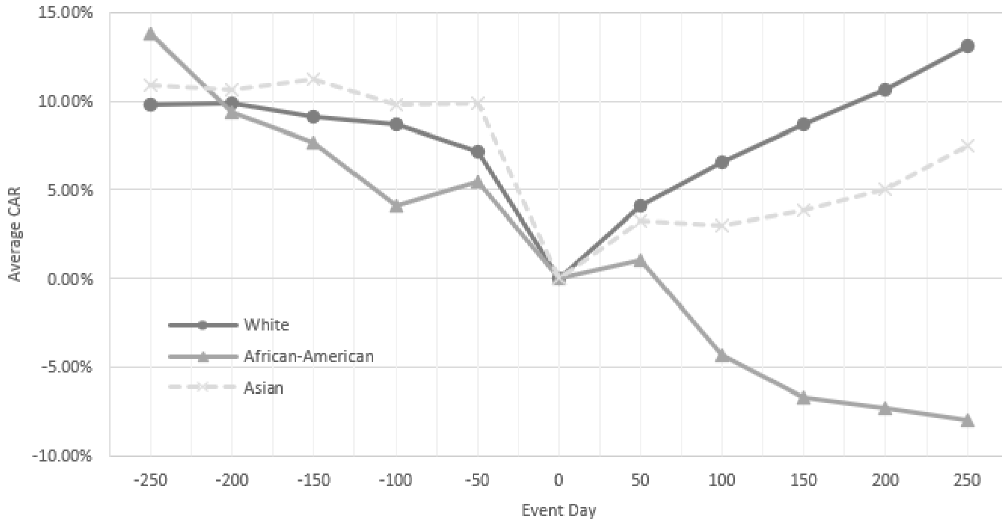

Notes: This figure displays the cross-sectional equal-weighted averages of the market-adjusted cumulative abnormal returns (CAR) of African-American, Asian, and other (White) insider purchases in event time, where the event is defined as the date of an insider purchase.

Figure 3.2: Returns from insider purchases - directors

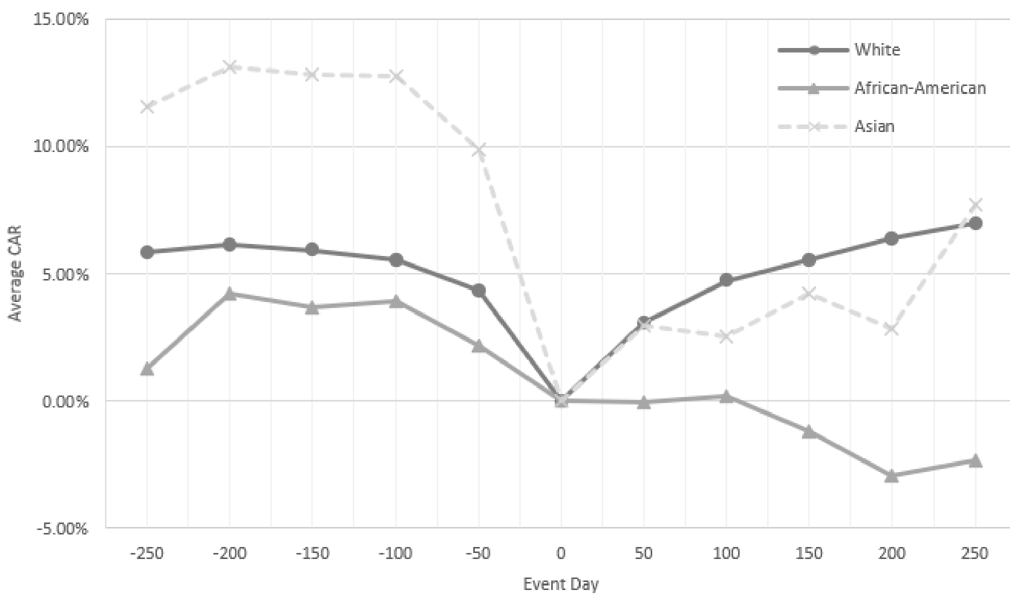

Notes: This figure displays the cross-sectional equal-weighted averages of the market-adjusted cumulative abnormal returns (CAR) of African-American, Asian, and other (White) insider purchases in event time, where the event is defined as the date of an insider purchase.

Table 3.1: Summary statistics

Variables	# of obs.	Mean	Std	25^{th} pctile	50^{th} pctile	75^{th} pctile
African-American	3201	0.011	0.104	0.000	0.000	0.000
log(mktcap)	3201	7.233	1.638	6.239	7.084	8.174
bm	3201	0.787	0.555	0.396	0.685	1.017
leverage	3201	3.595	6.673	0.733	1.653	5.522
age	3201	56.375	8.057	51.000	56.000	60.000
male	3201	0.944	0.230	1.000	1.000	1.000
tenure	3201	6.571	5.324	2.000	5.000	9.000
log(networksize)	3201	6.567	1.245	5.778	6.739	7.530
achievement	3201	0.361	0.480	0.000	0.000	1.000
log(compensation)	3201	7.559	1.022	6.830	7.531	8.241
shares owned	3201	0.017	0.057	0.001	0.002	0.007
phd	3201	0.031	0.174	0.000	0.000	0.000
cpa	3201	0.157	0.370	0.000	0.000	0.000
humanity	2490	0.040	0.269	0.000	0.000	0.000
employee	2490	0.320	0.994	0.000	0.000	0.000
diversity	2490	-0.131	1.105	-1.000	0.000	0.000

Notes: This table reports firm and top executive characteristics for the sample used in the analyses, including the number of insider purchase observations, average value, standard deviation, 25th percentile, median, and 75th percentile for key variables. African-American is an indicator variable taking the value of 1 if an insider has African American ethnicity and 0 otherwise. bm is the book value of equity divided by market value of equity. log(mktcap) is the natural logarithm of the market value of the firm. leverage is total liabilities divided by shareholders' equity. age is the age of an insider; male is an indicator variable taking the value of 1 if an insider is male and 0 otherwise; tenure is the number of years that an insider has spent at his current work. humanity, employee, and diversity are variables measuring a company's performance on human rights, employee satisfaction, and equality-related topics, respectively. log(networksize) is the logarithm of an individual's network size, achievement is an indicator for an executive's past achievements, log(compensation) is the natural logarithm of executive compensation, shares owned is the ratio of company's shares owned by an insider, and phd and cpa are dummy variables taking the value of 1 if an executive has a PhD degree and a CPA certificate, respectively.

Table 3.2: Summary statistics of market and size-value-momentum adjusted returns

	Non-African-American			African-American		
Return	Mean	Std	50^{th} pctile	Mean	Std	50^{th} pctile
market_adj_3m	5.5%	18.3%	2.9%	0.8%	16.5%	-0.2%
market_adj_6m	7.9%	26.7%	4.1%	-4.5%	21.1%	-1.4%
market_adj_9m	10.1%	34.3%	5.0%	-8.3%	28.2%	-5.0%
market_adj_12m	13.4%	43.3%	5.6%	-7.7%	38.5%	-6.0%
dgtw_3m	3.6%	14.8%	1.2%	0.6%	12.4%	1.1%
dgtw_6m	5.0%	24.2%	1.6%	-6.3%	17.9%	-5.8%
dgtw_9m	5.2%	29.9%	0.6%	-10.3%	24.2%	-7.0%
dgtw_12m	7.1%	38.6%	0.8%	-11.8%	32.0%	-8.9%

Notes: This table reports the summary statistics for three, six, nine, and 12 months' cumulative market and size-value-adjusted returns following an insider trade by decomposing the sample into two groups: purchases of non-African-American insiders and purchases of African-American insiders. Statistics include the average value, standard deviation, and median values. market_adj_Xm denotes the X-month market adjusted return. dgtw_Xm denotes the X-month size-momentum-value adjusted returns following Wermers et al. (1997).

Table 3.3: Nonparametric T-test for H_0: Median abnormal return is zero

	Non-African-American			African-American		
Return	H_0 vs. m(ret)>0	H_0 vs. m(ret)<0	H_0 vs. m(ret)!=0	H_0 vs. m(ret)>0	H_0 vs. m(ret)<0	H_0 vs. m(ret)!=0
market_adj_3m	0.000	1.000	0.000	0.832	0.261	0.522
market_adj_6m	0.000	1.000	0.000	0.900	0.168	0.337
market_adj_9m	0.000	1.000	0.000	0.973	0.054	0.108
market_adj_12m	0.000	1.000	0.000	0.973	0.054	0.108
dgtw_3m	0.000	1.000	0.000	0.500	0.625	1.000
dgtw_6m	0.000	1.000	0.000	0.946	0.100	0.200
dgtw_9m	0.040	0.960	0.079	0.946	0.100	0.200
dgtw_12m	0.025	0.977	0.050	0.946	0.100	0.200

Notes: This table reports the the p-values of the non-parametric analog of the single-sample t-test for top African-American executives and the top executives from other ethnicities. The null hypothesis H_0 is that the median of the differences is zero; no further assumptions are made about the distributions. This, in turn, is equivalent to the hypothesis that the true proportion of positive (negative) signs is one-half. m(ret) denotes the median return. market_adj_Xm denotes the X-month market adjusted return. dgtw_Xm denotes the X-month size-momentum-value adjusted returns following Wermers et al. (1997).

Table 3.4: Changes in insider trade profitability across ethnicity groups - mean difference

market_adj_3m	difference	4.6%
	t-stat	1.742
market_adj_6m	difference	12.5%
	t-stat	3.667
market_adj_9m	difference	18.4%
	t-stat	4.043
market_adj_12m	difference	21.0%
	t-stat	3.391
dgtw_3m	difference	3.0%
	t-stat	1.472
dgtw_6m	difference	11.2%
	t-stat	3.893
dgtw_9m	difference	15.5%
	t-stat	3.965
dgtw_12m	difference	19.0%
	t-stat	3.675

Notes: This table reports differences between the average value of cumulative returns following an insider purchase for top executives with African American ethnicity and the top executives from other ethnicities. The coefficients reported next to each row show the difference between the average cumulative return following an insider trade performed by non-African American executives and the average cumulative return following an insider trade performed by African American executives. The t-stats are reported below the differences. market_adj_Xm denotes the X-month market adjusted return. dgtw_Xm denotes the X-month size-momentum-value adjusted returns.

Table 3.5: Changes in insider trade profitability across ethnicity groups - median difference

market_adj_3m	z-stat	2.094
	p-value	0.036
	Prob(Non-AA return > AA return)	0.597
market_adj_6m	z-stat	2.873
	p-value	0.004
	Prob(Non-AA return > AA return)	0.634
market_adj_9m	z-stat	3.528
	p-value	0.001
	Prob(Non-AA return > AA return)	0.664
market_adj_12m	z-stat	3.128
	p-value	0.002
	Prob(Non-AA return > AA return)	0.645
dgtw_3m	z-stat	1.072
	p-value	0.284
	Prob(Non-AA return > AA return)	0.550
dgtw_6m	z-stat	2.945
	p-value	0.003
	Prob(Non-AA return > AA return)	0.637
dgtw_9m	z-stat	3.299
	p-value	0.001
	Prob(Non-AA return > AA return)	0.653
dgtw_12m	z-stat	2.942
	p-value	0.003
	Prob(Non-AA return > AA return)	0.637

Notes: This table reports differences between the median value of cumulative returns following an insider purchase for top executives with African American ethnicity and top executives of other ethnicities. The z-stat and p-value reported next to each row show the results of the hypothesis that cumulative returns following the purchases of African American insiders and the purchases of insiders from other ethnicities are from populations with the same distribution by using the Wilcoxon rank-sum test. Prob(Non-AA return>AA return) is an estimate of the probability that a random cumulative return draw from the non-African-American insider purchases is larger than a random cumulative return draw from the African-American insider purchases. market_adj_Xm denotes the X-month market adjusted return. dgtw_Xm denotes the X-month size-momentum-value adjusted returns.

Table 3.6: Baseline results - market adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
African-American	-0.049	-0.123*	-0.203**	-0.260**
	(0.035)	(0.067)	(0.093)	(0.132)
log(mktcap)	-0.015***	-0.026***	-0.036***	-0.061***
	(0.003)	(0.007)	(0.009)	(0.018)
bm	0.022**	0.038***	0.043**	0.035
	(0.010)	(0.013)	(0.019)	(0.027)
leverage	0.001	0.003***	0.004***	0.004**
	(0.001)	(0.001)	(0.001)	(0.002)
age	-0.001	-0.002	-0.005***	-0.007***
	(0.001)	(0.001)	(0.002)	(0.003)
male	0.026**	0.014	0.033	0.038
	(0.011)	(0.020)	(0.029)	(0.042)
tenure	-0.000	-0.001	-0.002	-0.004*
	(0.001)	(0.001)	(0.002)	(0.002)
log(networksize)	0.004	-0.003	-0.006	-0.001
	(0.004)	(0.007)	(0.009)	(0.014)
achievement	-0.007	-0.001	0.002	0.024
	(0.008)	(0.014)	(0.018)	(0.037)
log(compensation)	-0.003	0.004	0.009	0.012
	(0.006)	(0.007)	(0.010)	(0.015)
shares owned	0.099	0.393**	0.722***	0.618**
	(0.086)	(0.183)	(0.218)	(0.255)
phd	0.002	-0.010	-0.042	-0.118*
	(0.033)	(0.052)	(0.054)	(0.071)
cpa	-0.013	-0.021	-0.025	-0.038
	(0.010)	(0.015)	(0.021)	(0.027)
sample size	3186	3186	3186	3186
R-squared	0.078	0.107	0.119	0.140

Notes: This table reports results for regression 3.1 for the insider trading profitability of top executives. The dependent variable, ret_Xm, is the X-month cumulative market-adjusted return following an insider purchases. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.7: Baseline results - size-value-momentum adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
African-American	-0.046	-0.131**	-0.189***	-0.255**
	(0.028)	(0.057)	(0.073)	(0.106)
log(mktcap)	-0.015***	-0.025***	-0.029***	-0.051***
	(0.003)	(0.007)	(0.008)	(0.015)
bm	0.025***	0.031**	0.046**	0.040
	(0.009)	(0.013)	(0.020)	(0.027)
leverage	0.002**	0.002	0.001	-0.000
	(0.001)	(0.001)	(0.002)	(0.003)
age	-0.001***	-0.002**	-0.005***	-0.007***
	(0.001)	(0.001)	(0.001)	(0.002)
male	0.022**	0.010	0.029	0.030
	(0.011)	(0.017)	(0.024)	(0.036)
tenure	0.000	-0.001	-0.002	-0.004*
	(0.001)	(0.001)	(0.001)	(0.002)
log(networksize)	0.001	-0.006	-0.009	-0.009
	(0.003)	(0.007)	(0.008)	(0.012)
achievement	-0.002	0.007	0.012	0.030
	(0.008)	(0.013)	(0.017)	(0.030)
log(compensation)	-0.003	0.002	0.006	0.004
	(0.005)	(0.007)	(0.010)	(0.013)
shares owned	0.177*	0.411*	0.676***	0.565**
	(0.107)	(0.214)	(0.240)	(0.282)
phd	0.018	0.000	-0.022	-0.095
	(0.033)	(0.052)	(0.052)	(0.069)
cpa	-0.012	-0.022	-0.023	-0.038
	(0.009)	(0.014)	(0.020)	(0.025)
sample size	3137	3134	3134	3132
R-squared	0.084	0.101	0.107	0.134

Notes: This table reports results for regression 3.1 for the insider trading profitability of top executives. The dependent variable, ret_Xm, is the X-month cumulative size-value-momentum-adjusted return following an insider purchases. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.8: Risk averseness - trading intensity

Variables	(1) log(shares)
African-American	-1.043
	(0.829)
log(mktcap)	-0.002
	(0.062)
bm	- 0.120
	(0.128)
leverage	-0.003
	(0.012)
age	-0.011
	(0.008)
male	0.502*
	(0.253)
tenure	-0.013
	(0.017)
log(networksize)	0.195**
	(0.058)
achievement	0.178
	(0.195)
log(compensation)	0.198**
	(0.080)
shares owned	0.027*
	(0.014)
phd	-0.145
	(0.362)
cpa	-0.057
	(0.192)
sample size	3186
R-squared	0.228

Notes: This table reports the results of the regression for testing the risk-averseness of African-American top executives. The dependent variable is the purchasing intensity of top executives, measured as the natural logarithm of the number of shares traded (log(shares)). We control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.9: Risk averseness - probability of trading before earnings announcement

	Probit		Logit	
	(1)	(2)	(3)	(4)
Variables	30-days	45-days	30-days	45-days
African-American	-0.037	0.183	-0.091	0.489
	(0.183)	(0.269)	(0.482)	(0.627)
log(mktcap)	-0.018	-0.024	-0.034	-0.050
	(0.019)	(0.020)	(0.048)	(0.048)
bm	-0.133*	-0.089	-0.320	-0.180
	(0.077)	(0.076)	(0.218)	(0.207)
leverage	0.002	-0.002	0.005	-0.007
	(0.005)	(0.005)	(0.012)	(0.012)
age	0.002	-0.001	0.006	-0.002
	(0.004)	(0.004)	(0.011)	(0.010)
male	-0.019	-0.011	-0.058	-0.040
	(0.097)	(0.097)	(0.254)	(0.247)
tenure	-0.006	-0.008	-0.015	-0.019
	(0.007)	(0.006)	(0.018)	(0.017)
log(networksize)	0.046*	0.032	0.113*	0.070
	(0.026)	(0.025)	(0.069)	(0.064)
achievement	-0.046	-0.066	-0.145	-0.189
	(0.066)	(0.070)	(0.173)	(0.180)
log(compensation)	-0.072*	-0.046	-0.201*	-0.128
	(0.042)	(0.043)	(0.107)	(0.109)
shares owned	-0.415	-0.660	-1.061	-1.608
	(1.093)	(1.045)	(3.008)	(2.758)
phd	-0.229*	-0.139	-0.589	-0.324
	(0.136)	(0.129)	(0.365)	(0.324)
cpa	-0.069	-0.127	-0.191	-0.345
	(0.090)	(0.089)	(0.234)	(0.224)
sample size	48970	48970	48970	48970
R-squared	0.044	0.044	0.045	0.045

Notes: This table reports the results of the regression for testing the risk-averseness of African-American top executives. We report probit and logit regression results for the trading probability of insiders before the earnings announcement dates where the dependent variables, 30-days and 45-days trading, are dummy variables that take the value of 1 if an insider trading is executed within 30-day and 45-day window before the earnings announcement date, respectively. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.10: Risk averseness - option backdating

Variables	(1) Probit	(2) Logit
African-American	-0.051	-0.091
	(0.169)	(0.318)
log(mktcap)	-0.040**	-0.077**
	(0.016)	(0.029)
bm	0.094**	0.165**
	(0.042)	(0.073)
leverage	0.012***	0.020***
	(0.004)	(0.007)
age	-0.002	-0.003
	(0.003)	(0.005)
male	0.021	0.050
	(0.063)	(0.116)
tenure	0.013***	0.024***
	(0.004)	(0.007)
log(networksize)	0.001	0.003
	(0.016)	(0.028)
achievement	-0.031	-0.053
	(0.045)	(0.081)
log(compensation)	0.016	0.032
	(0.028)	(0.050)
shares owned	-0.524	-0.854
	(0.777)	(1.398)
phd	0.105	0.176
	(0.119)	(0.215)
cpa	0.045	0.078
	(0.054)	(0.096)
sample size	7846	7846
R-squared	0.036	0.036

Notes: This table reports the results of the regression for testing the risk-averseness of African-American top executives. We report probit and logit regression results for testing the backdating of option awards where the dependent variable is a dummy variable that takes the value of unity if the insider is in the backdating sample. For each firm-insider pair in our sample, we compute the number of times options were granted at one of the three lowest stock prices during a 51-day window centered on reported grant dates. We then divide this number by the total number of option grant dates by the firm-insider pair during our sample period. If the resulting ratio is greater than 10%, we classify the insider as having engaged in backdating. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.11: Regression results with humanity scores - market adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
African-American	-0.053	-0.077	-0.230*	-0.216
	(0.052)	(0.047)	(0.124)	(0.149)
humanity	-0.005	-0.057***	-0.026	-0.049*
	(0.012)	(0.020)	(0.023)	(0.027)
African-American x humanity	0.155***	0.203***	0.348**	0.490***
	(0.059)	(0.063)	(0.140)	(0.175)
log(mktcap)	-0.017***	-0.022***	-0.025***	-0.031***
	(0.003)	(0.005)	(0.008)	(0.011)
bm	0.022**	0.045***	0.069***	0.076**
	(0.010)	(0.015)	(0.023)	(0.032)
leverage	0.001	0.002**	0.002	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
age	-0.001	-0.001	-0.004***	-0.004**
	(0.001)	(0.001)	(0.001)	(0.002)
male	0.028**	0.019	0.050*	0.037
	(0.012)	(0.020)	(0.027)	(0.038)
tenure	-0.001	-0.001	-0.002	-0.003
	(0.001)	(0.001)	(0.002)	(0.002)
log(networksize)	0.004	0.001	0.003	0.008
	(0.004)	(0.005)	(0.007)	(0.008)
achievement	-0.004	0.003	0.013	0.003
	(0.009)	(0.014)	(0.018)	(0.028)
log(compensation)	-0.001	0.008	0.012	0.004
	(0.007)	(0.008)	(0.011)	(0.015)
shares owned	0.118	0.391*	0.686***	0.536**
	(0.098)	(0.201)	(0.253)	(0.268)
phd	-0.005	-0.000	0.014	-0.022
	(0.028)	(0.051)	(0.055)	(0.065)
cpa	-0.004	-0.008	-0.008	-0.007
	(0.010)	(0.014)	(0.020)	(0.025)
sample size	2481	2481	2481	2481
R-squared	0.077	0.106	0.108	0.111

Notes: This table reports results for regression 3.2 for the insider trading profitability of top executives. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.12: Regression results with humanity scores - size-value-momentum adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
African-American	-0.043	-0.080*	-0.209*	-0.191
	(0.050)	(0.044)	(0.122)	(0.147)
humanity	-0.007	-0.055***	-0.032	-0.060**
	(0.013)	(0.020)	(0.024)	(0.029)
African-American x humanity	0.162***	0.237***	0.325**	0.403**
	(0.058)	(0.061)	(0.134)	(0.168)
log(mktcap)	-0.014***	-0.018***	-0.017**	-0.021**
	(0.003)	(0.005)	(0.007)	(0.009)
bm	0.026**	0.042***	0.065***	0.075**
	(0.010)	(0.015)	(0.023)	(0.033)
leverage	0.002***	0.003***	0.004***	0.004**
	(0.001)	(0.001)	(0.001)	(0.002)
age	-0.001**	-0.002	-0.004***	-0.005***
	(0.001)	(0.001)	(0.001)	(0.002)
male	0.025**	0.021	0.049**	0.036
	(0.012)	(0.017)	(0.023)	(0.033)
tenure	-0.001	-0.000	-0.002	-0.002
	(0.001)	(0.001)	(0.001)	(0.002)
log(networksize)	0.001	-0.004	-0.004	-0.003
	(0.004)	(0.005)	(0.007)	(0.007)
achievement	0.003	0.007	0.014	0.001
	(0.009)	(0.014)	(0.018)	(0.026)
log(compensation)	-0.004	0.004	0.007	-0.004
	(0.006)	(0.008)	(0.009)	(0.013)
shares owned	0.190	0.446*	0.698***	0.569**
	(0.130)	(0.228)	(0.265)	(0.277)
phd	0.010	0.011	0.020	-0.015
	(0.030)	(0.051)	(0.053)	(0.062)
cpa	-0.003	-0.005	0.002	0.005
	(0.009)	(0.015)	(0.021)	(0.024)
sample size	2454	2452	2452	2452
R-squared	0.072	0.086	0.096	0.100

Notes: This table reports results for regression 3.2 for the insider trading profitability of top executives. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.13: Regression results with employee scores - market adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
African-American	-0.071	-0.116***	-0.316***	-0.303**
	(0.058)	(0.035)	(0.100)	(0.133)
employee	0.005	0.008	0.009	0.012
	(0.005)	(0.008)	(0.007)	(0.010)
African-American x employee	0.032*	0.057***	0.129***	0.140***
	(0.018)	(0.014)	(0.031)	(0.043)
log(mktcap)	-0.018***	-0.026***	-0.029***	-0.036***
	(0.003)	(0.005)	(0.008)	(0.011)
bm	0.022**	0.043***	0.069***	0.075**
	(0.010)	(0.015)	(0.023)	(0.032)
leverage	0.001	0.002**	0.002*	0.002
	(0.001)	(0.001)	(0.001)	(0.001)
age	-0.001	-0.001	-0.004***	-0.004**
	(0.001)	(0.001)	(0.001)	(0.002)
male	0.029**	0.018	0.048*	0.035
	(0.012)	(0.021)	(0.027)	(0.039)
tenure	-0.001	-0.001	-0.002	-0.003
	(0.001)	(0.001)	(0.002)	(0.002)
log(networksize)	0.004	0.001	0.004	0.009
	(0.004)	(0.005)	(0.007)	(0.008)
achievement	-0.004	0.003	0.013	0.003
	(0.009)	(0.014)	(0.018)	(0.028)
log(compensation)	-0.001	0.007	0.011	0.004
	(0.007)	(0.008)	(0.011)	(0.015)
shares owned	0.124	0.398*	0.693***	0.547**
	(0.099)	(0.203)	(0.254)	(0.270)
phd	-0.007	-0.005	0.009	-0.030
	(0.028)	(0.051)	(0.055)	(0.066)
cpa	-0.004	-0.007	-0.008	-0.006
	(0.010)	(0.015)	(0.021)	(0.025)
sample size	2481	2481	2481	2481
R-squared	0.078	0.103	0.110	0.112

Notes: This table reports results for regression 3.2 for the insider trading profitability of top executives. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.14: Regression results with employee scores - size-value-momentum adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
African-American	-0.047	-0.112***	-0.275**	-0.267*
	(0.058)	(0.035)	(0.116)	(0.143)
employee	0.007	0.008	0.009	0.013
	(0.006)	(0.008)	(0.007)	(0.009)
African-American x employee	0.016	0.052***	0.103***	0.117***
	(0.019)	(0.015)	(0.035)	(0.044)
log(mktcap)	-0.015***	-0.022***	-0.021***	-0.026***
	(0.003)	(0.005)	(0.007)	(0.009)
bm	0.025**	0.040***	0.064***	0.073**
	(0.010)	(0.015)	(0.023)	(0.033)
leverage	0.002***	0.004***	0.004***	0.004**
	(0.001)	(0.001)	(0.001)	(0.002)
age	-0.001**	-0.002	-0.004***	-0.005**
	(0.001)	(0.001)	(0.001)	(0.002)
male	0.027**	0.021	0.048**	0.035
	(0.012)	(0.017)	(0.023)	(0.034)
tenure	-0.001	-0.000	-0.001	-0.002
	(0.001)	(0.001)	(0.001)	(0.002)
log(networksize)	0.001	-0.004	-0.003	-0.002
	(0.004)	(0.005)	(0.007)	(0.008)
achievement	0.003	0.007	0.014	0.002
	(0.009)	(0.013)	(0.017)	(0.026)
log(compensation)	-0.004	0.004	0.007	-0.004
	(0.006)	(0.008)	(0.009)	(0.013)
shares owned	0.198	0.455**	0.706***	0.582**
	(0.131)	(0.231)	(0.267)	(0.280)
phd	0.007	0.005	0.015	-0.022
	(0.030)	(0.052)	(0.054)	(0.063)
cpa	-0.003	-0.004	0.003	0.007
	(0.009)	(0.015)	(0.021)	(0.024)
sample size	2454	2452	2452	2452
R-squared	0.074	0.084	0.097	0.101

Notes: This table reports results for regression 3.2 for the insider trading profitability of top executives. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.15: Regression results with diversity scores - market adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
African-American	-0.040	-0.071*	-0.226**	-0.207
	(0.049)	(0.042)	(0.105)	(0.127)
diversity	-0.001	0.003	0.001	-0.002
	(0.003)	(0.005)	(0.008)	(0.010)
African-American x diversity	0.000	0.014	0.055***	0.069***
	(0.021)	(0.011)	(0.019)	(0.023)
log(mktcap)	-0.016***	-0.025***	-0.026***	-0.032***
	(0.004)	(0.005)	(0.008)	(0.011)
bm	0.022**	0.043***	0.068***	0.074**
	(0.010)	(0.015)	(0.023)	(0.032)
leverage	0.001	0.002**	0.002*	0.002
	(0.001)	(0.001)	(0.001)	(0.001)
age	-0.001	-0.001	-0.004***	-0.004**
	(0.001)	(0.001)	(0.001)	(0.002)
male	0.028**	0.018	0.047*	0.033
	(0.012)	(0.021)	(0.027)	(0.038)
tenure	-0.001	-0.001	-0.002	-0.003
	(0.001)	(0.001)	(0.002)	(0.002)
log(networksize)	0.004	0.001	0.003	0.008
	(0.004)	(0.006)	(0.007)	(0.008)
achievement	-0.004	0.003	0.014	0.004
	(0.009)	(0.014)	(0.018)	(0.028)
log(compensation)	-0.001	0.007	0.012	0.005
	(0.007)	(0.008)	(0.011)	(0.015)
shares owned	0.119	0.390*	0.687***	0.538**
	(0.098)	(0.202)	(0.253)	(0.270)
phd	-0.005	-0.002	0.013	-0.024
	(0.028)	(0.052)	(0.055)	(0.067)
cpa	-0.004	-0.008	-0.008	-0.006
	(0.010)	(0.014)	(0.020)	(0.025)
sample size	2481	2481	2481	2481
R-squared	0.076	0.102	0.108	0.110

Notes: This table reports results for regression 3.2 for the insider trading profitability of top executives. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.16: Regression results with diversity scores - size-value-momentum adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
African-American	-0.029	-0.073*	-0.201*	-0.181
	(0.048)	(0.040)	(0.107)	(0.129)
diveristy	0.002	0.004	0.002	0.000
	(0.003)	(0.006)	(0.007)	(0.010)
African-American x diversity	-0.002	0.018	0.040**	0.049**
	(0.022)	(0.011)	(0.018)	(0.023)
log(mktcap)	-0.014***	-0.021***	-0.019***	-0.022**
	(0.003)	(0.005)	(0.007)	(0.009)
bm	0.025**	0.040***	0.063***	0.073**
	(0.010)	(0.015)	(0.023)	(0.033)
leverage	0.002***	0.003***	0.004***	0.004**
	(0.001)	(0.001)	(0.001)	(0.002)
age	-0.001**	-0.002	-0.004***	-0.005**
	(0.001)	(0.001)	(0.001)	(0.002)
male	0.026**	0.020	0.047**	0.033
	(0.011)	(0.017)	(0.023)	(0.033)
tenure	-0.001	-0.000	-0.002	-0.002
	(0.001)	(0.001)	(0.001)	(0.002)
log(networksize)	0.001	-0.004	-0.004	-0.003
	(0.004)	(0.005)	(0.007)	(0.008)
achievement	0.003	0.008	0.014	0.002
	(0.009)	(0.014)	(0.017)	(0.026)
log(compensation)	-0.004	0.003	0.006	-0.004
	(0.006)	(0.007)	(0.009)	(0.013)
shares owned	0.190	0.446*	0.698***	0.570**
	(0.130)	(0.229)	(0.265)	(0.279)
phd	0.010	0.009	0.019	-0.017
	(0.030)	(0.052)	(0.054)	(0.064)
cpa	-0.003	-0.005	0.003	0.007
	(0.009)	(0.015)	(0.021)	(0.024)
sample size	2454	2452	2452	2452
R-squared	0.072	0.082	0.095	0.098

Notes: This table reports results for regression 3.2 for the insider trading profitability of top executives. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.17: Insider Profitability for Asian-Americans - market adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
Asian-American	0.021	0.054	0.058*	0.078**
	(0.026)	(0.034)	(0.033)	(0.037)
log(mktcap)	-0.015***	-0.025***	-0.035***	-0.061***
	(0.004)	(0.007)	(0.010)	(0.019)
bm	0.023**	0.042***	0.51**	0.48
	(0.011)	(0.014)	(0.021)	(0.029)
leverage	0.000	0.002***	0.004***	0.004**
	(0.001)	(0.001)	(0.001)	(0.002)
age	-0.001	-0.002	-0.005***	-0.008***
	(0.001)	(0.001)	(0.002)	(0.003)
male	0.026**	0.016	0.036	0.050
	(0.012)	(0.022)	(0.033)	(0.048)
tenure	-0.000	-0.001	-0.001	-0.003
	(0.001)	(0.001)	(0.002)	(0.002)
log(networksize)	0.003	-0.008	-0.010	-0.005
	(0.004)	(0.008)	(0.011)	(0.017)
achievement	-0.003	0.001	0.001	0.029
	(0.009)	(0.014)	(0.018)	(0.041)
log(compensation)	-0.004	0.003	0.009	0.010
	(0.006)	(0.007)	(0.011)	(0.016)
shares owned	0.100	0.362*	0.736***	0.645**
	(0.090)	(0.188)	(0.214)	(0.255)
phd	-0.005	-0.008	-0.048	-0.139*
	(0.032)	(0.055)	(0.059)	(0.080)
cpa	-0.007	-0.010	-0.001	-0.005
	(0.011)	(0.017)	(0.024)	(0.032)
sample size	2674	2674	2674	2674
R-squared	0.082	0.122	0.134	0.159

Notes: This table reports results for regression 3.1 for the insider trading profitability of top Asian-American executives. The dependent variable, ret_Xm, is the X-month cumulative market-adjusted return following an insider purchases. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.18: Insider Profitability for Asian-Americans - size-value-momentum adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
Asian-American	-0.012	0.019	0.029	0.055
	(0.027)	(0.034)	(0.039)	(0.045)
log(mktcap)	-0.015***	-0.025***	-0.028***	-0.051***
	(0.003)	(0.007)	(0.009)	(0.015)
bm	0.028***	0.036**	0.054**	0.051*
	(0.009)	(0.014)	(0.021)	(0.030)
leverage	0.001**	0.001	0.000	-0.002
	(0.001)	(0.001)	(0.002)	(0.002)
age	-0.001**	-0.003**	-0.006***	-0.008***
	(0.001)	(0.001)	(0.002)	(0.003)
male	0.024**	0.013	0.033	0.043
	(0.011)	(0.019)	(0.027)	(0.041)
tenure	-0.000	-0.000	-0.001	-0.003
	(0.001)	(0.001)	(0.002)	(0.002)
log(networksize)	0.002	-0.009	-0.012	-0.011
	(0.004)	(0.008)	(0.010)	(0.015)
achievement	0.002	0.011	0.014	0.036
	(0.009)	(0.014)	(0.018)	(0.031)
log(compensation)	-0.004	0.002	0.006	0.002
	(0.005)	(0.008)	(0.011)	(0.014)
shares owned	0.186*	0.385*	0.693***	0.607**
	(0.108)	(0.221)	(0.234)	(0.280)
phd	0.012	-0.002	-0.031	-0.121
	(0.033)	(0.056)	(0.056)	(0.079)
cpa	-0.005	-0.006	0.002	-0.008
	(0.009)	(0.016)	(0.023)	(0.030)
sample size	2629	2626	2626	2624
R-squared	0.088	0.116	0.121	0.156

Notes: This table reports results for regression 3.1 for the insider trading profitability of top Asian-American executives. The dependent variable, ret_Xm, is the X-month cumulative size-value-momentum-adjusted return following an insider purchases. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.19: Nearest-neighbor matching results

market_adj_3m	difference	10.3%
	t-stat	3.433
market_adj_6m	difference	11.6%
	t-stat	2.974
market_adj_9m	difference	15.5%
	t-stat	1.631
market_adj_12m	difference	19.7%
	t-stat	1.807
dgtw_3m	difference	7.3%
	t-stat	2.703
dgtw_6m	difference	8.1%
	t-stat	1.653
dgtw_9m	difference	10.7%
	t-stat	0.811
dgtw_12m	difference	12.9%
	t-stat	0.843

Notes: This table reports the nearest-neighbor matching results for the insider trading profitability of top executives. The reported differences are the returns of non-African American executives minus the returns of African American executives following their insider purchases. Matching algorithm pairs each observation from one ethnicity group with the "closest" or the "most similar" observation from the other group. Similarity measure is determined based on several covariates including market value (log(mktcap)), book-to-market ratio (bm), leverage, gender (male), age, tenure, log(networksize), achievement, log(compensation), shares owned, phd, cpa, humanity, employee, diversity, and insider role and year dummies. The t-stats are reported below the differences. market_adj_Xm denotes the X-month market adjusted return. dgtw_Xm denotes the X-month size-momentum-value adjusted returns.

Table 3.20: Baseline results with directors - market adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
African-American	-0.022*	-0.044***	-0.073***	-0.074**
	(0.012)	(0.017)	(0.025)	(0.029)
log(mktcap)	-0.023***	-0.026***	-0.034***	-0.051***
	(0.005)	(0.005)	(0.007)	(0.009)
bm	0.007	0.021*	0.043**	0.052**
	(0.009)	(0.011)	(0.020)	(0.022)
leverage	-0.001	-0.000	-0.001	-0.000
	(0.001)	(0.001)	(0.001)	(0.001)
age	-0.000	-0.000	-0.001	0.000
	(0.000)	(0.001)	(0.001)	(0.001)
male	-0.003	0.001	0.005	0.006
	(0.007)	(0.010)	(0.013)	(0.016)
tenure	0.000	0.001	0.001	0.000
	(0.000)	(0.001)	(0.001)	(0.001)
log(networksize)	-0.000	-0.005	-0.004	-0.001
	(0.003)	(0.004)	(0.005)	(0.006)
achievement	-0.006	0.005	0.009	-0.004
	(0.006)	(0.008)	(0.011)	(0.013)
log(compensation)	0.004	-0.004	-0.006	-0.005
	(0.003)	(0.004)	(0.006)	(0.007)
log(shares owned)	0.003	0.005**	0.008**	0.007*
	(0.002)	(0.003)	(0.004)	(0.005)
phd	0.009	0.026	0.026	0.013
	(0.010)	(0.022)	(0.029)	(0.028)
cpa	0.021	0.016	0.015	0.012
	(0.014)	(0.015)	(0.020)	(0.024)
sample size	4288	4288	4287	4287
R-squared	0.082	0.080	0.099	0.111

Notes: This table reports results for regression 3.1 for the insider trading profitability of directors. The dependent variable, ret_Xm, is the X-month cumulative market-adjusted return following an insider purchases. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Table 3.21: Baseline results with directors- size-value-momentum adjusted return

Variables	(1) ret_3m	(2) ret_6m	(3) ret_9m	(4) ret_12m
African-American	-0.022*	-0.041**	-0.071***	-0.070**
	(0.012)	(0.016)	(0.024)	(0.028)
log(mktcap)	-0.021***	-0.025***	-0.032***	-0.047***
	(0.004)	(0.004)	(0.006)	(0.008)
bm	0.005	0.017*	0.033*	0.032*
	(0.007)	(0.010)	(0.017)	(0.017)
leverage	-0.001	-0.001	-0.002	-0.002
	(0.001)	(0.001)	(0.001)	(0.001)
age	-0.000	-0.000	-0.000	0.000
	(0.000)	(0.001)	(0.001)	(0.001)
male	0.003	0.003	0.009	0.010
	(0.008)	(0.011)	(0.013)	(0.016)
tenure	-0.000	0.000	-0.000	-0.000
	(0.000)	(0.001)	(0.001)	(0.001)
log(networksize)	0.001	-0.002	-0.004	-0.003
	(0.002)	(0.003)	(0.005)	(0.005)
achievement	-0.005	0.006	0.015	0.004
	(0.006)	(0.008)	(0.011)	(0.013)
log(compensation)	0.002	-0.005	-0.006	-0.004
	(0.003)	(0.004)	(0.006)	(0.006)
log(shares owned)	0.002	0.003	0.004	0.004
	(0.002)	(0.003)	(0.004)	(0.004)
phd	0.007	0.014	0.013	0.005
	(0.013)	(0.025)	(0.030)	(0.028)
cpa	0.013	0.009	-0.001	-0.006
	(0.010)	(0.011)	(0.014)	(0.018)
sample size	4223	4223	4218	4216
R-squared	0.068	0.063	0.078	0.090

Notes: This table reports results for regression 3.1 for the insider trading profitability of directors. The dependent variable, ret_Xm, is the X-month cumulative size-value-momentum-adjusted return following an insider purchases. In all columns, we control for year, insider role, and industry fixed effects. Standard errors clustered at the insider level are reported below coefficient estimates in parentheses. ***, ** and * indicate that the coefficient estimate is significant at the 1%, 5%, and 10% levels, respectively.

Bibliography

[1] Avidit Acharya, Matthew Blackwell, and Maya Sen. The political legacy of american slavery. *The Journal of Politics*, 78(3):621–641, 2016.

[2] Biljana Adebambo, Paul Brockman, and Xuemin Sterling Yan. Anticipating the 2007–2008 financial crisis: Who knew what and when did they know it? *Journal of financial and quantitative analysis*, 50(4):647–669, 2015.

[3] Alan G Ahearne, William L Griever, and Francis E Warnock. Information costs and home bias: an analysis of us holdings of foreign equities. *Journal of international economics*, 62(2):313–336, 2004.

[4] Donald WK Andrews. Tests for parameter instability and structural change with unknown change point. *Econometrica: Journal of the Econometric Society*, pages 821–856, 1993.

[5] Deniz Anginer, MP Narayanan, Cindy A Schipani, and H Nejat Seyhun. Should size matter when regulating firms-implications from backdating of executive options. *NYUJ Legis. & Pub. Pol'y*, 15:1, 2012.

[6] Andreas M Antonopoulos and Gavin Wood. *Mastering ethereum: building smart contracts and dapps*. O'reilly Media, 2018.

[7] Jasmina Arifovic and John Ledyard. Call market book information and efficiency. *Journal of Economic Dynamics and Control*, 31(6):1971–2000, 2007.

[8] Jasmina Arifovic and John Ledyard. Individual evolutionary learning, other-regarding preferences, and the voluntary contributions mechanism. *Journal of Public Economics*, 96(9-10):808–823, 2012.

[9] Jasmina Arifovic and John Ledyard. Learning to alternate. *Experimental Economics*, 21(3):692–721, 2018.

[10] Kenneth J Arrow. Some mathematical models of race discrimination in the labor market. *Racial discrimination in economic life*, pages 187–204, 1972.

[11] Kenneth J Arrow and Ron Borzekowski. Limited network connections and the distribution of wages. *Available at SSRN 632321*, 2004.

[12] Susan Athey, Ivo Parashkevov, Vishnu Sarukkai, and Jing Xia. Bitcoin pricing, adoption, and usage: Theory and evidence. 2016.

[13] S Burcu Avci, Cindy A Schipani, and H Nejat Seyhun. Manipulative games of gifts by corporate executives. *U. Pa. J. Bus. L.*, 18:1131, 2015.

[14] Sureyya Burcu Avci, Cindy A Schipani, H Nejat Seyhun, and Andrew Verstein. Insider giving. *Andrew, Insider Giving*, 71:21–02, 2021.

[15] Ian Ayres and Peter Siegelman. Race and gender discrimination in bargaining for a new car. *The American Economic Review*, pages 304–321, 1995.

[16] Robert Bartlett, Adair Morse, Richard Stanton, and Nancy Wallace. Consumer-lending discrimination in the fintech era. *Journal of Financial Economics*, 143(1):30–56, 2022.

[17] Gary S Becker. *The economics of discrimination*. University of Chicago press, 2010.

[18] Taylor A Begley and Amiyatosh Purnanandam. Color and credit: Race, regulation, and the quality of financial services. *Journal of Financial Economics*, 141(1):48–65, 2021.

[19] Marianne Bertrand and Sendhil Mullainathan. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American economic review*, 94(4):991–1013, 2004.

[20] John Bizjak, Michael Lemmon, and Ryan Whitby. Option backdating and board interlocks. *The Review of Financial Studies*, 22(11):4821–4847, 2009.

[21] Peter M Blau. A macrosociological theory of social structure. *American journal of sociology*, 83(1):26–54, 1977.

[22] Christiane Bode, Jasjit Singh, and Michelle Rogan. Corporate social initiatives and employee retention. *Organization Science*, 26(6):1702–1720, 2015.

[23] Lukas Bolte, Nicole Immorlica, and Matthew O Jackson. The role of referrals in inequality, immobility, and inefficiency in labor markets. *Immobility, and Inefficiency in Labor Markets (January 1, 2020)*, 2020.

[24] Vitalik Buterin et al. Ethereum white paper. *GitHub repository*, 1:22–23, 2013.

[25] Jing Cai and Adam Szeidl. Interfirm relationships and business performance. *The Quarterly Journal of Economics*, 133(3):1229–1282, 2018.

[26] Colin Camerer and Teck Hua Ho. Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874, 1999.

[27] Ying Cao, Feng Guan, Zengquan Li, and Yong George Yang. Analysts' beauty and performance. *Management Science*, 66(9):4315–4335, 2020.

[28] Timothy N Cason and Daniel Friedman. Price formation in double auction markets. *Journal of Economic Dynamics and Control*, 20(8):1307–1337, 1996.

[29] Lea Cassar and Stephan Meier. Nonmonetary incentives and the implications of work as a source of meaning. *Journal of Economic Perspectives*, 32(3):215–38, 2018.

[30] Christian Catalini and Joshua S Gans. Some simple economics of the blockchain. *Communications of the ACM*, 63(7):80–90, 2020.

[31] Ting Chen, Hagit Levy, Xiumin Martin, and Ron Shalev. Buying products from whom you know: personal connections and information asymmetry in supply chain relationships. *Review of Accounting Studies*, 26(4):1492–1531, 2021.

[32] Raj Chetty, John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan. Income segregation and intergenerational mobility across colleges in the united states. *The Quarterly Journal of Economics*, 135(3):1567–1633, 2020.

[33] Jonathan Chiu and Thorsten V Koeppl. Blockchain-based settlement for asset trading. *The Review of Financial Studies*, 32(5):1716–1753, 2019.

[34] Pavel Ciaian, Miroslava Rajcaniova, and d'Artis Kancs. The economics of bitcoin price formation. *Applied economics*, 48(19):1799–1815, 2016.

[35] Dave Cliff and Janet Bruten. Less than human: Simple adaptive trading agents for cda markets. *IFAC Proceedings Volumes*, 31(16):117–122, 1998.

[36] Lauren Cohen, Christopher Malloy, and Lukasz Pomorski. Decoding inside information. *The Journal of Finance*, 67(3):1009–1043, 2012.

[37] Lin William Cong and Zhiguo He. Blockchain disruption and smart contracts. *The Review of Financial Studies*, 32(5):1754–1797, 2019.

[38] Shaen Corbet, Brian Lucey, and Larisa Yarovaya. Datestamping the bitcoin and ethereum bubbles. *Finance Research Letters*, 26:81–88, 2018.

[39] Joshua D Coval and Tobias J Moskowitz. Home bias at home: Local equity preference in domestic portfolios. *The Journal of Finance*, 54(6):2045–2073, 1999.

[40] Sean Crockett. Price dynamics in general equilibrium experiments. *Journal of Economic Surveys*, 27(3):421–438, 2013.

[41] Sergio Currarini, Matthew O Jackson, and Paolo Pin. An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, 77(4):1003–1045, 2009.

[42] Melville Dalton. *Men who manage: Fusions of feeling and theory in administration*. Routledge, 2017.

[43] Kent Daniel, Mark Grinblatt, Sheridan Titman, and Russ Wermers. Measuring mutual fund performance with characteristic-based benchmarks. *The Journal of finance*, 52(3):1035–1058, 1997.

[44] Rajeev H Dehejia and Sadek Wahba. Propensity score-matching methods for nonexperimental causal studies. *Review of Economics and statistics*, 84(1):151–161, 2002.

[45] John Duffy and M Utku Ünver. Asset price bubbles and crashes with near-zero-intelligence traders. *Economic theory*, 27(3):537–563, 2006.

[46] David Easley and John O Ledyard. Theories of price formation and exchange in double oral auctions. In *The Double Auction Market Institutions, Theories, and Evidence*, pages 63–98. Routledge, 2018.

[47] David Easley, Maureen O'Hara, and Soumya Basu. From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics*, 134(1):91–109, 2019.

[48] Ido Erev and Alvin E Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, pages 848–881, 1998.

[49] Scott L Feld. The focused organization of social ties. *American journal of sociology*, 86(5):1015–1035, 1981.

[50] Caroline Flammer. Does corporate social responsibility lead to superior financial performance? a regression discontinuity approach. *Management Science*, 61(11):2549–2568, 2015.

[51] Caroline Flammer and Jiao Luo. Corporate social responsibility as an employee governance tool: Evidence from a quasi-experiment. *Strategic Management Journal*, 38(2):163–183, 2017.

[52] Sean Foley, Jonathan R Karlsen, and Tālis J Putniņš. Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? *The Review of Financial Studies*, 32(5):1798–1853, 2019.

[53] Kenneth R French and James M Poterba. Investor diversification and international equity markets, 1991.

[54] Daniel Friedman. A simple testable model of double auction markets. *Journal of Economic Behavior & Organization*, 15(1):47–70, 1991.

[55] Neil Gandal, JT Hamrick, Tyler Moore, and Tali Oberman. Price manipulation in the bitcoin ecosystem. *Journal of Monetary Economics*, 95:86–96, 2018.

[56] Herbert J Gans. The possibility of a new racial hierarchy in the twenty-first-century united states. In *The Inequality Reader*, pages 304–313. Routledge, 2018.

[57] Steven Gjerstad and John Dickhaut. Price formation in double auctions. *Games and economic behavior*, 22(1):1–29, 1998.

[58] Steven Gjerstad, Jason M Shachat, et al. Individual rationality and market efficiency. Technical report, Purdue University, Department of Economics, 2007.

[59] Dhananjay Dan K Gode and Shyam Sunder. Double auction dynamics: structural effects of non-binding price controls. *Journal of economic dynamics and control*, 28(9):1707–1731, 2004.

[60] Dhananjay K Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of political economy*, 101(1):119–137, 1993.

[61] Dhananjay K Gode and Shyam Sunder. What makes markets allocationally efficient? *The Quarterly Journal of Economics*, 112(2):603–630, 1997.

[62] Benjamin Golub and Matthew O Jackson. How homophily affects the speed of learning and best-response dynamics. *The Quarterly Journal of Economics*, 127(3):1287–1338, 2012.

[63] John M Griffin and Amin Shams. Is bitcoin really untethered? *The Journal of Finance*, 75(4):1913–1964, 2020.

[64] Judith K Hellerstein and David Neumark. Workplace segregation in the united states: Race, ethnicity, and skill. *The review of economics and statistics*, 90(3):459–477, 2008.

[65] Randall A Heron and Erik Lie. Does backdating explain the stock price pattern around executive stock option grants? *Journal of Financial Economics*, 83(2):271–295, 2007.

[66] Charles A Holt. Industrial organization: A survey of laboratory research. *The handbook of experimental economics*, 349:402–03, 1995.

[67] Karel Hrazdil, Fereshteh Mahmoudian, and Jamal A Nazari. Executive personality and sustainability: Do extraverted chief executive officers improve corporate social responsibility? *Corporate Social Responsibility and Environmental Management*, 28(6):1564–1578, 2021.

[68] Tien-Shih Hsieh, Jeong-Bon Kim, Ray R Wang, and Zhihong Wang. Seeing is believing? executives' facial trustworthiness, auditor tenure, and audit fees. *Journal of Accounting and Economics*, 69(1):101260, 2020.

[69] Gur Huberman, Jacob Leshno, and Ciamac C Moallemi. An economic analysis of the bitcoin payment system. *Columbia Business School Research Paper*, (17-92), 2019.

[70] Steven Huddart, Bin Ke, and Charles Shi. Jeopardy, non-public information, and insider trading around sec 10-k and 10-q filings. *Journal of Accounting and Economics*, 43(1):3–36, 2007.

[71] Herminia Ibarra. Personal networks of women and minorities in management: A conceptual framework. *Academy of management Review*, 18(1):56–87, 1993.

[72] Herminia Ibarra. Race, opportunity, and diversity of social circles in managerial networks. *Academy of management journal*, 38(3):673–703, 1995.

[73] A Can Inci, MP Narayanan, and H Nejat Seyhun. Gender differences in executives' access to information. *Journal of Financial and Quantitative Analysis*, 52(3):991–1016, 2017.

[74] Yannis M Ioannides and Linda Datcher Loury. Job information networks, neighborhood effects, and inequality. *Journal of economic literature*, 42(4):1056–1093, 2004.

[75] Jeffrey F Jaffe. Special information and insider trading. *The Journal of Business*, 47(3):410–428, 1974.

[76] Alan D Jagolinzer, David F Larcker, Gaizka Ormazabal, and Daniel J Taylor. Political connections and the informativeness of insider trades. *The Journal of Finance*, 75(4):1833–1876, 2020.

[77] Rosabeth Moss Kanter. *The change masters: innovation for productivity in the American mode.* Simon and schuster, 1983.

[78] Alexander K Karajvanov. *Blockchains, collateral and financial contracts.* Simon Fraser University, Department of Economics, 2021.

[79] Shoji Kasahara and Jun Kawahara. Effect of bitcoin fee on transaction-confirmation process. *arXiv preprint arXiv:1604.00103*, 2016.

[80] Seoung Kyun Kim, Zane Ma, Siddharth Murali, Joshua Mason, Andrew Miller, and Michael Bailey. Measuring ethereum network peers. In *Proceedings of the Internet Measurement Conference 2018*, pages 91–104, 2018.

[81] Sanders Korenman and Susan C Turner. Employment contacts and minority-white wage differences. *Industrial Relations: A Journal of Economy and Society*, 35(1):106–122, 1996.

[82] Ladislav Kristoufek. What are the main drivers of the bitcoin price? evidence from wavelet coherence analysis. *PloS one*, 10(4):e0123923, 2015.

[83] Philipp Krüger. Corporate goodness and shareholder wealth. *Journal of financial economics*, 115(2):304–329, 2015.

[84] Joanna N Lahey. Age, women, and hiring an experimental study. *Journal of Human resources*, 43(1):30–56, 2008.

[85] Josef Lakonishok and Inmoo Lee. Are insider trades informative? *The Review of Financial Studies*, 14(1):79–111, 2001.

[86] Marie Lalanne and Paul Seabright. The old boy network: Gender differences in the impact of social networks on remuneration in top executive jobs. *Available at SSRN 1952484*, 2011.

[87] Kevin Lang and Ariella Kahn-Lang Spitzer. Race discrimination: An economic perspective. *Journal of Economic Perspectives*, 34(2):68–89, 2020.

[88] Jennifer Lee and Frank D Bean. Reinventing the color line immigration and america's new racial/ethnic divide. *Social Forces*, 86(2):561–586, 2007.

[89] Quan-Lin Li, Jing-Yu Ma, Yan-Xia Chang, Fan-Qi Ma, and Hai-Bo Yu. Markov processes in blockchain systems. *Computational Social Networks*, 6(1):1–28, 2019.

[90] Tao Li, Donghwa Shin, and Baolian Wang. Cryptocurrency pump-and-dump schemes. *Available at SSRN 3267041*, 2021.

[91] Po-Hsuan Lin, Alexander L Brown, Taisuke Imai, Joseph Tao-yi Wang, Stephanie W Wang, and Colin F Camerer. Evidence of general economic principles of bargaining and trade from 2,000 classroom experiments. *Nature Human Behaviour*, 4(9):917–927, 2020.

[92] Ilse Lindenlaub and Anja Prummer. Gender, social networks and peformance. Technical report, Working Paper, 2016.

[93] Valentina Marano and Tatiana Kostova. Unpacking the institutional complexity in adoption of csr practices in multinational enterprises. *Journal of Management Studies*, 53(1):28–54, 2016.

[94] Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.

[95] Friederike Mengel. Gender differences in networking. *The Economic Journal*, 130(630):1842–1873, 2020.

[96] Philip Mirvis. Employee engagement and csr: Transactional, relational, and developmental approaches. *California management review*, 54(4):93–117, 2012.

[97] Malte Möser and Rainer Böhme. Trends, tips, tolls: A longitudinal study of bitcoin transaction fees. In *International conference on financial cryptography and data security*, pages 19–33. Springer, 2015.

[98] MP Narayanan, Cindy A Schipani, and HNejat Seyhun. The economic impact of backdating of executive stock options. *Michigan Law Review*, 105(8):1597–1641, 2007.

[99] MP Narayanan and H Nejat Seyhun. The dating game: do managers designate option grant dates to increase their compensation? *The Review of Financial Studies*, 21(5):1907–1945, 2008.

[100] Virginia E O'Leary and Jeannette R Ickovics. Cracking the glass ceiling: Overcoming isolation and alienation. 1992.

[101] Lin Peng, Siew Hong Teoh, Yakun Wang, and Jiawen Yan. Face value: Trait inference, performance characteristics, and market outcomes for financial analysts. *Journal of Accounting Research Forthcoming*, 2021.

[102] Edmund S Phelps. The statistical theory of racism and sexism. *The american economic review*, 62(4):659–661, 1972.

[103] Charles R Plott. Properties of disequilibrium adjustment in double auction markets. *Handbook of experimental economics results*, 1:16–21, 2008.

[104] Devin G Pope and Justin R Sydnor. What's in a picture? evidence of discrimination from prosper. com. *Journal of Human resources*, 46(1):53–92, 2011.

[105] Amiyatosh Purnanandam and H Nejat Seyhun. Do short sellers trade on private information or false information? *Journal of Financial and Quantitative Analysis*, 53(3):997–1023, 2018.

[106] Enrichetta Ravina. Love & loans: The effect of beauty and personal characteristics in credit markets. *Available at SSRN 1107307*, 2019.

[107] William Robert Reed. On the practice of lagging variables to avoid simultaneity. *Oxford Bulletin of Economics and Statistics*, 77(6):897–905, 2015.

[108] H Nejat Seyhun. Insiders' profits, costs of trading, and market efficiency. *Journal of financial Economics*, 16(2):189–212, 1986.

[109] H Nejat Seyhun. The information content of aggregate insider trading. *Journal of Business*, pages 1–24, 1988.

[110] H Nejat Seyhun. Overreaction or fundamentals: Some lessons from insiders' response to the market crash of 1987. *The Journal of Finance*, 45(5):1363–1388, 1990.

[111] H Nejat Seyhun. Why does aggregate insider trading predict future stock returns? *The Quarterly Journal of Economics*, 107(4):1303–1331, 1992.

[112] H Nejat Seyhun. *Investment intelligence from insider trading.* MIT press, 2000.

[113] Jie Shen and John Benson. When csr is a social norm: How socially responsible human resource management affects employee work behavior. *Journal of management*, 42(6):1723–1746, 2016.

[114] John F Shortle, James M Thompson, Donald Gross, and Carl M Harris. *Fundamentals of queueing theory*, volume 399. John Wiley & Sons, 2018.

[115] Vernon L Smith. An experimental study of competitive market behavior. *Journal of political economy*, 70(2):111–137, 1962.

[116] Douglas O Staiger and James H Stock. Instrumental variables regression with weak instruments, 1994.

[117] Shinji Takagi. Pegging to a currency basket. *Finance and Development*, 23(3):41, 1986.

[118] Linda L Tesar and Ingrid M Werner. Home bias and high turnover. *Journal of international Money and Finance*, 14(4):467–492, 1995.

[119] Robert M Townsend. *Distributed Ledgers: Design and Regulation of Financial Infrastructure and Payment Systems.* MIT Press, 2020.

[120] Michiel van de Leur and Mikhail Anufriev. Timing under individual evolutionary learning in a continuous double auction. *Journal of Evolutionary Economics*, 28(3):609–631, 2018.

[121] Roger Waldinger and Michael I Lichter. How the other half works. In *How the Other Half Works.* University of California Press, 2003.

[122] Robert B Wilson. On equilibria of bid-ask markets. In *Arrow and the ascent of modern economic theory*, pages 375–414. Springer, 1987.

[123] David Yermack. Is bitcoin a real currency? an economic appraisal. In *Handbook of digital currency*, pages 31–43. Elsevier, 2015.

Appendix A

Chapter 1

A.1 Additional tables and figures

Figure A.1: Average ETH price in terms of USD, EUR and CNY basket of currencies (BoC)

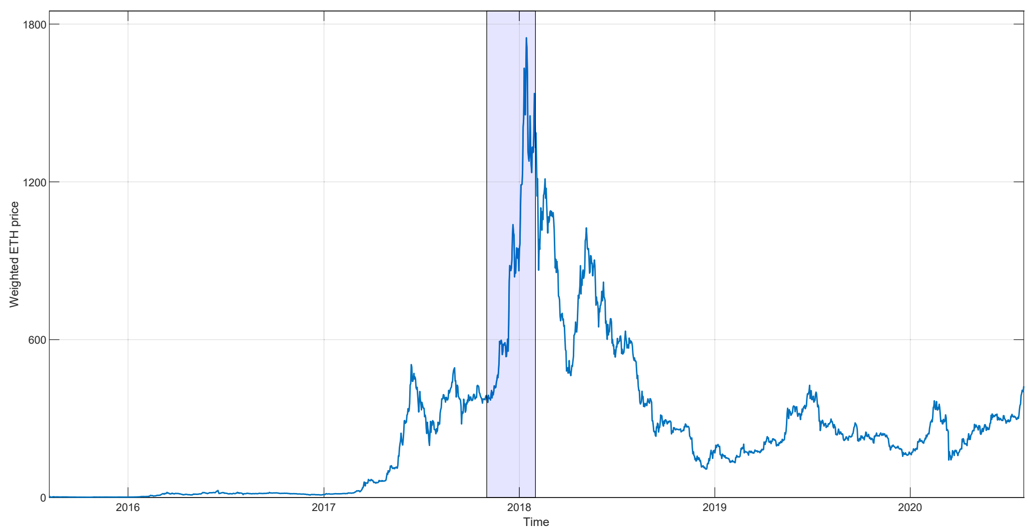

Notes: The Figure plots the daily average weighted exchange rate between ETH and a basket of currencies (BoC) consisting of US dollars (USD), Euro (EUR), and Chinese Yuan (CNY) (see Appendix A.2). The shaded region highlights the peak period between Nov 1, 2017 and Jan 31, 2018.

Figure A.2: New sender accounts

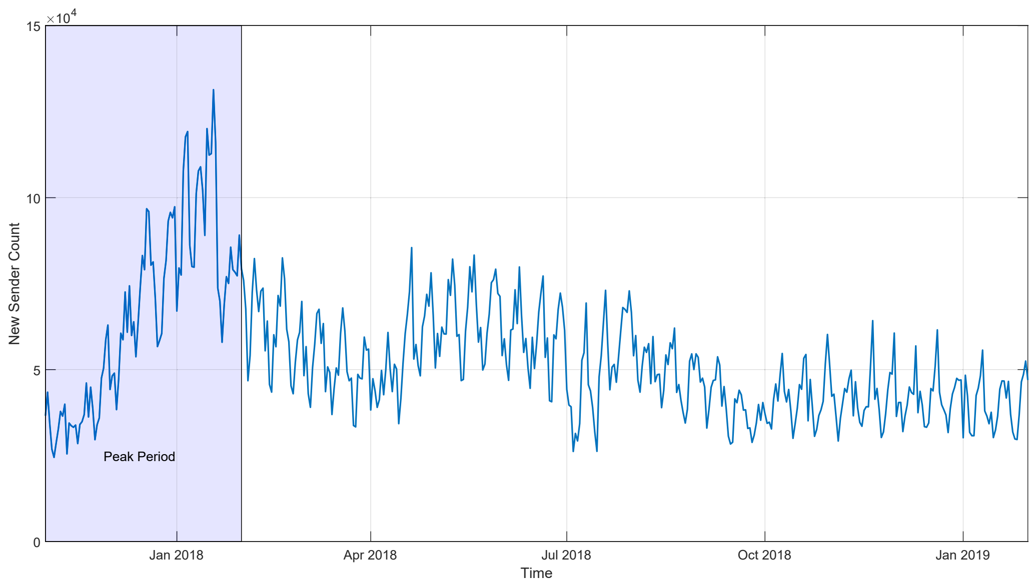

Notes: The Figure plots the daily number of new sender accounts, defined as the daily number of transactions with nonce value equal to 0. The nonce field in the transaction-level data records the past transaction count of a sender address. The shaded region highlights the peak period between Nov. 1, 2017 and Jan. 31, 2018.

Table A.1: Summary statistics

Variables	# of obs.	Mean	Std	25^{th} pctile	50^{th} pctile	75^{th} pctile
	Full sample, Nov. 1, 2017 - Jan. 31, 2019					
marginal gas price, 10^{-8} ETH	457	0.82	0.89	0.43	0.55	0.76
marginal gas price, 10^{-5} USD	457	0.45	0.76	0.11	0.18	0.42
marginal gas price, 10^{-5} BoC	457	0.56	0.95	0.14	0.23	0.52
lowest 5 pctile gas price, 10^{-8} ETH	457	0.89	0.96	0.46	0.58	0.80
lowest 5 pctile gas price, 10^{-5} USD	457	0.49	0.84	0.12	0.20	0.47
lowest 5 pctile gas price, 10^{-5} BoC	457	0.62	1.05	0.15	0.25	0.59
median gas price, 10^{-8} ETH	457	1.67	1.39	0.98	1.24	1.58
median gas price, 10^{-5} USD	457	0.94	1.45	0.23	0.44	1.03
median gas price, 10^{-5} BoC	457	1.19	1.81	0.30	0.54	1.32
gas supply, 10^{10} gas/day	457	4.60	0.23	4.52	4.69	4.75
blockchain utilization	457	0.82	0.09	0.76	0.84	0.90
regular transactions share	457	0.52	0.08	0.47	0.51	0.55
ETH price in USD, 10^3 USD	457	0.46	0.29	0.21	0.44	0.67
ETH price (BoC), 10^3 BoC	457	0.58	0.36	0.27	0.55	0.84
	Peak period, Nov. 1, 2017 - Jan 31, 2018					
	# of obs.	Mean	Std	25^{th} pctile	50^{th} pctile	75^{th} pctile
marginal gas price, 10^{-8} ETH	92	1.38	1.20	0.50	0.77	2.09
marginal gas price, 10^{-5} USD	92	1.13	1.32	0.20	0.65	1.45
marginal gas price, 10^{-5} BoC	92	1.42	1.66	0.25	0.80	1.83
lowest 5 pctile gas price, 10^{-8} ETH	92	1.55	1.32	0.56	0.90	2.32
lowest 5 pctile gas price, 10^{-5} USD	92	1.28	1.47	0.22	0.77	1.56
lowest 5 pctile gas price, 10^{-5} BoC	92	1.60	1.84	0.28	0.95	1.98
median gas price, 10^{-8} ETH	92	2.98	1.90	1.54	2.24	3.87
median gas price, 10^{-5} USD	92	2.46	2.48	0.58	1.75	2.87
median gas price, 10^{-5} BoC	92	3.08	3.12	0.72	2.10	3.75
gas supply, 10^{10} gas/day	92	4.36	0.28	4.12	4.34	4.62
blockchain utilization	92	0.82	0.15	0.67	0.87	0.94
regular transactions share	92	0.61	0.09	0.53	0.62	0.70
ETH price in USD, 10^3 USD	92	0.69	0.33	0.41	0.69	1.02
ETH price (BoC), 10^3 BoC	92	0.87	0.41	0.50	0.86	1.25

Notes: summary statistics of the blockchain variables aggregated to the daily level. The full sample (peak period) statistics are reported in the upper (lower) part of the Table respectively. BoC denotes "basket of currencies" – see Section 1.4.3 and Appendix A.2.

Table A.2: Alternative marginal gas price definition, lowest 5-th percentile (BoC)

	lowest 5-th percentile, ETH (1)	(2)	lowest 5-th percentile, BoC (3)	(4)
blockchain utilization, B_t	2.239***	0.118	3.336***	0.970*
	(0.619)	(0.305)	(0.794)	(0.493)
utilization > 90%, TR_t	-	32.38***	-	27.16***
	-	(4.062)	-	(4.928)
regular transactions share, R_t	4.549***	2.658***	8.503***	6.952***
	(0.982)	(0.667)	(0.787)	(0.759)
ETH price (BoC), X_t	-0.517***	-0.452***	-	-
	(0.163)	(0.127)	-	-
sample size	457	457	457	457
R-squared	0.313	0.597	0.657	0.761

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable "lowest 5-th percentile" gas price is the natural logarithm of the bottom 5-th percentile of gas prices (in 10^{-8} ETH or 10^{-5} BoC, weighted basket of currencies consisting of USD, EUR, and CNY) in each block, averaged across all blocks created on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization $>$ 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price (BoC)" is the daily average ETH price in terms of a basket of currencies consisting of US Dollars, Euro, and Chinese Yuan, weighted by the country/region specific number active nodes. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

Table A.3: Median gas price – Basket of currencies (BoC)

	median gas price, ETH (1)	(2)	median gas price, BoC (3)	(4)
blockchain utilization, B_t	0.792*	-0.582**	2.198***	0.585
	(0.480)	(0.268)	(0.699)	(0.519)
utilization > 90%, TR_t	-	24.38***	-	18.79***
	-	(3.817)	-	(5.026)
regular transactions share, R_t	3.955***	2.590***	9.353***	8.281***
	(0.752)	(0.587)	(0.734)	(0.835)
ETH price (BoC), X_t	-0.075	-0.028	-	-
	(0.129)	(0.111)	-	-
sample size	457	457	457	457
R-squared	0.371	0.580	0.676	0.727

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable "median gas price" is the natural logarithm of the median gas price (in 10^{-8} ETH or 10^{-5} BoC, weighted basket of currencies consisting of USD, EUR, and CNY) in each block, averaged across all blocks created on date t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price (BoC)" is the daily average ETH price in terms of a basket of currencies consisting of US Dollars, Euro, and Chinese Yuan, weighted by the country/region specific number active nodes. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

Table A.4: Non-linear effect of blockchain utilization – quadratic model

	gas price					
	marginal		lowest 5-th percentile		median	
	ETH	USD	ETH	USD	ETH	USD
blockchain utilization (demeaned)	3.770***	5.273***	3.711***	5.255***	1.758***	3.633***
	(0.772)	(0.855)	(0.768)	(0.856)	(0.651)	(0.802)
blockchain utilization squared	17.86***	23.48***	17.45***	23.09***	11.47**	17.30***
	(4.627)	(4.888)	(4.581)	(4.861)	(3.969)	(4.497)
regular transactions share, R_t	3.828***	7.334***	3.966***	7.652***	3.546***	8.723***
	(0.932)	(0.747)	(0.916)	(0.749)	(0.725)	(0.764)
ETH price in USD, X_t	-0.732***	-	-0.687***	-	-0.110	-
	(0.208)	-	(0.200)	-	(0.160)	-
sample size	457	457	457	457	457	457
R-squared	0.389	0.714	0.404	0.721	0.427	0.712

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable “marginal gas price" is the natural logarithm of the minimum observed gas price (in 10^{-8} ETH or 10^{-5} USD) over all included transactions in a block, averaged across all blocks created on day t. The dependent variable “lowest 5-th percentile" gas price is the natural logarithm of the bottom 5-th percentile of gas prices (in 10^{-8} ETH or 10^{-5} USD) in each block, averaged across all blocks created on day t. The dependent variable “median gas price" is the natural logarithm of the median gas price (in 10^{-8} ETH or 10^{-5} USD) in each block, averaged across all blocks recorded on day t. “Blockchain utilization (demeaned)" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply, centered around its mean to avoid multicollinearity. “Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. “ETH price in USD" is the daily average ETH price in 10^3 USD. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

Table A.5: Gas price in levels instead of log

	marginal gas price, ETH		marginal gas price, USD	
	(1)	(2)	(3)	(4)
blockchain utilization, B_t	3.303***	0.376*	2.073***	0.241
	(0.950)	(0.216)	(0.526)	(0.164)
utilization $>$ 90%, TR_t	-	51.03***	-	37.62***
	-	(7.838)	-	(6.859)
regular transactions share, R_t	4.550***	1.691*	5.266***	3.453***
	(1.507)	(0.997)	(1.356)	(0.676)
ETH price in USD, X_t	-0.605**	-0.486***	-	-
	(0.291)	(0.179)	-	-
sample size	457	457	457	457
R-squared	0.266	0.630	0.469	0.723

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable "marginal gas price" is the minimum observed gas price (in 10^{-8} ETH or 10^{-5} USD) over all included transactions in a block, averaged across all blocks created on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization $>$ 90%" is the product of blockchain utilization, B_t and a binary variable D_t that equals 1 if $B_t > 0.9$ and zero otherwise. D_t is also included in columns (2) and (4) separately. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price in USD" is the daily average ETH price in 10^3 USD. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

Table A.6: Robustness – including gas supply

dependent variable: marginal gas price	OLS regression		IV regression	
	ETH	USD	ETH	USD
blockchain utilization, B_t	0.077	0.768	0.100	0.705
	(0.363)	(0.584)	(0.356)	(0.565)
utilization > 90%, TR_t	33.09***	28.342***	32.58***	27.305***
	(4.219)	(5.106)	(4.510)	(5.934)
regular transactions share, R_t	2.542***	6.671***	2.997**	7.146***
	(0.659)	(0.739)	(0.973)	(1.203)
gas supply, G_t	0.012	0.104	0.018	0.111
	(0.133)	(0.228)	(0.130)	(0.216)
ETH price in USD, X_t	-0.652***	-	-0.744***	-
	(0.159)	-	(0.190)	-
sample size	457	457	456	456
R-squared	0.595	0.760	0.594	0.759

Notes: OLS and IV regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The IV first stage regresses the regular transactions share on its first lag and the total number of new sender accounts together with other exogenous variables. The dependent variable "marginal gas price" is the natural logarithm of the minimum observed gas price (in 10^{-8} ETH or 10^{-5} USD) over all included transactions in a block, averaged across all blocks created on day t. "Lowest 5-th percentile" gas price is the natural logarithm of the bottom 5-th percentile of gas prices in each block, averaged across all blocks created on day t. "Median Gas Price" is the natural logarithm of the median gas price in each block, averaged across all blocks recorded on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. D_t is also included in all columns separately. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. "Gas supply" is the sum of the gas limits of all blocks created in day t. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price in USD" is the daily average ETH price in 10^3 USD. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

Table A.7: Robustness – using lagged ETH price in USD

dep. variable: marginal gas price, ETH	OLS regression	IV regression
blockchain utilization, B_t	0.086	0.114
	(0.307)	(0.303)
utilization > 90%, TR_t	32.89***	32.37***
	(4.041)	(4.367)
regular transactions share, R_t	2.488***	2.913***
	(0.667)	(0.981)
lagged ETH price in USD	-0.630***	-0.716***
	(0.157)	(0.188)
sample size	457	456
R-squared	0.593	0.592

Notes: OLS and IV regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The IV first stage regresses the regular transactions share on its first lag and the total number of new sender accounts together with other exogenous variables. The dependent variable "marginal gas price" is the natural logarithm of the minimum observed gas price (in 10^{-8} ETH) over all included transactions in a block, averaged across all blocks created on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. D_t is also included in all columns separately. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "lagged ETH price in USD" is the first lag of the daily average ETH price in 10^3 USD. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

Table A.8: Peak period (November 2017 - January 2018) – USD

	gas price in USD					
	marginal		lowest 5-th percentile		median	
	(1)	(2)	(3)	(4)	(5)	(6)
blockchain utilization, B_t	4.020***	2.802***	4.188***	3.143***	3.228***	3.372***
	(0.848)	(0.818)	(0.791)	(0.795)	(0.607)	(0.767)
utilization > 90%, TR_t	-	17.837***	-	16.623***	-	10.768**
	-	(3.789)	-	(3.746)	-	(4.708)
regular transactions share, R_t	5.335***	4.188***	5.253***	4.226***	5.173***	4.895***
	(1.547)	(0.951)	(1.472)	(0.934)	(1.316)	(1.102)
sample size	92	92	92	92	92	92
R-squared	0.833	0.898	0.852	0.904	0.832	0.852

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable "marginal gas price" is the natural logarithm of the minimum observed gas price (in 10^{-5} USD) over all included transactions in a block, averaged across all blocks created on day t. "Lowest 5-th percentile" gas price is the natural logarithm of the bottom 5-th percentile of gas prices in each block, averaged across all blocks created on day t. "Median Gas Price" is the natural logarithm of the median gas price in each block, averaged across all blocks recorded on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

Table A.9: Peak period – Basket of currencies (BoC)

	gas price in ETH					
	marginal		lowest 5-th percentile		median	
	(1)	(2)	(3)	(4)	(5)	(6)
blockchain utilization, B_t	2.759**	0.501	2.895***	0.830	1.515**	0.636
	(1.127)	(0.741)	(1.043)	(0.695)	(0.676)	(0.584)
utilization > 90%, TR_t	-	17.38***	-	16.15***	-	9.610***
	-	(3.117)	-	(3.019)	-	(2.759)
regular transactions share, R_t	3.346**	1.124	3.177**	1.139	1.959*	1.038
	(1.439)	(0.928)	(1.356)	(0.860)	(1.032)	(0.771)
ETH price (BoC), X_t	-0.236	0.069	-0.199	0.078	0.279	0.381**
	(0.389)	(0.258)	(0.362)	(0.238)	(0.232)	(0.186)
sample size	92	92	92	92	92	92
R-squared	0.641	0.854	0.683	0.867	0.734	0.820

Notes: OLS regressions with daily-level data including a constant. Newey-West standard errors reported in the parentheses. The dependent variable "marginal gas price" is the natural logarithm of the minimum observed gas price (in 10^{-8} ETH) over all included transactions in a block, averaged across all blocks created on day t. The dependent variable "lowest 5-th percentile" gas price is the natural logarithm of the bottom 5-th percentile of gas prices (in 10^{-8} ETH) in each block, averaged across all blocks created on day t. The dependent variable "median gas price" is the natural logarithm of the median gas price (in 10^{-8} ETH) in each block, averaged across all blocks recorded on day t. "Blockchain utilization" is the ratio between the total gas requirement for all transactions on day t and the total day-t gas supply. "Utilization > 90%" is the product of the blockchain utilization, B_t and a binary variable that equals 1 if $B_t > 0.9$ and zero otherwise. "Regular transactions share" is the fraction of regular transactions in all day-t transactions. We include dummies for the system-wide gas supply changes on December 10-12, 2017, January 3, and January 21, 2019. "ETH price (BoC)" is the daily average ETH price in terms of a basket of currencies consisting of US Dollars, Euro, and Chinese Yuan, weighted by the country/region specific number active nodes. *, **, *** denote 10%, 5%, and 1% significance level, respectively.

A.2 Additional details

Gas supply events
In December 2017, the gas price increased significantly, associated with very high activity in the Ethereum platform. To compensate, the Ethereum protocol began increasing the block gas limit on December 10, 2017 (marked with a vertical dashed line) to allow higher number of transactions to be included in each block, see Figure 1.6. The upper panel of Figure 1.6 shows the daily average block gas limit. During a 3-day period, the block gas limit was increased by 18% boosting the network service rate. A second event affecting gas supply occurred in January 2019 when the Ethereum developers engaged in a planned increase in the cryptographic difficulty with the intention to switch the consensus algorithm from Proof-of-Work, PoW to Proof-of-Stake, PoS.[1] On January 3 and 21, 2019, the Ethereum blockchain protocol increased the cryptographic difficulty twice. The lower panel of Figure 1.6 displays the corresponding daily block count before and after the difficulty increases. Together, these difficulty increases lowered the daily gas supply by about 17%. We control for these system-wide gas supply events in the regression specifications by using dummy variables which equal 1 at the event dates and zero otherwise.

Transactions order within blocks
By design, each block of the Ethereum blockchain has a capacity limit measured in gas, called the block *gas limit*. The sum of the gas requirements of all included transactions in a block cannot exceed the block gas limit. This implies that miners maximize their profits by ordering the submitted transactions by their gas price.[2] We checked all blocks in our data to verify whether the transactions are always sorted in descending gas price order. We confirmed that this is the case in 85% of the blocks in our data, with the rest of the blocks featuring only minor exceptions from descending gas price order. If a user sets a low gas price, her transaction request may not be written to the next block. Transaction requests that fail to be executed join the 'pending transactions' pool and wait to be recorded in later blocks. This means that Ethereum users face a trade-off between waiting costs and transaction fees and are competing to obtain higher priority among other waiting transactions in terms of the *gas price* and not in terms of the total transaction fee.[3]

[1]In a Proof-of-Stake system, there is no costly competition among miners and instead the block creator is chosen by an algorithm based on the user's "stake", or total ETH balance. In order to change the consensus mechanism from PoW to PoS, the blockchain developers algorithmically increase the cryptographic difficulty which, in theory, would make mining less profitable and provide incentives for the introduction of a Proof-of-Stake consensus mechanism. A switch to PoS has not yet occurred in Ethereum as of December 2020, though it remains planned.

[2]The miners need to solve the Proof-of-Work cryptographic problem as quickly as possible to be selected as the successful creator of the next block and collect the associated transaction fees and block reward. The miners use algorithms which sort and select the transactions to include in the current block in descending gas price order.

[3]This process can also be thought as follows: the successful miner sells space in the block that they mined. Users are buyers who submit bids to purchase space in the block. Assuming that there are only N units of space in the current block and that each user needs one unit on average, then only the users submitting the N highest gas price bids will have their transactions recorded in the block.

Construction of basket of currencies and weighted ETH price

In Section 4 we argued that users may be concerned about the size of transaction fees in terms of their local currency, instead of or in addition to the fee's ETH value. To capture the possible effect of the ETH price in terms of conventional currencies on the gas price bidding choice of users, we create a basket of currencies consisting of the US dollar, the Euro, and the Chinese Yuan and define a weighted exchange rate between ETH and this basket. The selection of currencies and their associated weights was determined by the number of active Ethereum nodes (node refers to a computer running special software and being active part of the Ethereum network) during our study period.[4] We then compute the weighted exchange rate between the ETH cryptocurrency and the defined basket of currencies using the geometric average method (Takagi, 1986),

$$\text{ETH price}_t = \left[\frac{E_t^{\text{USD,ETH}}}{E_{base}^{\text{USD,ETH}}}\right]^{\text{W}^{\text{USD}}} \text{x} \left[\frac{E_t^{\text{EUR,ETH}}}{E_{base}^{\text{EUR,ETH}}}\right]^{\text{W}^{\text{EUR}}} \text{x} \left[\frac{E_t^{\text{CNY,ETH}}}{E_{base}^{\text{CNY,ETH}}}\right]^{\text{W}^{\text{CNY}}}$$

Above, ETH price$_t$ denotes the weighted exchange rate between ETH and the basket of currencies at time t, $\text{E}_t^{\text{USD,ETH}}$, $\text{E}_t^{\text{EUR,ETH}}$, and $\text{E}_t^{\text{CNY,ETH}}$ are the prices of ETH in US Dollars, Euro, and Chinese Yuan at time t, respectively; $\text{E}_{base}^{\text{USD,ETH}}$, $\text{E}_{base}^{\text{EUR,ETH}}$, and $\text{E}_{base}^{\text{CNY,ETH}}$ are the prices of ETH in US Dollars, Euro, and Chinese Yuan in the base period, respectively; and the weights W^{USD}, W^{EUR}, and W^{CNY} equal the shares of active nodes running in the US, Euro-zone countries, and China, respectively. We chose August 7, 2015 as the base date, i.e., we set each ratio $\frac{E_t^{\text{c,ETH}}}{E_{base}^{\text{c,ETH}}}$ for $c = USD, EUR, CNY$ equal to 1 at $t = base$.[5]

[4] We use data from Kim et al. (2018) about the number of active nodes on the Ethereum network in 2018. We only consider locations with share of active nodes greater than 10%. This filtering retains the USA (43.2% share), China (12.9%) and the Euro-zone countries (11.7%). Scaling the shares to 100% yields weights of 63.7%, 19%, and 17.3% for the US, China, and the Euro-zone, respectively.

[5] August 7, 2015 is the first date on which ETH was valued positively against conventional currencies according to min-api.cryptocompare.com.

Appendix B

Chapter 2

B.1 Tables

B.1.1 IEL performance

Table B.1: IEL additional performance

draws	E	Q	ED	ρ_B	ρ_S	I	P	BPS	α
135	91.3	10.9	79.9	0.55	0.59	0.239	106.1	0.48	0.18
140	91.5	10.9	81.4	0.55	0.6	0.242	106.1	0.48	0.18
145	91.5	10.9	81.7	0.55	0.6	0.244	106.2	0.48	0.18
150	91.5	10.9	81.9	0.55	0.6	0.242	106.1	0.48	0.18
155	91.5	10.9	82.5	0.55	0.6	0.244	106.2	0.48	0.18
160	91.6	11	83.2	0.55	0.6	0.245	106.2	0.48	0.18
165	91.6	10.9	83.7	0.55	0.6	0.244	106.1	0.48	0.18
170	91.6	11	84.3	0.55	0.6	0.249	106.2	0.48	0.18
175	91.7	11	84.1	0.55	0.6	0.251	106.2	0.48	0.18

Notes: E = efficiency, Q = quantity, ED = draw of last trade, ρ_B = buyers' correlation coefficient, ρ_S = sellers' correlation coefficient, I = inefficiency measurement, P = price, BPS = buyers' profit split, α= price volatility.

B.1.2 ZI performance

Table B.3: ZI additional performance

draws	E	Q	ED	ρ_B	ρ_S	I	P	Π_B	Π_S	α
135	79.2	8.2	111	0.52	0.52	0.03	109.7	223	314	0.16
140	79.9	8.3	115	0.52	0.53	0.03	109.7	225	317	0.16
145	80.8	8.5	119	0.53	0.53	0.04	109.6	229	319	0.16
150	81.6	8.6	123	0.53	0.54	0.04	109.5	231	322	0.16

155	82.3	8.7	127	0.54	0.54	0.04	109.7	232	326	0.16
160	83.1	8.9	130	0.54	0.55	0.05	109.7	234	329	0.16
165	83.8	9	134	0.54	0.55	0.05	109.6	237	331	0.16
170	84.4	9.1	138	0.55	0.56	0.05	109.8	238	335	0.16
175	85	9.2	142	0.55	0.56	0.05	109.7	240	336	0.16
180	85.6	9.3	145	0.56	0.57	0.06	109.7	242	339	0.16
185	86.1	9.4	150	0.56	0.57	0.06	109.7	244	340	0.16
190	86.6	9.5	153	0.56	0.58	0.07	109.7	244	343	0.16
195	87.1	9.6	157	0.57	0.58	0.07	109.7	246	344	0.16
200	87.6	9.6	160	0.57	0.58	0.08	109.7	247	346	0.16
225	89.4	10	177.4	0.58	0.6	0.088	109.7	253	353	0.15
250	90.9	10.3	196.4	0.59	0.61	0.109	109.5	259	357	0.15
275	92	10.6	212.3	0.6	0.62	0.129	109.7	261	363	0.15
300	92.9	10.8	228	0.61	0.63	0.15	109.8	263	366	0.15
325	93.7	10.9	244.9	0.62	0.64	0.167	109.8	265	370	0.15
350	94.4	11.1	257.8	0.62	0.65	0.188	109.8	267	372	0.15
375	94.8	11.2	274.2	0.63	0.65	0.203	109.8	269	373	0.15
400	95.2	11.4	289.1	0.63	0.66	0.224	109.9	270	376	0.15
425	95.6	11.5	302.6	0.63	0.66	0.241	109.9	271	377	0.15
450	95.9	11.6	315.4	0.64	0.67	0.255	109.9	272	379	0.15
475	96.2	11.6	328.4	0.64	0.67	0.273	109.9	273	379	0.15
500	96.4	11.7	341.8	0.64	0.67	0.283	109.9	273	380	0.15
550	96.8	11.8	368	0.65	0.68	0.318	109.9	275	381	0.15
600	97.1	11.9	390.4	0.65	0.69	0.352	109.9	276	382	0.15
650	97.3	12	416.1	0.65	0.69	0.364	109.9	277	383	0.15
700	97.5	12.1	436.4	0.66	0.69	0.393	110	276	385	0.14
750	97.6	12.2	459.5	0.66	0.7	0.421	109.9	278	384	0.14

Notes: E = efficiency, Q = quantity, ED = draw of last trade, ρ_B = buyers' correlation coefficient, ρ_S = sellers' correlation coefficient, I = inefficiency measurement, P = price, Π_B = buyers' profits, Π_S = sellers' profits.

Table B.4: ZI: price dynamics: average price of the first 10 trades

	Order of Trade									
draws	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
20	107.3	107.4	106.2	105.1	105.9	116.5				
25	107.7	108	107.3	105.8	104.8	102.6	92.9			
30	107.8	108.7	108.4	107.5	107	109.1	106.1			
35	108	109	108.6	107.8	107.7	107	108.9	111		
40	108.3	109	109	108.6	108.3	106.1	104.2	105.6		
45	108.5	109.4	109.2	108.8	108.1	108.3	108.7	102.3		
50	108	109.4	109.3	109.1	108.7	108.5	108.7	105.8		
55	108.2	109.4	109.2	109.7	109	108.4	109.3	108.3	110.4	
60	108.1	109.3	109.6	109.6	109	108.9	108.2	108.3	109.9	123.9

65	108.2	109	109.3	109.7	109.4	108.8	108.5	107.7	109.2	108.9
70	107.8	109.3	109.5	109.5	109.2	109.1	108.8	107.8	108.5	113.1
75	107.6	109.1	109.5	109.6	109.5	109.1	109.3	109.6	108.1	110.9
80	108.3	109.2	109.6	109.4	109.9	109.5	109.4	108.9	108.2	107.1
85	107.9	109.2	109.8	109.7	109.7	109.4	109.4	109.3	109.2	108.9
90	108.5	109.2	109.4	109.7	109.7	109.8	109.4	109.3	109.2	109.8
95	108.2	109.2	109.4	109.6	109.8	109.9	109.9	109.3	107.9	108.4
100	108.6	109.5	109.5	109.6	110	109.7	109.5	109.4	109.6	108.3
105	108.2	109.3	109.7	109.9	110	109.6	109.8	109.8	109.2	109.1
110	108.3	109.3	109.6	109.7	109.7	109.7	109.6	109.8	109.5	108.8
115	108.2	109.1	110.3	109.9	109.6	109.7	109.8	109.3	109.4	109.6
120	108.1	109.3	109.9	109.7	109.7	110.1	109.8	109.3	109.5	109.6
125	108.1	109.3	109.5	110.1	109.8	109.6	109.8	109.7	109.8	110
130	108.1	109.3	109.7	110.1	109.7	110.1	109.9	109.8	109.7	109.7
135	108.1	109.7	110	109.7	109.7	109.9	110	109.5	110	109.4
140	108.1	109.5	109.9	110.2	109.8	109.8	110	109.6	110	110
145	107.9	109.1	109.7	109.9	109.6	110	109.7	109.6	109.7	109.9
150	107.3	109.4	109.8	109.7	109.6	109.8	110	109.8	109.5	109.9
155	108.5	109	109.7	109.9	109.8	110	109.6	110.1	109.8	110.1
160	108.4	109	109.6	109.8	109.9	110	109.9	110	109.9	109.9
165	107.9	109.5	109.5	109.7	109.8	109.7	109.9	109.9	110	109.6
170	108.1	109.7	109.6	109.8	109.9	110.1	109.6	110	110	110
175	108	109.2	109.7	109.9	109.9	110.1	109.7	109.9	110	109.8
180	108.1	109.5	109.5	109.8	110.1	110	109.8	109.8	109.5	110
185	107.8	109.4	109.3	110	110.1	109.8	110.1	109.7	109.8	110
190	108	109.2	109.6	109.6	110	109.8	110.1	110	110.1	110.4
195	108	109.2	109.6	109.6	110	109.8	110.1	109.9	109.9	110
200	108.1	109.5	109.5	109.8	109.8	109.8	109.9	109.9	110.1	110

Table B.2: IEL: price dynamics: average price of the first 10 trades

	Order of Trade									
draws	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
20	104.1	103.6	103.8	104.2	104.7	104.5	104.3	97.2		
25	103.4	103.4	104.2	104.5	104.3	104.9	104.6	104.1	106.5	98
30	103.8	103.6	104	104.4	105.2	105.2	105.6	106.3	106.8	110.5
35	104	103.3	103.8	104.5	105	105.4	105.7	106.5	106.7	109.5
40	104.1	104.1	104.1	104.4	104.6	105.4	105.7	106.4	107.7	107.8
45	103.1	103.6	104.1	104.5	105	105.6	106.4	106.7	107.4	107.9
50	104.4	103.7	103.7	104.3	104.9	105.4	106.3	106.7	107.7	108.2
55	103.9	103.6	104.1	104.3	105.1	105.7	105.8	107	107.9	108.5
60	104	103.8	103.9	104.1	104.5	105.5	106.2	107.3	107.8	108.2
65	104.1	103.6	104	104.3	105.1	105.6	106.3	107.2	107.7	108.7
70	103.9	103.6	104.3	104.5	104.8	105.5	106.3	107.1	108	108.9
75	103.8	103.8	104	104.7	105.1	105.6	106.2	107.3	108	108.8
80	104.1	104	104	104.2	105.3	105.7	106.4	107.4	108.2	108.8
85	103.6	104.1	104.3	104.6	105.2	105.7	106.1	107.3	108.4	109.1
90	103.5	103.6	104.1	103.7	104.9	105.4	106.4	107.3	108.4	109.1
95	104.1	103.9	103.8	104.5	104.9	105.8	106.5	107.2	108.2	109
100	103.8	103.9	104.4	104.7	105.2	105.8	106.4	107.3	108.4	109
105	103.7	103.9	103.6	104.5	104.9	105.6	106.7	107.2	108.6	109.1
110	103.6	104	104	104.4	105	105.5	106.2	107.1	108.1	109
115	103.7	103.9	104.2	104.1	104.3	105.6	106.4	107.4	108.2	109.3
120	103.9	104	103.9	104.6	104.9	105.4	106.4	107.2	108.5	109.1
125	103.9	103.5	104	104.5	105.2	105.5	106.4	107.4	108.7	109.1
130	103.5	103.9	104.2	104.3	105.2	105.6	106.3	107.2	108.1	109.2

B.2 Robustness

B.2.1 IEL parameters: J and μ

Table B.5: Robustness: IEL parameters - efficiency (SFU, 40 draws)

J	μ=0.0033	μ=0.033	μ=0.25
50	69.8	70.2	70.0
100	74.0	73.9	73.9
200	75.7	75.9	75.6
300	76.2	76.1	76.4
400	76.9	76.7	76.7
500	76.8	76.7	76.8

Table B.6: Robustness: IEL parameters - price (SFU, 40 draws)

J	μ=0.0033	μ=0.033	μ=0.25
50	105.4	105.6	105.4
100	104.9	104.9	104.7
200	104.3	104.3	104.4
300	104.4	104.3	104.4
400	104.2	104.0	104.3
500	103.9	104.0	104.0

Table B.7: IEL performance: additional $J = 300, \mu = 0.033$

draws	E	Q	ED	ρ_B	ρ_S	I	P	BPS	α
20	49.9	4.7	16	0.37	0.35	0.01	103.5	0.49	0.21
25	58.9	5.8	21	0.42	0.4	0.01	103.5	0.50	0.21
30	65.8	6.7	26	0.45	0.44	0.03	103.7	0.50	0.21
35	71.7	7.5	30	0.47	0.46	0.04	104	0.50	0.2
40	76.2	8.2	35	0.49	0.49	0.06	104.3	0.50	0.2
45	80.1	8.8	39	0.51	0.51	0.08	104.4	0.50	0.2
50	83.2	9.3	43	0.52	0.53	0.11	104.5	0.50	0.2
55	85.6	9.8	47	0.53	0.54	0.13	104.8	0.49	0.19
60	87.6	10.1	51	0.53	0.55	0.17	104.9	0.49	0.19
65	89.2	10.5	55	0.54	0.57	0.2	105.1	0.49	0.19
70	90.6	10.8	58	0.55	0.58	0.23	105.2	0.49	0.19
75	91.7	11	61	0.55	0.58	0.26	105.3	0.49	0.19
80	92.5	11.2	65	0.56	0.59	0.29	105.4	0.49	0.19
85	93.1	11.3	68	0.56	0.6	0.32	105.4	0.49	0.18
90	93.7	11.5	71	0.56	0.61	0.34	105.4	0.50	0.18
95	94.3	11.6	73	0.56	0.61	0.37	105.5	0.49	0.18
100	94.7	11.7	76	0.57	0.62	0.4	105.6	0.49	0.18
105	95	11.8	79	0.57	0.62	0.42	105.7	0.49	0.18
110	95.2	11.9	81	0.57	0.62	0.44	105.7	0.49	0.18
115	95.4	11.9	83	0.57	0.62	0.45	105.6	0.49	0.18
120	95.7	12	85	0.57	0.63	0.46	105.8	0.49	0.18
125	95.8	12	87	0.57	0.63	0.48	105.8	0.49	0.18
130	96	12.1	89	0.58	0.63	0.49	105.8	0.49	0.18

Notes: E = efficiency, Q = quantity, ED = draw of last trade, ρ_B = buyers' correlation coefficient, ρ_S = sellers' correlation coefficient, I = inefficiency measurement, P = price, BPS = buyers' profit split, α= price volatility.

B.2.2 NI parameters: %NI and $S_n, J = 300$

Table B.8: Robustness: IEL+NI search (55 draws) vs MobLab

	S_N=200			S_N=250			S_N=300		
%NI	e	p	d	e	p	d	e	p	d
0.05	85.9	105.1	51.20	86.1	106.2	38.29	85.8	107.1	29.05
0.10	86.6	105.1	51.86	86.4	107.7	23.95	86.2	109.5	10.81
0.15	86.9	105.3	50.06	86.6	108.9	15.09	86.5	112.1	1.43
0.20	87.2	105.1	53.49	87.0	109.9	10.14	86.7	114.2	1.65
0.25	87.6	105.1	55.12	87.4	111.2	6.17	86.9	116.2	8.38
0.30	88.0	105.2	55.92	87.6	112.4	4.42	87.2	118.6	25.05
0.35	88.2	105.4	54.65	88.0	113.4	6.01	87.6	121.4	56.39
0.40	88.6	104.9	63.64	88.4	114.8	10.47	87.8	123.3	84.50
iel only	85.6	104.8	55.19						
zi only	50.8	109.4	1680.93						

Notes: e=efficiency, p=price, d= distance = $(e - e*)^2 + (p - p*)^2$ where (e^*, p^*) are the MobLab values.

Table B.9: Robustness: IEL+NI search (40 draws) vs SFU

	S_N=200			S_N=250			S_N=300		
%NI	e	p	d	e	p	d	e	p	d
0.01	76.5	104.3	61.76	76.3	104.3	57.48	76.2	104.8	54.37
0.02	76.6	104.2	64.22	76.6	104.6	63.29	76.5	105.3	60.15
0.03	76.6	104.2	64.22	76.4	105.0	58.28	76.5	106.0	60.09
0.04	76.6	104.3	63.96	76.6	105.4	62.29	76.4	106.3	58.17
0.05	77.0	104.4	72.93	76.8	105.7	66.73	76.6	106.9	63.50
0.10	77.3	104.7	79.63	77.2	107.3	78.55	76.8	109.4	78.99
0.20	78.3	104.9	106.29	78.1	109.7	114.32	77.7	114.5	158.36
iel only	76.2	104.2	55.65						
zi only	40.3	109.1	1880.00						

Notes: e=efficiency, p=price, d= distance = $(e - e*)^2 + (p - p*)^2$ where (e^*, p^*) are the MobLab values.

B.3 IEL

We investigate market outcomes under a simple evolutionary mechanism of *individual learning*, which reinforces successful and discourages unsuccessful strategies. We call this model of individual behavior, *Individual Evolutionary Learning*, IEL. It is based on an individual (not social) evolutionary process. It is well suited for applications in environments with large strategy spaces (subsets of real line) such as our continuous double auction environment. See Arifovic and Ledyard (2004) for a discussion of the advantages of IEL over other commonly used models of individual learning, such as reinforcement learning (Erev and Roth, 1998) and Experience-Weighted Attraction learning (Camerer and Ho, 1999), in the environments with large strategy spaces.

In our double auction market environment, agents can submit multiple orders during each trading round. The evolution of the orders within a round is modeled by the Individual Evolutionary Learning (IEL) algorithm and involves the following steps:

- specification of a space of strategies (or messages);
- limiting this space to a small remembered set of strategies for every trader;
- choosing one message from the remembered set on the basis of its performance measure;
- evolving the remembered set using experimentation and replication.

B.3.1 Messages

In our continuous double auction environment, a trader is randomly selected without replacement. A selected trader submits her bid or ask, depending on her role. We assume that a message, $\epsilon_{b,r,t}$ ($\epsilon_{s,r,t}$) represents a potential bid (or ask) order price from buyer b (or seller s) at the rth random selection draw in trading session t. In our treatments, we do not allow a violation of the individual rationality (IR) constraint; that is, we require $\epsilon_{b,r,t} \leq$ V$_{b,t,k+1}$ and $\epsilon_{s,r,t} \geq$ C$_{s,t,k}$ where k denotes the number of units a buyer or a seller holds as submitting a new trade order.[1] In all treatments, we require that possible orders belong to the interval $[s_l, s_u]$.[2]

B.3.2 Individual remembered set

Even if there is a continuum of possible messages, every agent will be restricted at every time to choose between a limited number of them. The remembered set of messages (bids) available for submission at time t by buyer b is denoted by B$_{b,r,t}$. The remembered set of messages (asks) available for submission at time t by seller s is denoted by A$_{s,r,t}$. We initialize these sets by randomly selecting J items from $[s_l, V_{b,k+1,0}]$ for the buyers and $[C_{s,k,0}, s_u]$ for the sellers. In our simulations in this paper, $J = 100$.

In each trading period, the remembered set of each agent is updated whenever there is a change in the order book, but the number of messages in the remembered set is fixed and equals to J. Some of the messages in the remembered set might be identical, so that an agent may be choosing from J or less possible alternatives.

[1] When a buyer already has k units, she bids for the next unit she buys; that is, she will consider her valuation for the k+1th unit. When a seller, on the other hand, has k units, she asks for the next unit she sells; that is, she will consider her cost for the kth unit.

[2] These limits are set by the trading platform and are not parameters of any IEL agent.

The remembered set used at r^{th} random selection draw in trading period t is updated after a change is observed in order book by subsequent application of two procedures: experimentation and replication. During the experimentation stage, any message from the old remembered set can be replaced with a small probability by some new message. In such a way for every buyer and seller the intermediate remembered set s are formed. More specifically, each message is removed from the remembered set with a small probability of experimentation, ρ, or remains in the old remembered set with probability 1-ρ. In case a message is removed, it is replaced by a new message drawn from the truncated normal distribution on $[s_l, V_{b,k+1,0}]$ with mean m and standard deviation equal to $\max\{1, 0.1(V_{b,k+1,0} - s_l)\}$ for buyers and symmetrically for sellers. In our simulations $\rho = 0.033$.

At the replication stage two randomly chosen messages from the just-formed (intermediate) remembered set are compared one with another, and the best of them occupies a place in a new remembered set , $B_{b,r',t}$, for a buyer or $A_{s,r',t}$ for a seller.[3] For every agent such a process is independently repeated J times (with replacement), in order to fill all the places in the new remembered set . The comparison is made according to a performance measure which is defined below. Therefore, during replication, we increase an amount of "successful" messages in the remembered set at the expense of less successful ones.

B.3.3 Calculating the foregone payoffs

How good is a given message? Indeed, only the message which has actually been used last random selection draw delivers a known payoff given by

$$U_b(p) = \begin{cases} V_{b,k+1,t} - p, & \text{if buyer b traded at price p} \\ 0, & \text{if buyer b did not trade} \end{cases} \tag{B.1}$$

$$U_s(p) = \begin{cases} p - C_{s,k,t}, & \text{if seller s traded at price p} \\ 0, & \text{if seller s did not trade} \end{cases} \tag{B.2}$$

An agent who is learning would also like to infer *foregone* payoffs from alternative strategies. To do this, every agent applies a counterfactual analysis. Notice that this is a boundedly rational reasoning, since our agent ignores the analogous learning process of all the other agents.

The calculation of foregone payoff is similar to above calculation; however, the price of transaction is notional and depends on the amount of information which is available to the agent. We use *open book* treatment in all of our treatments. Under the open book treatment each agent uses the full information about all bids, asks, and prices from the previous trades

[3] r' denotes the next random selection draw in which buyer b (or seller s) is selected and the state of the order book has changed.

made in a given trading period. Only the identities of bidders are not known preventing direct access to the behavioral strategies used by others.

Under the open book treatment, an agent knows the state of the order book at every moment of the previous state of a trading session. Assuming that her arrival time does not change, the agent can find a price of a (notional) transaction, $p^*_{\cdot,r,t}(\epsilon)$, for any alternative message $\epsilon_{\cdot}$ and compute his own payoff using regular surplus formula provided above. Thus, the foregone payoffs are given by

$$U_{b,r,t}(\epsilon_{b,r,t} \mid \Im_{b,r,t}) = \begin{cases} V_{b,k+1,t} - p^*_{b,r,t}(\epsilon_{b,r,t} \mid \Im_{b,r,t}), & \text{if order } \epsilon_{b,r,t} \text{ of buyer b transacts} \\ 0, & \text{otherwise} \end{cases} \tag{B.3}$$

$$U_{s,r,t}(\epsilon_{s,r,t} \mid \Im_{s,r,t}) = \begin{cases} p^*_{s,r,t}(\epsilon_{s,r,t} \mid \Im_{s,r,t}) - C_{s,k,t}, & \text{if order } \epsilon_{s,r,t} \text{ of seller s transacts} \\ 0, & \text{otherwise} \end{cases} \tag{B.4}$$

where $\Im_{b,r,t}$ and $\Im_{s,r,t}$ are information sets of buyer b and seller s at random selection draw of r in trading round t.

B.3.4 Selection of a message from the remembered set

In each trading period, traders submit trade orders according to random sequential bidding. That is, a trader is randomly selected from a uniform distribution without replacement and the selected trader submits a message from her remembered set of messages with a *selection probability*. The selection probability is based upon foregone payoffs from the previous change in the state of the market.[4] For example, for buyer b the selection probability of each particular message $\epsilon_{b,r,t}$ from remembered set $\mathrm{B}_{b,r,t}$ is computed as

$$\pi_{b,r,t}(\epsilon_{b,r,t}) = \frac{\mathrm{U}_{b,r,t}(\epsilon_{b,r,t}|\Im_{r,t})}{\sum_{\epsilon \in \mathrm{B}_{b,r,t}} \mathrm{U}_{b,r,t}(\epsilon|\Im_{r,t})}$$

where $\Im_{r,t}$ is an information set available at random selection order r in trading round t. Under individual rationality all messages have non-negative performances. This guarantees that the selection probability is always between 0 and 1.

[4]Within a trading round, market state changes when there is a change in the order book. This could be due to a successful trade, a new best bid (highest bid) or ask (lowest ask), or a cancelled trade order.

B.3.5 What is IEL doing that is different from previous versions used in battle of the sexes, VCM, call markets etc.?

The remembered strategy set varies as the holdings vary. Because IEL buys 1 unit at a time and because the payoff of the next unit depends on how many units IEL currently holds, both the remembered sets and the foregone utilities of their items need to change as the holdings change.

After a buy bid is filled for an IEL buyer, suppose IEL now holds k units of the commodity. All strategies b such that $b \geq V_{k+1}$ are filtered out of A, the set of remembered strategies. Then new strategies are randomly picked from $[s_l, V_{k+1}]$ until A again has J items.

The forgone utilities vary as the holdings and the book change. Foregone utility depends on V_{k+1} and O. V_{k+1} depends on holdings. So the foregone utility of each strategy changes when holdings change.

Timing In past uses of IEL, agents took an action once each round. There was no timing decision to be made. In the CDA, IEL must interact asynchronously in continuous time. This requires a timing decision - when to submit orders. For the CDA simulations, we use the Gode and Sunder methodology of random, sequential bidding. IEL does not change its timing in reaction to changes in the book or to changes in its holdings.

B.4 SFU experiment

B.4.1 Screen shots

Figure B.1: Flex-e-Markets online double auction trading platform (buyer)

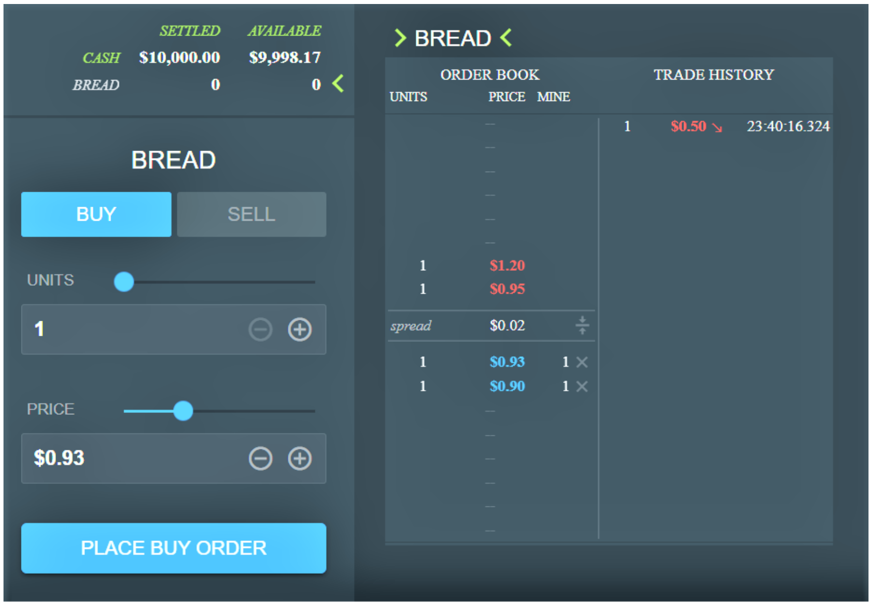

Figure B.2: Flex-e-Markets online double auction trading platform (seller)

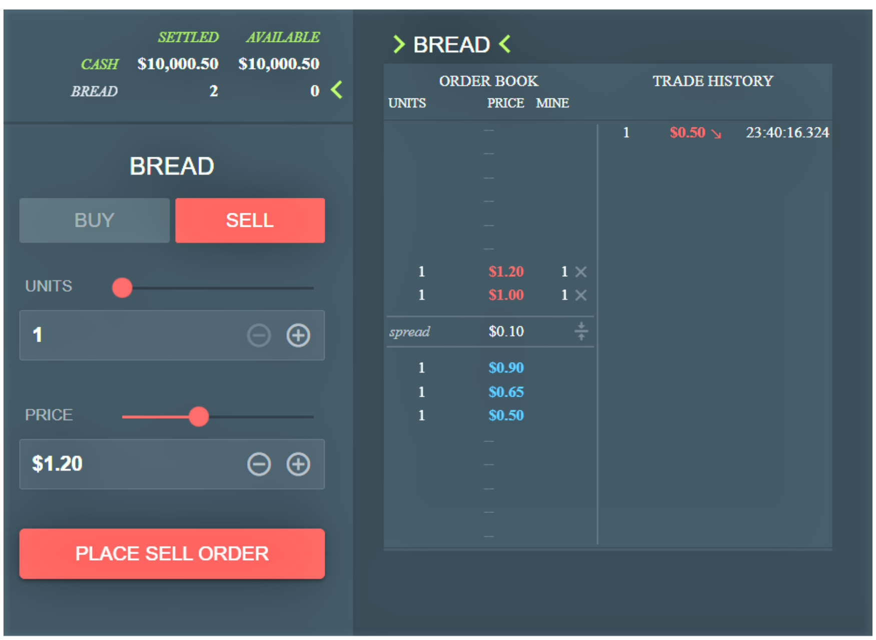

B.4.2 Values and costs

Table B.10: Values and costs for SFU experiements

	Subject #									
Round	1	2	3	4	5	6	7	8	9	10
1	val1+43	val2+43	val3+43	val4+43	val5+43	cost1+43	cost2+43	cost3+43	cost4+43	cost5+43
2	val2-9	val3-9	val4-9	val5-9	val1-9	cost2-9	cost3-9	cost4-9	cost5-9	cost1-9
3	val3+25	val4+25	val5+25	val1+25	val2+25	cost3+25	cost4+25	cost5+25	cost1+25	cost2+25
4	val4+3	val5+3	val1+3	val2+3	val3+3	cost4+3	cost5+3	cost1+3	cost2+3	cost3+3
5	val5-11	val1-11	val2-11	val3-11	val4-11	cost5-11	cost1-11	cost2-11	cost3-11	cost4-11
6	cost1+11	cost2+11	cost3+11	cost4+11	cost5+11	val1+11	val2+11	val3+11	val4+11	val5+11
7	cost2+13	cost3+13	cost4+13	cost5+13	cost1+13	val2+13	val3+13	val4+13	val5+13	val1+13
8	cost3-29	cost4-29	cost5-29	cost1-29	cost2-29	val3-29	val4-29	val5-29	val1-29	val2-29
9	cost4+36	cost5+36	cost1+36	cost2+36	cost3+36	val4+36	val5+36	val1+36	val2+36	val3+36
10	cost5-14	cost1-14	cost2-14	cost3-14	cost4-14	val5-14	val1-14	val2-14	val3-14	val4-14

Notes: val and cost denote buyer valuation and seller cost arrays, respectively, where val1=[150, 135, 120], val2=[147, 132, 117], val3=[144, 129, 114], val4=[141, 126, 111], val5=[138, 123, 108], cost1=[50, 75, 100], cost2=[55, 80, 105], cost3=[60, 85, 110], cost4=[65, 90, 115], cost5=[70, 95, 120].