Measuring misconduct in financial markets

Jonathan R. Karlsen

A thesis submitted in partial fulfilment of the requirements for the degree
of Doctor of Philosophy

Discipline of Finance
Faculty of Business
University of Technology Sydney
July 2021

Statement of originality

This is to certify that to the best of my knowledge, the content of this thesis is my own work. This thesis has not been submitted previously for a higher degree or qualification at any other university or institute of higher learning. I certify that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis, and sources used have been acknowledged. This research is supported by the Australian Government Research Training Program.

Jonathan R. Karlsen

Date: July 20, 2021

To my family

Acknowledgments

This PhD has been an incredible journey and I have in no small amount Prof. Tālis Putniņš and Assoc. Prof. Sean Foley to thank for that. I am grateful for the time, effort, and patience they devoted to instilling in me a high level of academic rigor and an appreciation for good research. It has been a privilege to stand on the shoulders of these giants and I struggle to put into words the admiration and respect I feel for both. I am forever grateful to Tālis and Sean for pushing me to produce what I am proud to call my thesis.

I also thank the UTS finance department for their support during the last four years. The UTS finance faculty is a tightly knit family of excellent researchers and staff and I struggle to think of a better place to write a PhD thesis. In Particular, I would like to thank Prof. David Michayluk and Dr. Ester Félez Viñas for their friendship when being far from home became tough as well as Andrea Greer and Mala Kapachi for their excellent help with everything admin related. I was also fortunate to spend a considerable amount of time with the Financial Conduct Authority in London and I thank them for their resources and expertise. A special thanks goes to the Chief Economist Department and in particular Peter O'Neill, Matteo Aquilina, Carla Ysusi, Andrea Pirrone, and Plamen Ivanov for their valuable insights. I am also grateful to my fellow research students in Australia and abroad, especially Junquin Khan, Man Nguyen, Anirudh Dhawan, Inji Allahverdiyeva, Nihad Aliyev, Christian Neumeier, Steve Freel, and Edward Curran for making Sydney and London feel like home.

I gratefully acknowledge funding from the Rozetta (formerly CMCRC) institute, the Financial Conduct Authority, International Research Scholarship, UTS Finance Discipline Groups, and UTS Business School Editing of Thesis Fund.

I dedicated this thesis to my family. The last couple of years have been some of the best but at times also challenging and my family and friends were always there to pick me up in the bleakest moments. I thank Marius Münier and Marc Bohmann for wonderful distractions including surf trips to remote beaches, motorbike rides in the Australian outback, and most importantly their friendship. Lastly, I thank the reader (whoever you are) for picking up this thesis. I hope you enjoy reading it as much as I enjoyed writing it.

Preface

Some of the work in this thesis is published and was presented as joint work. A version of Chapter 2 and Chapter 3 is published in the *Review of Financial Studies* as an article coauthored with Prof. Tālis Putniņš and Assoc. Prof Sean Foley. I thank the two anonymous referees, the RFS FinTech sponsoring editors (Itay Goldstein, Wei Jiang, and Andrew Karolyi), Michael Weber (discussant), David Easley, Maureen O'Hara, Paolo Tasca, as well as the conference/seminar participants at the RFS FinTech Conference, the RFS FinTech Workshop of Registered Reports, the Behavioral Finance and Capital Markets Conference, the UBS Equity Markets Conference, the University of Technology Sydney, the Australian Federal Police, AUSTRAC, and the Financial Conduct Authority (UK) for their useful comments and suggestions. I also thank Tristan Blakers, Adrian Manning, Luke Anderson, Yaseen Kadir, Evans Gomes, and Joseph Van Buskirk for assistance relating to data. This paper has also been presented at various conferences and seminar series, including the RFS FinTech Conference, the 7th Behavioural Finance and Capital Markets Conference, the FCA Internal Seminar Series, the UBS Equity Markets Conference, and the Vietnam Symposium in Banking and Finance.

Chapters 4 and 5 are part of working papers coauthored with Prof. Tālis Putniņš and Assoc. Prof Sean Foley and were presented at the 12th International Accounting and Finance Doctoral Symposium (IAFDS) and the Rozetta (formerly CMCRC) workshop seminar series. I thank Dr. Ortenca Kume, Prof. Terry Walter, and Prof. Petko Kalev for their useful comments and suggestions.

Table of contents

List of tables

List of figures

Abstract

Fair and efficient financial markets facilitate economic growth by providing funding for firms, investments for individuals, and risk allocation mechanisms. New technologies such as blockchains have the potential to improve these markets by providing more efficient settlements and a range of applications implemented in smart contracts. However, for financial markets and financial technologies to deliver their full potential, they must not only be efficient but also operate with a high level of integrity. Although much research has been devoted to examining financial market efficiency, market integrity—or how fair and free of misconduct markets are—has received far less attention. This thesis helps bridge this gap in knowledge by providing empirical evidence on the prevalence, characteristics, and determinants of illegal activity in both traditional financial markets and cryptocurrencies.

This thesis begins by examining the illegal activity in bitcoin—a cryptocurrency that presents significant challenges for law enforcement given its anonymity, decentralization, and popularity among darknet market criminals. The thesis uses the transaction data from the bitcoin blockchain and a hand-collected sample of individuals who use bitcoin for transactions in illegal goods and services, including in darknet markets and darknet forums. The thesis then estimates the total amount of illegal activity involving bitcoin payments by using two empirical models that rely on different assumptions. The key finding is that approximately one-quarter of all bitcoin users are involved in illegal activity and are responsible for one-half of all transactions in the bitcoin network, equivalent to a total value of $76 billion per year.

The second issue examined is the characteristics of illegal bitcoin users, the determinants of illegal activity, and the topology of their network. Illegal users have characteristics that suggest they use bitcoin as a payment system rather than for speculation or investment; they conduct transactions with many other users, often with the same user repeatedly, hold fewer bitcoins than legal users, and make small and frequent transactions. They also commonly use services that obscure their activity. The proportion of illegal bitcoin usage decreases with its mainstream popularity and the number of alternative (or shadow) coins, which provide technological innovation in improved user and transaction anonymity.[1] The users in the illegal network are very heterogeneous in the number of counterparties they transact with. For example, darknet markets have many more transactional counterparties than darknet market participants.

Third, this thesis examines criminal activity in traditional financial markets. The thesis develops measures of market integrity based on the estimated frequency of insider trading and market manipulation, validating the measures using a hand-collected sample of prosecuted insider trading and closing price manipulation cases. The measures of insider trading are based on abnormal returns,

[1] This thesis refers to anonymity in bitcoin as a continuum (i.e. not binary). Bitcoin does not offer "perfect anonymity" but is more anonymous than, for example, Visa and less anonymous than some other cryptocurrencies, for example, Monero or Z-cash.

abnormal volumes, and abnormal order imbalances in the target companies ahead of announced merger and acquisition (M&A) events. The measures of closing price manipulation are based on abnormal day-end order imbalances, day-end stock price volatility, and overnight return reversals. This thesis combines the insider trading and market manipulation proxies to form a market integrity index that can be used to track market integrity over time and make comparisons between countries. The results for the US suggest that higher regulatory resourcing and whistleblower programs tend to increase market integrity.

Finally, the thesis uses the market integrity index to compare the integrity of financial markets around the world and test the determinants of market integrity. Developed countries exhibit high levels of integrity because of their resourceful regulatory bodies, low corruption levels, and rule-abiding societies. Large, liquid stocks are vulnerable to insider trading because insiders can hide their trades in large order flows, while small, illiquid stocks attract manipulation because their closing prices are easier to move. However, regulatory interventions such as whistleblower schemes, enforcement, increased penalties, and cooperation across jurisdictions can effectively deter misconduct. Market designs such as fragmentation of trading across competing trading venues, dark trading restrictions, and colocation services for algorithmic and high-frequency traders tend to improve market integrity. These market design features are especially useful in jurisdictions with lower regulatory resources.

This thesis has implications for law enforcement, regulators, and legislators around the world. Financial markets commonly face integrity issues when there are lucrative opportunities for criminal activity. Government institutions around the world can use the tools developed in this thesis to conduct market surveillance, detect and discourage criminal activity, and improve market integrity in traditional financial markets and cryptocurrencies.

Chapter 1

Introduction

1.1 Background and motivation

Together with market efficiency, market integrity (or fair markets) is often quoted as one of the main objectives among regulators. The largest regulatory bodies who safeguard more than 50% of the world's listed domestic equity all have either "integrity" or "fair" in their mission statements:[2]

Maintaining fair, orderly and efficient markets

(U.S. Securities and Exchange Commission, SEC)

We protect and enhance the integrity of the UK financial system

(Financial Conduct Authority, FCA)

ASIC's vision is for a fair, strong and efficient financial system for all Australians

(Australian Securities and Investments Commission, ASIC)

Protect investors, foster fair and efficient markets, and contribute to the stability of the financial system

(Ontario Securities Commission, OSC)

But what do regulators really mean by "market integrity"? Considering the regulatory attention it receives, it is surprising how ambiguously market integrity is defined (see, e.g., Austin, 2016). Market integrity is also much less studied than market efficiency, perhaps because it is so difficult to measure. The academic literature often defines market integrity as markets free of insider trading and market manipulation.[3] Similarly, regulators also often quote insider trading and market manipulation as the main culprits of poor market integrity.[4] Both types of misconduct can negatively affect the function of financial markets.[5]

The ways that market integrity can be harmed are abundant. A study[6] by FMSB (FICC Markets Standards Board) on a 200-year dataset of misconduct, called "a history of human greed" by chairman Mark Yallop,[7] finds 250 different types of financial market misconduct over the past 200 years, but "they could have easily gone to 4,000," he adds. The number of ways that markets can be manipulated

[2] See databank.worldbank.org for data on listed market equity.

[3] See: Comerton-Forde and Rydge, 2006; Bhattacharya, Daouk, Jorgenson, Kehr, 2000.

[4] See, e.g., the December 2015 markets article from the Australian regulator (ASIC): "Market integrity matters! You can play an important role in keeping our markets clean."

[5] Insider trading can for example increase the cost of capital (Bhattacharya and Daouk, 2002) and market manipulation may cause inaccurate prices (Comerton-Forde and Putniņš, 2011).

[6] See FMSB (FICC Markets Standards Board) 2017 annual report.

[7] Source: businessinsider.com.

is perhaps put best by the US court in a now famous case of market manipulation: "The methods and techniques of manipulation are limited only by the ingenuity of man."[8] Despite the high focus on market integrity and the detrimental effects it can have on financial markets, there are still no good measures of it.

New financial instruments such as cryptocurrencies have provided criminals with new opportunities and given rise to new avenues of criminal financing and misconduct. Illegal activities that were observed only in financial markets such as "pump-and-dump" schemes, where manipulators intentionally inflate a price using misleading information (see, e.g., Dhawan and Putniņš, 2020) or "wash trading," where manipulators clear their own order book to create a false indication of trading activity, have emerged in cryptocurrencies (Aloosh and Li, 2019). Cryptocurrencies create a lucrative opportunity for manipulation because of their inherent anonymous properties and their low levels of regulation. Unlike traditional financial markets, where the identities of traders are available to law enforcement agencies, cryptocurrencies obscure identities with alphanumerical account IDs (or so-called "addresses"). This level of anonymity for individual manipulators allows them to openly coordinate with each other regarding the time and date of price manipulation in a cryptocurrency without fear of detection. Illegal activity in cryptocurrencies is not only constrained to the activity seen in traditional financial markets but has created avenues for black markets which previously relied on cash as a secure means of transaction. These activities are harmful to the individuals involved and overshadow the benefits of cryptocurrencies and blockchain technology.

Compromised market integrity is costly—cryptocurrencies facilitate illegal drugs and weapons trade, human trafficking, and even assassinations, all of which bear substantial costs to human lives. In traditional financial markets, the resources that go into detection, evidence collection, and prosecution of misconduct are astounding; the US SEC regulatory budget alone is almost $2 billion per year and increased by more than 15% (around $160 million) after the Dodd–Frank act of 2010 to facilitate fair and efficient markets.[9] In 2003, the Bush administration increased the budget by a record 40% "to protect investors, root out fraud, and instill corporate social responsibility" in response to criticism on the agency's failure to respond to corporate corruption scandals.[10]

The threat posed by cryptocurrencies as a payment system for darknet crimes is also given serious consideration by federal enforcement agencies, who collaborate globally to close darknet sites, make arrests, and seize bitcoins and other cryptocurrencies associated with illegal activity. The illegal environment responds to this unwanted attention by developing cryptocurrencies such as Monaro, ZCash, or Dash (previously DarkCoin) with increasingly advanced privacy features. Some of these features, for example, include creating a new address for every transaction ("stealth addresses"), having

[8] See Cargill, Incorporated v. Hardin (1971).

[9] See sec.gov for the SEC regulatory budgets from 1995 to 2020.

[10] See the *New York Times* article "Bush Proposes Big Increase In S.E.C. Budget" from February 2003.

a groups of unrelated networks members "sign off" on transactions ("ring signatures"), or encrypting the addresses and amounts involved in the transaction ("zero-knowledge proofs").

Despite the significant interest in market integrity, many aspects are not yet well understood—how to measure it, its level around the world, and its determinants. This thesis seeks to enhance our understanding by providing evidence on the level and determinants of market integrity in traditional financial markets and cryptocurrencies.

1.1.1 What is market integrity?

Insider trading and market manipulation are two types of market misconduct that harm market integrity. While market manipulation is the act of moving a security's price away from its fundamental value, insider trading refers to the illegal buying or selling of a security based on nonpublic, material information.

For there to be insider trading, there must be nonpublic information to trade on. Common sources of material nonpublic information are M&A announcements that cause price movements when disclosed to the public; any individual in possession of the information can profit by trading on it before it becomes public. Perhaps the most famous example of insider trading is the case of Ivan Boesky, who after being caught by the US Federal Reserve for buying 5% (almost 2 million shares) of the food producer "Carnation" in a Nestle takeover, cooperated with the Federal Reserve System (FED) to build cases against other insider traders. The case ended an era where insider trading legislation existed yet was rarely enforced.[11]

Figure 1.1 illustrates a typical example of insider trading from a prosecution case.[12] The case involves Gary D. Force, an individual who buys shares in DSC Communications Corp. before the news about its acquisition became public. On June 4, 1998, Force instructs his broker, Chad Connor, to buy 50 thousand shares in DSC valued at $18.44. Two days later, on June 6, DSC announces that it will be acquired, and Chad sells half of Force's investment (25 thousand shares) for $28.80 and the other half for $29 the following day. The investment yields a return of more than $0.5 million (or 57%). The price impact of insider trading is noticeable in the example; Force's purchase of $1 million creates a significantly positive return on day -2 and a red flag for the authorities.

[11] For more information on the Carnation case, see Chakravarty and McConnell (1999).

[12] US District Court case file 05 CV 5411.

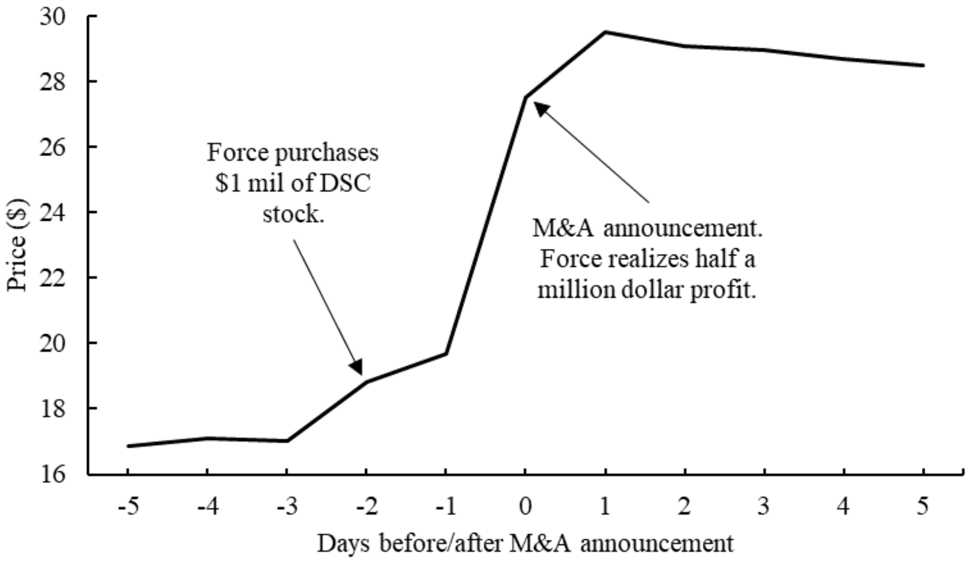

Figure 1.1
An example of insider trading
This figure illustrates the closing price for shares of DSC Communications Corp. ("DSC") in a ten-day interval around the M&A announcement on June 4, 1998. The text indicates when Force purchased shares in the company and when the M&A is made public.

Another form of market misconduct is closing price manipulation. As opposed to insider trading, closing price manipulation is the act of artificially moving the closing price away from its natural level and this does not require any knowledge of future price movements. To achieve this, manipulators often submit aggressive buy or sell orders until the price reaches their desired level immediately prior to market close. Figure 1.2 below illustrates closing price manipulation from a prosecution case.[13]

On August 20, 1999, Moises Saba Masri ("Saba") instructed his broker, Albert Meyer Sutton ("Sutton"), to manipulate the closing price of TV Azteca above $5 to avoid losses in a put options portfolio purchased earlier. Through a succession of seven buy orders in the last minutes of the day's trading, Sutton managed to set the closing price above $5. Sutton's seven buy orders made up 75% of all buy side activity on the day.

[13] US District Court, S.E.C. v. Masri, 523 F.Supp.2d 361 (2007).

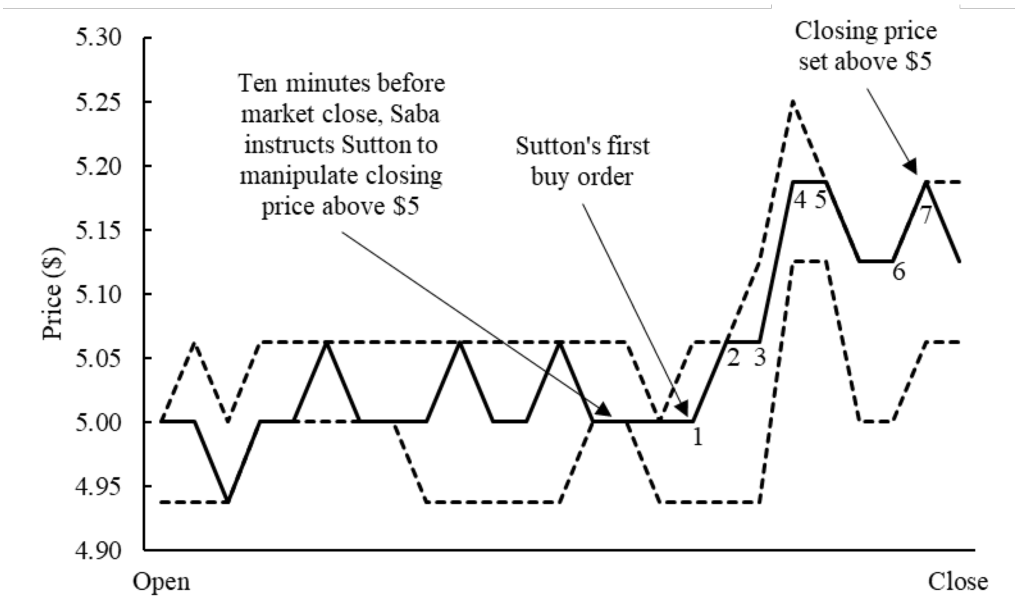

Figure 1.2
An example of market manipulation
This figure illustrates the quoted bid (bottom dashed line), the closing price (middle solid line), and the quoted ask price (top dashed line) for shares of TV Azteca ("TZA") on August 20, 1999. The quoted text indicates when Saba called Sutton and instructed him to manipulate the stock price and when Sutton submits his first and last buy order. Numbers one through seven below the black line indicate Sutton's seven buy orders.

The two examples of market misconduct negatively affect the function of financial markets. Well-functioning financial markets help grow economies by facilitating funding for companies and investments for individuals. Novel financial technologies such as cryptocurrencies represent a first step in how blockchain technology can improve financial systems. However, neither can live up to their full potential if compromised by low levels of market integrity. The focus of this thesis is on market integrity using insider trading in M&A announcements, closing price manipulation, and illegal black market activity in cryptocurrencies to measure it. These types of misconduct are suitable because of their considerably detrimental effects on market integrity.

There are also other advantages to using M&A events; they are very frequent, allowing for a large sample, despite meticulous filtering, and they usually create a positive stock price return. Hence, insider trading is constrained to buying the stock, which leads to a lower measurement error. M&A events are also (as opposed to earning announcements) unscheduled and, therefore, are particularly interesting when measuring insider trading. Their unscheduled nature means that they attract higher concentrations of insider trading because the time and date of the event is nonpublic, thus only known by insiders.

The type of manipulation studied in this thesis, closing price manipulation, is especially harmful to market integrity because it results in unrepresentative closing prices. Closing prices are frequently used to measure stock indices, mutual fund NAVs, derivative instruments, broker performance, and incentive devices such as stock options, which use them to ascertain when directors meet target goals. Thus, closing price manipulation has spillover effects to other financial instruments that use the price as an important reference when computing their value.

Finally, illegal activity in cryptocurrencies has facilitated human trafficking, illegal drug and weapons trade, and even assassinations. Cryptocurrencies have given black market criminals a more efficient way to transact privately through a digital means of payment and online black markets have emerged. These dark web markets facilitate the exchange of illegal goods and services for bitcoin, with activity migrating from the less-efficient cash-based black markets. Figure 1.3 below illustrates an example of illegal activity in bitcoin.

On March 27, 2013, Ross Ulbricht contacts a purported Hells Angels member via encrypted messages on the dark web to arrange an assassination. In the first conversation (Panel A), he provides details about the soon-to-be victim, including his name (Blake Krakoff), age (34), address (White Rock Beach), and details about his family. The parties agree on a price of 1670 BTC for the assassination and Ross Ulbricht sends the agreed amount in a bitcoin transaction on April 1, 2013 (Panel B). The Hells Angels member confirms that the assassination was successful on the evening of April 1, 2013 (Panel C).

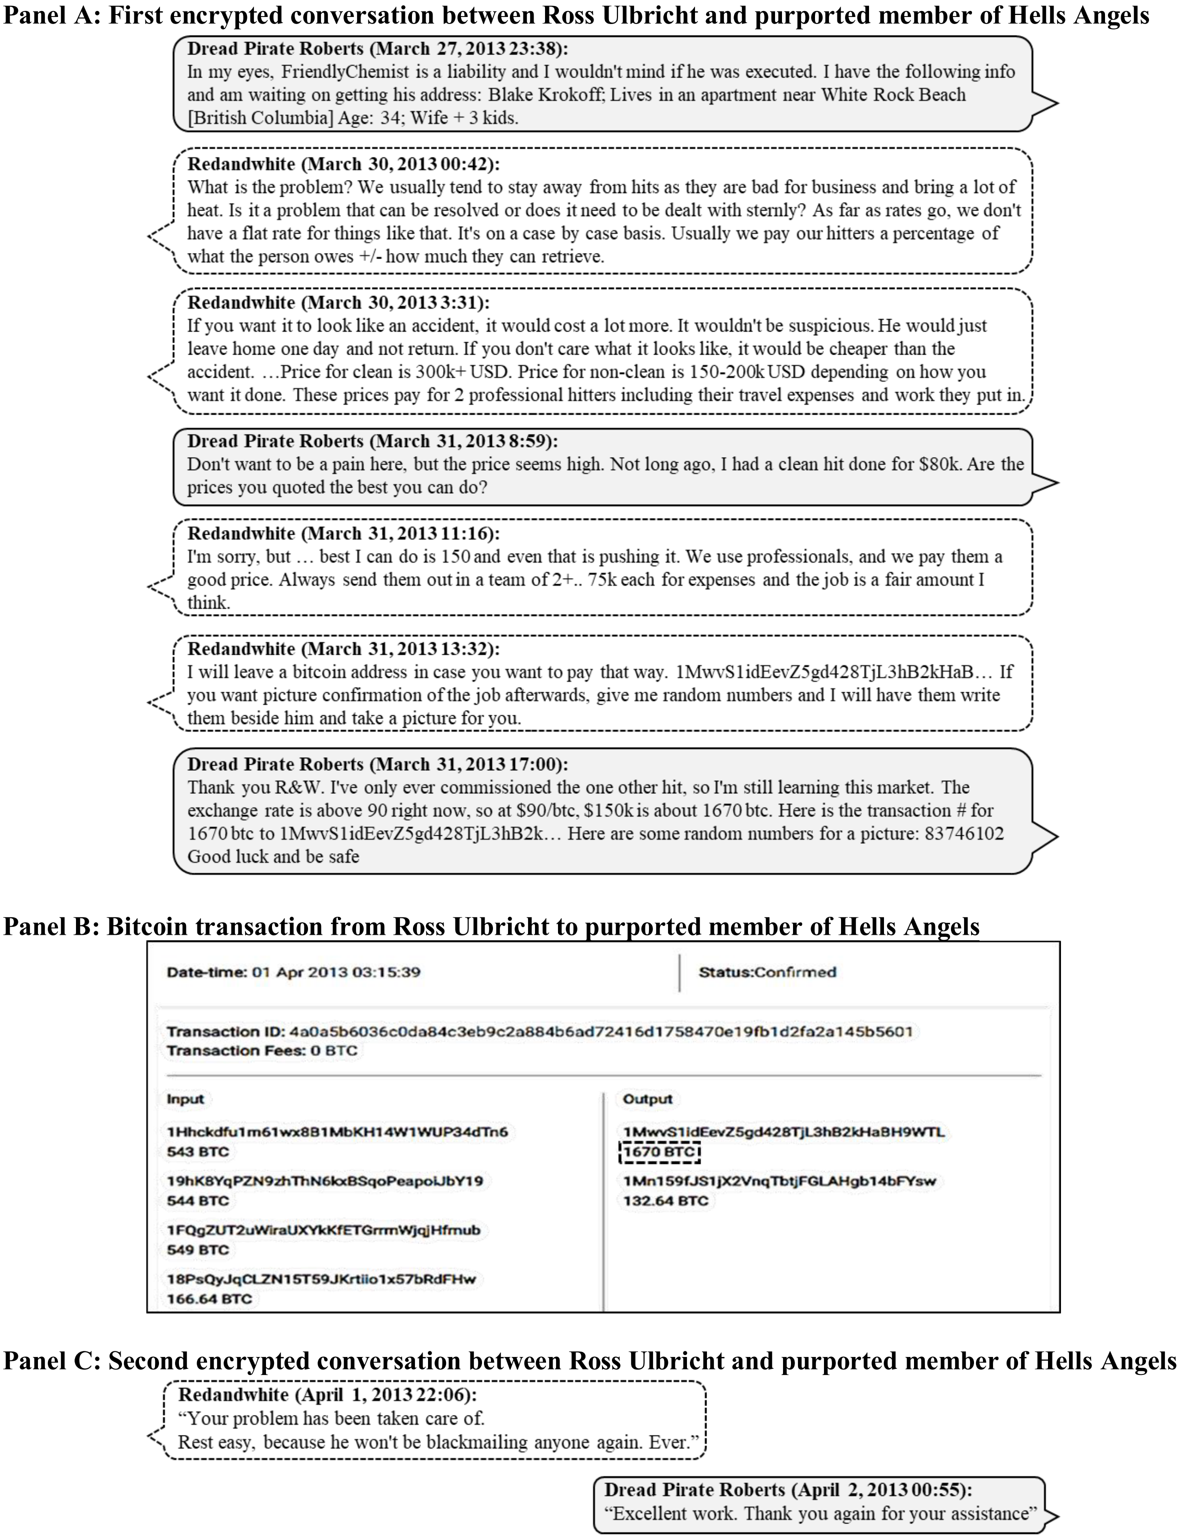

Panel A: First encrypted conversation between Ross Ulbricht and purported member of Hells Angels

Dread Pirate Roberts (March 27, 2013 23:38):
In my eyes, FriendlyChemist is a liability and I wouldn't mind if he was executed. I have the following info and am waiting on getting his address: Blake Krokoff; Lives in an apartment near White Rock Beach [British Columbia] Age: 34; Wife + 3 kids.

Redandwhite (March 30, 2013 00:42):
What is the problem? We usually tend to stay away from hits as they are bad for business and bring a lot of heat. Is it a problem that can be resolved or does it need to be dealt with sternly? As far as rates go, we don't have a flat rate for things like that. It's on a case by case basis. Usually we pay our hitters a percentage of what the person owes +/- how much they can retrieve.

Redandwhite (March 30, 2013 3:31):
If you want it to look like an accident, it would cost a lot more. It wouldn't be suspicious. He would just leave home one day and not return. If you don't care what it looks like, it would be cheaper than the accident. …Price for clean is 300k+ USD. Price for non-clean is 150-200k USD depending on how you want it done. These prices pay for 2 professional hitters including their travel expenses and work they put in.

Dread Pirate Roberts (March 31, 2013 8:59):
Don't want to be a pain here, but the price seems high. Not long ago, I had a clean hit done for $80k. Are the prices you quoted the best you can do?

Redandwhite (March 31, 2013 11:16):
I'm sorry, but … best I can do is 150 and even that is pushing it. We use professionals, and we pay them a good price. Always send them out in a team of 2+.. 75k each for expenses and the job is a fair amount I think.

Redandwhite (March 31, 2013 13:32):
I will leave a bitcoin address in case you want to pay that way. 1MwvS1idEevZ5gd428TjL3hB2kHaB… If you want picture confirmation of the job afterwards, give me random numbers and I will have them write them beside him and take a picture for you.

Dread Pirate Roberts (March 31, 2013 17:00):
Thank you R&W. I've only ever commissioned the one other hit, so I'm still learning this market. The exchange rate is above 90 right now, so at $90/btc, $150k is about 1670 btc. Here is the transaction # for 1670 btc to 1MwvS1idEevZ5gd428TjL3hB2k… Here are some random numbers for a picture: 83746102 Good luck and be safe

Panel B: Bitcoin transaction from Ross Ulbricht to purported member of Hells Angels

Date-time: 01 Apr 2013 03:15:39 | **Status:**Confirmed

Transaction ID: 4a0a5b6036c0da84c3eb9c2a884b6ad72416d1758470e19fb1d2fa2a145b5601
Transaction Fees: 0 BTC

Input	Output
1Hhckdfu1m61wx8B1MbKH14W1WUP34dTn6 543 BTC	1MwvS1idEevZ5gd428TjL3hB2kHaBH9WTL 1670 BTC
19hK8YqPZN9zhThN6kxBSqoPeapoiJbY19 544 BTC	1Mn159fJS1jX2VnqTbtjFGLAHgb14bFYsw 132.64 BTC
1FQgZUT2uWiraUXYkKfETGrmWjqjHfmub 549 BTC	
18PsQyJqCLZN15T59JKrtiio1x57bRdFHw 166.64 BTC	

Panel C: Second encrypted conversation between Ross Ulbricht and purported member of Hells Angels

Redandwhite (April 1, 2013 22:06):
"Your problem has been taken care of.
Rest easy, because he won't be blackmailing anyone again. Ever."

Dread Pirate Roberts (April 2, 2013 00:55):
"Excellent work. Thank you again for your assistance"

Figure 1.3
An example of illegal activity in bitcoin
This figure illustrates an alleged contract killing using bitcoin as a means of payment. Panel A shows an encrypted conversation on the dark web where Ross Ulbricht (pseudonym: *Dread Pirate Roberts*) contracts a Hells Angels member (pseudonym: *Redandwhite*) to assassinate Blake Krokoff (pseudonym: *FriendlyChemist*). Panel B shows a bitcoin transaction from Ross Ulbricht to the Hells Angels member for the arranged fee (1670 BTC) and Panel C shows a second encrypted conversation where the Hells Angels member confirms the assassination. The encrypted conversation in Panels A and B is an excerpt from the full conversation found in *Wired* article "Read the Transcript of Silk Road's Boss Ordering 5 Assassinations" from February 2, 2015.

1.1.2 Why market integrity matters

Market integrity is important for the functioning of financial markets. Insider trading and market manipulation both discourages investor participation (who prefer to trade in cleaner markets), thereby increasing transaction costs and decreasing liquidity. This increases the cost of capital and discourages firms from listing their stocks. Liquidity is also important for price discovery, and inaccurate prices lead to suboptimal resource allocation and wealth redistribution, which lowers market efficiency and creates deadweight economic loss (Pirrong, 1995). Market integrity is especially vulnerable in new technologies such as cryptocurrencies because their design and low regulation pose a lucrative opportunity for undetected criminal activity.

Financial markets and new financial technologies can, when functioning correctly, yield great economic benefits. Financial markets grow economies by efficiently allocating resources, and one of many uses for blockchain technology is the efficient settlement of payment provided by cryptocurrencies. Therefore, understanding market integrity is essential for these markets to meet their full potential.

1.2 Purpose and contribution

The purpose of this thesis is to enhance our understanding of market integrity, which scholars have studied much less than other aspects of markets but is of critical importance to both the investing public and regulators, whose mandate is to protect them. This thesis provides novel evidence of the prevalence and determinants of such misconduct, as well as developing the tools to combat such behavior in both traditional financial markets and cryptocurrencies. Regulators can use the models presented to keep a finger on the "pulse" of market integrity in their markets and use the findings to formulate an appropriate regulatory response. Other market participants, such as investors or firms, may find the model useful when measuring the level of integrity in the markets they invest and raise capital.

The first issue addressed in this thesis is the quantity of illegal activity in cryptocurrencies. Chapter 2 measures the illegal activity in bitcoin using a hand-collected sample of known darknet market participants and transaction data from the blockchain. Approximately one-quarter (25%) of all users and almost half of all bitcoin transactions (46%) are estimated to be illegal. Illegal users hold close to half (49%) of all bitcoin in circulation, and they account for over one-fifth (23%) of the transacted dollar value on the blockchain. Our most recent estimate indicates that there were 27 million illegal bitcoin users, making 37 million transactions valued at $76 billion and holding more than $7 billion of bitcoin in 2017.

The second issue addressed in this thesis is the characteristics of illegal activity in cryptocurrencies. Chapter 3 finds that bitcoin users who engaged in illegal activity tend to make frequent small transactions with few counterparties and hold fewer bitcoin. They are more active immediately following a darknet market seizure or scam and make frequent use of tumbling services and wash trades to conceal their transactions. Illegal bitcoin users become especially active when there are more operational darknet markets, the market value of alternative shadow coins is low, and there is less mainstream interest in bitcoin.

The third issue addressed in this thesis is how to measure market integrity in traditional financial markets. Chapter 4 uses a sample of US prosecutions cases from 1996 to 2016 to develop and validate the measures of insider trading, market manipulation, and market integrity. Cumulative abnormal returns, increased trading volumes, and positive order imbalances indicate insider trading, while manipulation exhibits intraday order imbalances, stock price volatility, and price reversals. Applying the measure to US exchange trading data indicates that market integrity has improved in US markets because of successful legislation, such as whistleblower programs.

The fourth issue addressed in this thesis is quantifying the level of market integrity around the world and understanding its determinants. Using trading data from 25 markets and the integrity measures developed in this thesis, Chapter 5 measures the level of market integrity around the world, its determinants, and how regulation and market design affects it. The chapter finds that developed countries are among the market leaders in market integrity because of their resourceful regulatory bodies, low corruption levels, and generally law-abiding societies. Large, liquid stocks are vulnerable to insider trading while small, illiquid stocks attract market manipulators, but regulatory intervention and market design can deter misconduct and improve market integrity.

1.3 Structure of this thesis

Chapter 2 uses the bitcoin blockchain to quantify the illegal activity in cryptocurrencies. Chapter 3 draws on the methods developed in Chapter 2 to analyze the characteristics of illegal bitcoin activity and the determinants of its detection. Chapter 4 creates, validates, and examines new market integrity metrics in US stock markets. Chapter 5 applies the measures developed in Chapter 4 to characterize integrity in stock markets around the world. Chapter 6 concludes.

Chapter 2

Quantifying illegal activity in bitcoin

2.1 Introduction

Cryptocurrencies have grown rapidly in price, popularity, and mainstream adoption. Over 1,800 cryptocurrencies exist with market capitalization exceeding $300 billion as at July 2018. Bitcoin, the largest cryptocurrency, accounts for around half of the total market capitalization. The numerous online cryptocurrency exchanges and markets have daily dollar volume of around $50 billion.[14] Over 170 "cryptofunds" have emerged (hedge funds that invest solely in cryptocurrencies), attracting around $2.3 billion in assets under management.[15] Recently, bitcoin futures have commenced trading on the CME and CBOE, catering to institutional demand for trading and hedging bitcoin.[16] What was once a fringe asset is quickly maturing.

The rapid growth in cryptocurrencies and the anonymity that they provide users has created considerable regulatory challenges. An application for a $100 million cryptocurrency Exchange Traded Fund (ETF) was rejected by the US Securities and Exchange Commission (SEC) in March 2017 (and several more rejected in 2018) amid concerns including the lack of regulation. The Chinese government banned residents from trading cryptocurrencies and made initial coin offerings (ICOs) illegal in September 2017. Central bank heads, such as the Bank of England's Mark Carney, have publicly expressed concerns about cryptocurrencies. While cryptocurrencies have many potential benefits including faster and more efficient settlement of payments, regulatory concerns center around their use in illegal trade (drugs, hacks and thefts, illegal pornography, even murder-for-hire), potential to fund terrorism, launder money, and avoid capital controls. There is little doubt that by providing a digital and anonymous payment mechanism, cryptocurrencies such as bitcoin have facilitated the growth of online "darknet" marketplaces in which illegal goods and services are traded. The recent FBI seizure of over $4 million worth of bitcoin from one such marketplace, the "Silk Road," provides some idea of the scale of the problem faced by regulators.

This chapter seeks to quantify and characterize the illegal trade facilitated by bitcoin. In doing so, we hope to better understand the nature and scale of the "problem" facing this nascent technology. We develop new methods for identifying illegal activity in bitcoin. These methods can also be used in

[14] SEC Release No. 34-79103, March 10, 2017; and https://coinmarketcap.com.

[15] Source: financial research firm Autonomous Next and cnbc.com.

[16] Bitcoin futures commenced trading on the CME (Chicago Mercantile Exchange) on December 18, 2017 and on the Chicago Board Options Exchange (CBOE) on December 10, 2017. A bitcoin futures contract on CBOE is for one bitcoin, whereas on CBOE it is five bitcoins. At a price of approximately $20,000 per bitcoin at the time the CME bitcoin futures launched, one CME bitcoin futures contract has a notional value of around $100,000.

analyzing many other blockchains. Several recent seizures of bitcoin by law enforcement agencies (including the US FBI's seizure of the "Silk Road" marketplace), combined with the public nature of the blockchain, provide us with a unique laboratory in which to analyze the illegal ecosystem that has evolved in the bitcoin network. Although individual identities are masked by the pseudo-anonymity of a 26-35 character alpha-numeric address, the public nature of the blockchain allows us to link bitcoin transactions to individual "users" (market participants) and then further identify the users that had bitcoin seized by authorities. Bitcoin seizures (combined with a few other sources) provide us with a sample of users known to be involved in illegal activity. This is the starting point for our analysis, from which we apply two different empirical approaches to go from the sample to the estimated population of illegal activity.

Our first approach exploits the trade networks of users known to be involved in illegal activity ("illegal users"). We use the bitcoin blockchain to reconstruct the complete network of transactions between market participants. We then apply a type of network cluster analysis to identify two distinct communities in the data—the legal and illegal communities. Our second approach exploits certain characteristics that distinguish between legal and illegal bitcoin users. We use these characteristics in simultaneous equation models that identify the illegal activity while accounting for the non-randomness of the sample of known illegal users. For example, we measure the extent to which individual bitcoin users take actions to conceal their identity and trading records, which predicts involvement in illegal activity.

We find that illegal activity accounts for a substantial proportion of the users and trading activity in bitcoin. For example, approximately one-quarter of all users (26%) and close to one-half of bitcoin transactions (46%) are associated with illegal activity. Furthermore, approximately one-fifth (23%) of the total dollar value of transactions and approximately one-half of bitcoin holdings (49%) through time are associated with illegal activity using our algorithms. Our estimates suggest that in April 2017, there are an estimated 27 million bitcoin market participants that use bitcoin primarily for illegal purposes. These users annually conduct around 37 million transactions, with a value of around $76 billion, and collectively hold around $7 billion worth of bitcoin.

To give these numbers some context, a report to the US White House Office of National Drug Control Policy estimates that drug users in the United States in 2010 spend in the order of $100 billion annually on illicit drugs.[17] Using different methods, the size of the European market for illegal drugs is estimated to be at least €24 billion per year.[18] While comparisons between such estimates and ours are

[17] The report, prepared by the RAND Corporation, estimates the user of cocaine, crack, heroin, marijuana, and methamphetamine, and is available at (www.rand.org/t/RR534). A significant share of the illegal activity involving bitcoin is likely associated with buying/selling illegal drugs online (e.g., Soska and Christin, 2015), which is what motivates the comparison with the size of the market for illegal drugs.

[18] The estimate is from the European Monitoring Centre for Drugs and Drug Addiction / Europol "EU Drug Markets Report" for the year 2013 (http://www.emcdda.europa.eu/attachements.cfm/att_194336_EN_TD3112366ENC.pdf).

imprecise for a number of reasons and the illegal activity captured by our estimates is broader than just illegal drugs, they do provide a sense that the scale of the illegal activity involving bitcoin is not only meaningful as a proportion of bitcoin activity, but also in absolute dollar terms.

We also uncover that the use of bitcoin in illegal trade varies through time. Since 2016, the *proportion* of bitcoin activity associated with illegal trade has declined, although the *absolute amount* has continued to increase. We attribute the declining share of illegal activity to two main factors. The first is the rapid growth in mainstream and speculative interest in bitcoin, which mechanically decreases the illegal share. For example, we find that the proportion of illegal activity in bitcoin is inversely related to the Google search intensity for the keyword "bitcoin." The second factor is the emergence of alternative "shadow" cryptocurrencies that are more opaque and better at concealing a user's activity (e.g., Dash, Monero, and ZCash). We find that the emergence of such shadow cryptocurrencies is also associated with a decrease in the proportion of illegal activity in bitcoin. Despite the emergence of alternative cryptocurrencies and numerous darknet marketplace seizures by law enforcement agencies, the *amount* of illegal activity involving bitcoin at the end of our sample in April 2017 remains close to its all-time high.

This chapter also makes a methodological contribution. The techniques developed in this chapter can be used in cryptocurrency surveillance in a number of ways, including monitoring trends in illegal activity, its response to regulatory interventions, and how its characteristics change through time. The methods can also be used to identify key bitcoin users (e.g., "hubs" in the illegal trade network) which, when combined with other sources of information, can be linked to specific individuals. The techniques in this chapter can also be used to study other types of activity in bitcoin or other blockchains.

The chapter contributes to a few areas of recent literature. We add to the literature on the economics of cryptocurrencies and applications of blockchain technology to securities markets by showing that one of the major uses of cryptocurrencies as a payment system is in settings where anonymity is valued (e.g., illegal trade).[19] The chapter also contributes to the computer science literature that analyzes the degree of anonymity in bitcoin.[20] We exploit algorithms from this literature to identify individual users in the data, and we add new methods to the literature that go beyond observing individuals, to identification of communities and estimation of populations of users. Finally, the chapter is also related to studies of darknet marketplaces and the online drug trade, including papers from computer science and drug policy.[21] We contribute to this literature by quantifying the amount of illegal activity that involves bitcoin, rather than studying a single market (e.g., Silk Road) or indirect lower-bound measures of darknet activity such as the feedback left by buyers. Empirically, we confirm that

[19] See: Malinova and Park, 2016; Khapko and Zoican, 2017; Yermack, 2017; Huberman et al., 2017; Basu et al., 2019.

[20] See: Meiklejohn et al., 2013; Ron and Shamir, 2013; Androulaki et al., 2013; Tasca et al., 2018.

[21] See: Soska and Christin, 2015; Barratt et al., 2016a; Aldridge and Décary-Hétu, 2016; Van Buskirk et al., 2016.

the estimated population of illegal activity is several times larger than what can be "observed" through studying known darknet marketplaces and their customers.

The next section provides institutional details about bitcoin and the blockchain, darknet marketplaces in which illegal goods and services are bought/sold using bitcoin, and law enforcement efforts to monitor and disrupt illegal online activity. Section 2.3 describes the blockchain data used in this chapter and Chapter 3. Section 2.4 explains three approaches that we use to construct a sample of illegal activity and characterizes that sample. The sample forms the input to our empirical methods in Section 2.5 that quantify the total amount of illegal activity and its trends. Section 2.6 concludes.

2.2 Institutional details

2.2.1 The structure of the bitcoin blockchain

Bitcoin is an international currency, not associated with any country or central bank, backed only by its limited total supply and the willingness of bitcoin users to recognize its value.[22] Bitcoins are "mined" (created) by solving cryptographic puzzles that deterministically increase in difficulty and once solved can be easily verified. Each solution results in a new "block" and provides the miner with the "block reward" (currently 12.5 bitcoins), which incentivizes the miner. The difficulty of the cryptographic puzzles is adjusted after every 2,016 blocks (approximately 14 days) by an amount that ensures the time between blocks remains ten minutes on average.

As well as expanding the supply of bitcoin, each block confirms a collection of recent transactions (transactions since the last block). Each block also contains a reference to the last block, thereby forming a "chain", giving rise to the term "blockchain". The blockchain thus forms a complete and sequential record of all transactions and is publicly available to any participant in the network.

Bitcoins are divisible to the "Satoshi", being one hundred millionth of one bitcoin (currently worth less than two hundredths of a cent). Each bitcoin holding (or parcel) is identified by an address, analogous to the serial number of a banknote. Unlike banknotes, bitcoin does not have to be held in round units (e.g., 5, 10, 50). Due to the revelation of the private key, unless a holding of bitcoin with a given address is exactly spent in a transaction, the "change" from the transaction is returned to a new address forming a new parcel of bitcoin.

A bitcoin "user" (a participant in the network) stores the addresses associated with each parcel of bitcoin that they own in a "wallet". Similar to a conventional cash wallet, a bitcoin wallet balance is the sum of the balances of all the addresses inside the wallet. While individual bitcoin addresses are designed to be anonymous, it is possible to link addresses belonging to the same wallet when more than one address is used to make a purchase.

[22] As of January 2017, over 16 million bitcoins had been mined out of a maximum of 21 million. This maximum limit is built into the protocol (Nakamoto, 2008).

2.2.2 Darknet marketplaces and their microstructure

The "darknet" is a network like the internet, but that can only be accessed through particular communications protocols that provide greater anonymity than the internet. The darknet contains online marketplaces, much like EBay, but with anonymous communications, which also makes these marketplaces less accessible than online stores on the internet. Darknet marketplaces are particularly popular for trading illegal goods and services because the identities of buyers and sellers are concealed. The darknet is estimated to contain approximately 30,000 domains (Lewman, 2016).

To access a darknet marketplace, a user is generally required to establish an account (usually free) at the marketplace to browse vendor products (Martin, 2014a; Van Slobbe, 2016). Similar to the way PayPal propelled EBay, the secure, decentralized, and anonymous nature of cryptocurrencies has played an important role in the success of darknet marketplaces. While bitcoin is the most widespread cryptocurrency used in such marketplaces, other currencies have occasionally been adopted, either due to their popularity (such as *Ethereum*) or improved anonymity (such as *Monero*). Despite the availability of alternate currencies on some marketplaces, the vast majority of transactions on the darknet are still undertaken in bitcoin.[23]

A user that wants to buy goods or services on a darknet marketplace must first acquire cryptocurrency (typically from an online exchange or broker) and then deposit this in an address belonging to the darknet marketplace (often termed a "hot wallet"). These funds are held in "escrow" by the marketplace. Vendor prices on darknet markets are often quoted inclusive of a marketplace fee. The escrow system also assists marketplace administrators in mediating disputes between buyers and sellers and minimizing scams in which money is collected without the intention of ever shipping any goods (Aldridge and Décary-Hétu, 2014; Christin, 2013). Funds are released when the vendor indicates the goods have been sent. In some marketplaces, the funds are held until the buyer indicates that the goods have been received. The escrow function of the darknet marketplaces sometimes leads to "exit scams", whereby a marketplace ceases operations but does not return bitcoin held in escrow. Many such scams have been perpetrated by marketplaces in the last five years, including *Sheep Marketplace* (2013), *Pirate Market* (2014), *Evolution* (2015), and *Nucleus* (2016).

The evolution of dark marketplaces allows sellers of illegal goods and services to reach global audiences (Van Buskirk et al., 2016). This internationalization of illegal trade necessitates more complex methods of communications and logistics to avoid detection. To this end, buyers placing an order with an online seller typically communicate using PGP (Pretty Good Privacy) encryption, which encodes and decodes messages using a pair of public and private keys (Cox, 2016). On some (typically more recent) marketplaces, this functionality is built into the site. Logistically, items are typically delivered by mail and the process by which this occurs has been widely documented (Christin, 2013;

[23] A recent estimate from a darknet marketplace operator identified bitcoin as accounting for 98% of transactions: https://www.wired.com/2017/01/monero-drug-dealers-cryptocurrency-choice-fire/.

Van Hout and Bingham, 2013; Lavorgna, 2016; Van Slobbe, 2016). Many methods are used to minimize the chance of such deliveries being intercepted by law enforcement, including professional logos, vacuum sealed bags, posting small quantities of product, and including a (fake) return address (Christin, 2013; Basu, 2014; Tzanetakis et al., 2016). Customers are advised by marketplaces to avoid using their real name or address to minimize the risk of being caught by law enforcement agencies (Martin, 2014b).

After receiving their goods, buyers are encouraged to leave feedback about the seller, commenting on the arrival (or otherwise) of the goods, their quality, and overall service (Van Slobbe, 2016). Such feedback is paramount for developing a reputation in a marketplace that is primarily based on trust between participants, with few ramifications for "scamming" purchasers (Aldridge and Décary-Hétu, 2014; Tzanetakis et al., 2016).

To get a sense of how a buyer navigates a darknet marketplace, Figure 2.1 provides screenshots from one of the first darknet marketplaces, "Silk Road". Panel A provides an example of the "Drugs" page illustrating that a wide variety of illegal drugs, weapons, and forgeries can be purchased using bitcoin. Panel B provides an example of information about individual items and sellers. Clicking on the appropriate headings, one can obtain further information about the items (detailed description, insurance/refund policies, available postage methods and locations, security and encryption, and so on) and about the seller (their rating from buyers, detailed feedback from buyers, history of sales, and so on). Panel C shows the interface for depositing bitcoin to Silk Road's escrow account, how to transfer bitcoins to a given seller, and how to withdraw bitcoins from escrow.

By providing an anonymous, digital method of payment, bitcoin did for darknet marketplaces what PayPal did for EBay—provide a reliable, scalable, and convenient payment mechanism. What was also required was an anonymous way of hosting and accessing those illegal marketplaces. This issue is solved through the use of The Onion Router (TOR), originally developed by the US Navy. By routing the message through several nodes, the TOR network obfuscates the path (and hence the IP address) of a message sent between two clients.

The combination of TOR for covert communications and bitcoin for covert payments has led to the proliferation of darknet marketplaces. The most well-known marketplace was the "Silk Road" started in 2011. Since its shutdown by the FBI in 2013, numerous other marketplaces have sprung up (see Table A1 in Appendix A for a list). Despite frequent shutdowns, seizures and scams, measures of darknet marketplace activity indicate steady growth in the number of market participants and products (Matthews et al., 2017). For example, one of the largest marketplaces in 2017, "AlphaBay", had over 350,000 items available for sale in categories such as drugs, weapons, malware, and illegal pornography.

Panel A: Example of illegal drugs that can be purchased with bitcoin on the Silk Road marketplace

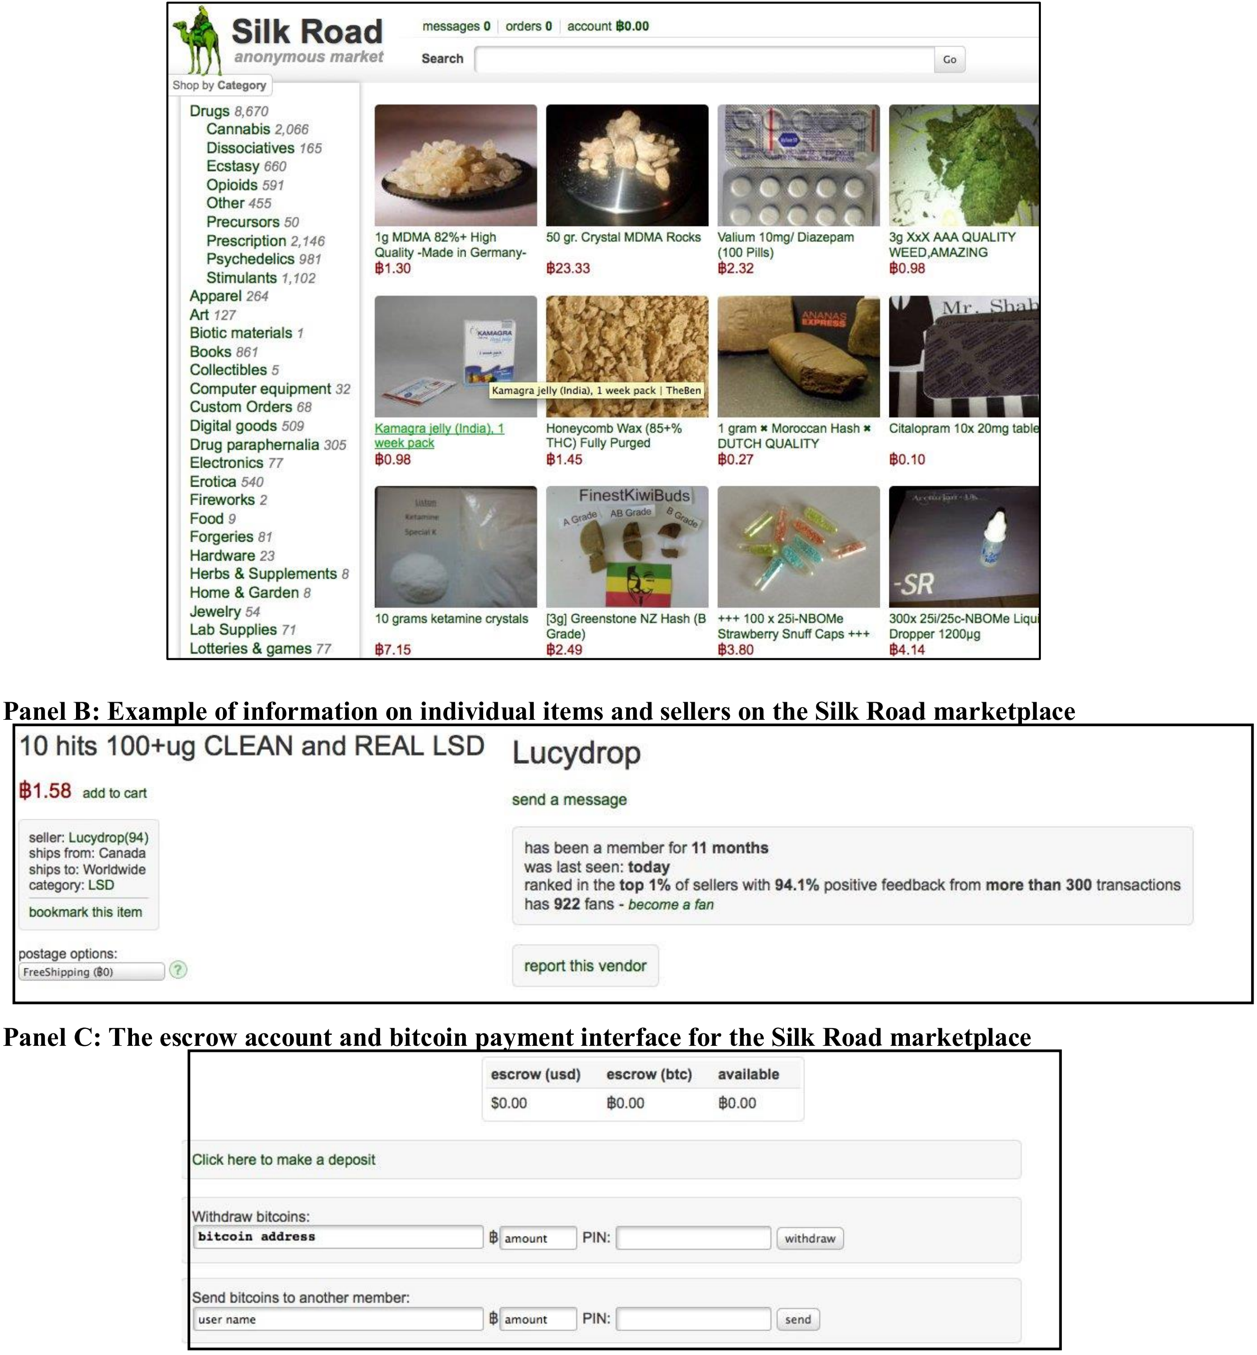

Silk Road
anonymous market

messages 0 | orders 0 | account ฿0.00

Search Go

Shop by Category

Drugs	8,670
Cannabis	2,066
Dissociatives	165
Ecstasy	660
Opioids	591
Other	455
Precursors	50
Prescription	2,146
Psychedelics	981
Stimulants	1,102
Apparel	264
Art	127
Biotic materials	1
Books	861
Collectibles	5
Computer equipment	32
Custom Orders	68
Digital goods	509
Drug paraphernalia	305
Electronics	77
Erotica	540
Fireworks	2
Food	9
Forgeries	81
Hardware	23
Herbs & Supplements	8
Home & Garden	8
Jewelry	54
Lab Supplies	71
Lotteries & games	77

1g MDMA 82%+ High Quality -Made in Germany-
฿1.30

50 gr. Crystal MDMA Rocks
฿23.33

Valium 10mg/ Diazepam (100 Pills)
฿2.32

3g XxX AAA QUALITY WEED,AMAZING
฿0.98

Kamagra jelly (India), 1 week pack | TheBen

Kamagra jelly (India), 1 week pack
฿0.98

Honeycomb Wax (85+% THC) Fully Purged
฿1.45

1 gram ✖ Moroccan Hash ✖ DUTCH QUALITY
฿0.27

Citalopram 10x 20mg table
฿0.10

10 grams ketamine crystals
฿7.15

[3g] Greenstone NZ Hash (B Grade)
฿2.49

+++ 100 x 25i-NBOMe Strawberry Snuff Caps +++
฿3.80

300x 25i/25c-NBOMe Liqui Dropper 1200µg
฿4.14

Panel B: Example of information on individual items and sellers on the Silk Road marketplace

10 hits 100+ug CLEAN and REAL LSD

฿1.58 add to cart

seller: Lucydrop(94)
ships from: Canada
ships to: Worldwide
category: LSD

bookmark this item

postage options:
FreeShipping (฿0)

Lucydrop

send a message

has been a member for **11 months**
was last seen: **today**
ranked in the **top 1%** of sellers with **94.1%** positive feedback from **more than 300** transactions
has **922** fans - *become a fan*

report this vendor

Panel C: The escrow account and bitcoin payment interface for the Silk Road marketplace

escrow (usd)	escrow (btc)	available
$0.00	฿0.00	฿0.00

Click here to make a deposit

Withdraw bitcoins:
bitcoin address ฿ amount PIN: withdraw

Send bitcoins to another member:
user name ฿ amount PIN: send

Figure 2.1
Screenshots from one of the first illegal darknet marketplaces, Silk Road 1
Panel A provides an example of the "Drugs" page from Silk Road. It illustrates the wide variety of illegal goods that can be purchased using bitcoin, including a vast array of illegal drugs, weapons, and forgeries. Panel B provides an example of information about individual items and sellers. Clicking on the appropriate headings, one can obtain further information about the item for sale (detailed product description, insurance/refunds, postage methods and locations, security and encryption, etc.) and about the seller (detailed feedback and ratings from buyers, history of sales, etc.). Panel C shows the interface for depositing bitcoin to Silk Road's escrow account, transferring bitcoins to a given seller, and withdrawing bitcoins from escrow. Screenshot source: www.businessinsider.com.au.

2.2.3 Surveillance and cryptocurrency seizures from darknet marketplaces

Cryptocurrencies have proven effective not only in facilitating illegal trade, but also in the detection of illegal activity due to the public nature of the blockchain. Even though bitcoin has been used extensively in illegal activity, some argue that the blockchain actually makes it easier for law enforcement to detect illegal activity, despite the currency's anonymity. Koshy, Koshy, and McDaniel (2014) show that by monitoring transactions transmitted from computers to the blockchain, they are able to link individual transactions to the IP address of the sender. Meiklejohn et al. (2013) describe how tracing a bitcoin theft on the blockchain to bitcoin exchanges could be used by authorities with subpoena powers to potentially identify perpetrators. Yermack (2017) hypothesizes that the growing popularity of bitcoin will inevitably lead to a growing market for de-anonymizing technologies, leading to increased transparency of the users making transactions on the blockchain. In response to these pressures, supporters of the anonymity provided by cryptocurrencies are actively developing new currencies that challenge law enforcement's detection methods. Such currencies include *Monero*, which hides user's public keys among a group of public keys that contain the same amount (known as "Ring Signatures"), and *ZCash* (launched in 2016), which uses zero-knowledge proofs that hide sender, recipient, and transaction amount (Noether, 2015; Ben-Sasson et al., 2014).

Recently, law enforcement agencies have been successful in seizing bitcoin from a number of darknet marketplaces. For example, the Silk Road marketplace was raided by the FBI on October 2, 2013, seizing bitcoin from customer and supplier escrow accounts (hot wallets) and from the owner/operator, Ross William Ulbricht. After the closure of the Silk Road, law enforcement agencies successfully seized bitcoin from several other illegal sites/individuals (see Table A2 of Appendix A). Numerous darknet sites were raided and shut down in "Operation Onymous"; an international collaboration between US and European law enforcement agencies that targeted illegal darknet sites. Despite the seizures, illegal darknet marketplaces continue to operate, with many new ones created after each seizure.

The seized bitcoin from these operations allows us to identify bitcoin users (customers, suppliers, and marketplace operators) involved in illegal activity. These observations provide a starting point from which to estimate the extent of illegal activity involving bitcoin.

Law enforcement agencies use a number of strategies to detect illegal activity on the darknet, ranging from cyber-surveillance to forensic analysis. Given that detected illegal activity feeds into our identification techniques, it is important to understand law enforcement strategies. Christin (2013) and Kruithof et al. (2016) describe a number of such strategies, including: infiltrating the TOR network to determine individual IP addresses, decoding the financial infrastructure of bitcoin to identify individuals, and using traditional forensic and investigative techniques on seized packages. Law enforcement agencies monitor suspicious packages passing through the postal service. Agencies also order drugs on darknet marketplaces to investigate the return address on the package. For example, an unusual amount of outgoing mail from a large Australian drug dealer led authorities to seize over 24,000 in bitcoin, along with a wide array of drugs and cash. Investigators also sometimes pose as suppliers to gather addresses of customers, thereby revealing their identities. Finally, by conducting major seizures, agencies can create distrust in the online trade of illegal drugs among participants (Van Slobbe, 2016; Christin, 2013). Large-scale initiatives such as "Operation Onymous", in which law enforcement agencies shut down several illegal marketplaces and made 17 arrests across 17 countries, can discourage illegal online activity by increasing the risk of detection (Franklin, Paxson, Perrig, and Savage, 2007).

2.3 Data and descriptive statistics

We extract the complete record of bitcoin transactions from the public bitcoin blockchain, from the first block on January 3, 2009, to the end of April 2017. For each transaction, we collect the unique transaction hash, the transaction amount, the fee, the sender and recipient addresses, the timestamp, and the block number.

2.3.1 Identifying users in transaction-level bitcoin data

The data that make up the bitcoin blockchain reveal "addresses" (identifiers for parcels of bitcoin) but not the "users" (individuals) that control those addresses. A user typically controls several addresses. This one-to-many mapping occurs partly as a result of various activities that users employ to preserve their anonymity and partly due to transaction mechanics (e.g., when a user receives "change" in a transaction, the change is given a new address).[24] We map addresses to individual users with the Union-Find algorithm developed by Cormen, Leiserson, Rivest, and Stein (2001) and Ron and Shamir (2013) and used in several related papers such as Meiklejohn et al. (2013). This algorithm transforms the transaction-level data into user-level data, linking each transaction to the associated users.

The following illustrates how the Union-Find algorithm works. A transaction usually involves several addresses from one user. For example, the payer ("sender") of bitcoin might send bitcoin from multiple addresses and receive change to a new address. Because a user must control the private key of each address from which bitcoin is sent in a given transaction, in the first step of the algorithm all of the sender's addresses in a given transaction are associated with one user. Transitivity is then used to link the addresses of a user across multiple transactions. For example, suppose two separate transactions are observed; one in which bitcoin is sent from addresses A and B and another in which bitcoin is sent from addresses B and C. The first transaction identifies that addresses A and B belong to one user, while the second identifies that B and C belong to one user. By transitivity, all three addresses (A, B, and C) belong to the same user.

[24] For example, individuals can send bitcoin to a "tumbling" service which then returns the bitcoin (minus a fee) to a new address, or by sending bitcoin to oneself using a newly generated address as the recipient of the transaction (Ron and Shamir, 2013).

None of the existing algorithms that cluster bitcoin addresses by user has perfect accuracy.[25] The Union-Find algorithm is the most widely used approach, primarily because the errors it makes (too little clustering of addresses rather than too much clustering) are conservative in most applications (Meiklejohn et al, 2013). The Union-Find algorithm might fail to cluster together two sets of addresses controlled by one user if that user never makes a transaction that uses an address from both sets. In such instances, two or more address clusters might in fact correspond to one user.[26] In contrast, the Union-Find algorithm (unlike other approaches such as those that exploit the change from transactions) is very unlikely to make the opposite and more severe error of incorrectly clustering together sets of addresses that involve *more* than one user. The Union-Find algorithm is a suitable choice in our application because too little clustering (and thus having instances where two or more clusters correspond to one actual user) is unlikely to have severe consequences for our empirical methods, whereas incorrectly joining multiple users into a single cluster would be far more problematic.[27]

The Union-Find algorithm's tendency to join too few addresses together into clusters adds bias to some of the measures in this chapter. First, the sample of known illegal users, which is the starting point for our empirical analysis, will not contain all of the addresses controlled by those users. We therefore start with a smaller sample than would be the case if the clustering algorithm had 100% accuracy. Second, measures of the number of users will be upward biased because in some cases, two or more of the clusters identified by the Union-Find algorithm will in fact be controlled by one real user. Consequently, our estimates of the total number of bitcoin users, the number of illegal users, and the number of legal users are all likely upward biased. For similar reasons, measures such as the number of transactions per user or holdings per user are likely downward biased. This bias is less of an issue when we quantify users as percentages of the total number of users. For example, our estimates of the percentage of users that are involved in illegal activity will be less biased than the absolute number of users. This bias is even less of an issue when we quantify the number of transactions, volume, or holdings of various groups of users because these measures do not rely on knowing the number of users in each group.

2.3.2 Filters and data transformations

Our blockchain dataset consists of 465,093 blocks containing 219.6 million bitcoin transactions (unique transaction hashes). In the raw blockchain data, one transaction can have several recipients. For example, in a single transaction, Alice could send five bitcoins to Bob, two bitcoins to Charlie, and 0.1

[25] For example, Androulaki et al. (2013) examine two approaches using simulations and find that many, but not all, users can be correctly identified by clustering algorithms even when they try to enhance privacy by creating new addresses.

[26] Meiklejohn et al. (2013) empirically find that this error is "not too common" in bitcoin blockchain analysis.

[27] For example, if a single user appears in the data as two or more clusters, all of those clusters could be correctly classified with the user's actual type (illegal or legal), whereas if a legal and illegal user are incorrectly clustered together, there is no way to assign a correct classification to the cluster.

bitcoins to the miner of the block as a transaction fee. We split these raw "compound" transactions up into their components such that each transaction has only one sender and one receiver. In the previous example, Alice's compound transaction would become three separate transactions: one with Bob, one with Charlie, and one being a transaction fee sent to the block miner. Among other things, this allows us to separate transaction fees and block rewards from other transactions. After splitting compound transactions into their components, we have 815.4 million transactions.

In this study, we are primarily interested in quantifying the amount of illegal trade that uses bitcoin. We therefore remove transaction fees and block rewards from the sample to avoid distorting the transaction counts. This step removes 208.3 million transactions. We also remove currency conversion transactions (conversions between bitcoin and fiat currency or other cryptocurrencies), by removing bitcoin exchanges and their 88.4 million transactions. These transactions do not involve trade in the sense of buying or selling goods or services and would therefore inflate our measures of transaction activity.[28] For similar reasons, we remove 71.1 million transactions that reflect the "change" given back to a user in a given transaction. These transactions are akin to paying for a $30 product with a $50 bill and receiving $20 of change back (our processed dataset would record this scenario as one $30 transaction rather than two transactions).[29]

We also exclude transactions that have a value of less than $1 on the day of the transaction. Such transactions reflect negligible transfers of value and are therefore used for purposes such as messages, test transactions, and tips.[30] Failure to exclude these transactions could significantly skew our data, particularly measures of the proportion of transactions.

After applying these filters, we are left with 302.8 million transactions, each having one sender and one receiver. Throughout much of the chapter we consider user-level statistics such as the number of transactions per user. Such measures naturally use double-counted volume as both the sender and receiver sides of each transaction are counted. Using double-counted volume, our transaction count is doubled to 605.7 million bitcoin transactions.

[28] The exchanges and miners are identified via "Wallet Explorer." Wallet Explorer joins transactions into "wallets" (the equivalent of our "users") using a similar procedure to the one described above and then classifies wallets by type either on the basis of (i) having observed an address being advertised as part of a given entity (e.g., a known address from a bitcoin exchange), or (ii) having identified an entity's wallet by sending a small amount of bitcoin to the entity, where that address is linked to the larger wallet of the entity (similar to Meiklejohn et al., 2013). See https://www.walletexplorer.com.

[29] The bitcoin protocol forces a user to spend the entire balance of a bitcoin address when the address is used in a transaction. Therefore, when a user has say 50 bitcoins in address A and wants to send another user say 30 bitcoins, the compound transaction would have two components: one transaction sending 30 bitcoins from address A to the other user and another transaction sending the remaining 20 bitcoins from address A to a new address B that is controlled by the sending user. The latter of the two transactions is the "change" and is removed from our sample.

[30] These 144.7 million small transactions represent 17.8% of transactions, but less than 0.0001% of total bitcoin volume.

2.3.3 Descriptive statistics of user-level variables

Our sample has a total of approximately 106 million bitcoin users, who collectively conduct approximately 606 million transactions, transferring around $1.9 trillion. For each user, we calculate a collection of variables that characterize features of their bitcoin transaction activity (e.g., transaction count, transaction size, transaction frequency, and number of counterparties). We also calculate a range of user-level variables that are more specific indicators of the nature of the activity in which a user is likely to be engaged, such as the number of illegal darknet marketplaces that operate at the time the user transacts, the extent to which the user engages in transactions designed to conceal their activity, and the degree of interest in bitcoin at the time the user transacts (using Google search intensity). The detailed definitions of these variables are reported in Table 2.1.

Table 2.1
Variable definitions
This table defines the variables that we compute for each bitcoin user. The third column, *DCE equation*, specifies whether the variable is used in the first equation of the DCE model (the equation modelling whether the user is involved in illegal activity, *I*), the second equation of the DCE model (the equation modelling whether a user that is involved in illegal activity is "detected", e.g., seized by law enforcement agencies, *D*), both equations (*I* & *D*), or as a control variable (*C*).

Variable	Definition	DCE equation
Transaction Count	The total number of bitcoin transactions involving the given user (where the user is a sender and/or recipient of bitcoin).	C
Transaction Size	Average USD value of the user's transactions. The transaction size is converted from bitcoin to USD using end of day USD/BTC conversion rates. Exchange rates prior to July 18, 2010 are not available and therefore values before this date are converted at the July 18, 2010 exchange rate.	I & D
Transaction Frequency	The number of bitcoin transactions made by the user per month. This is computed as *Transaction Count* divided by *Existence Time*.	I & D
Counterparties	The total number of other users with which the given user has transacted.	C
Holding Value	The average USD value of the user's bitcoin holdings. The average is computed from the holding balances recorded at the end of each of the user's bitcoin transactions. Holding values are converted from bitcoin to USD using end of day USD/BTC conversion rates. Exchange rates prior to July 18, 2010 are not available and therefore values before this date are converted at the July 18, 2010 exchange rate.	C
Concentration	Concentration is a measure of the tendency for a user to transact with one or many counterparties. It ranges from 1 for a highly concentrated user who transacts with only one counterparty, to 0 for a user that has many transactions, each with a different counterparty. Formally, it is computed using an adaptation of a normalized Herfindahl–Hirschman Index: $Concentration = \begin{cases} 1 - \left[\frac{\left[\left(\frac{C}{T}\right) - \left(\frac{1}{T}\right) \right]}{1 - \left(\frac{1}{T}\right)} \right] & \text{if} \quad T > 1 \\ 1 & \text{if} \quad T = 1 \end{cases}$ where T is *Transaction Count* and C is *Counterparties* (the total number of other users with which the given user has transacted).	I & D
Existence Time	Number of months the bitcoin user is active in the bitcoin network. Measured as the number of months from the user's first transaction until the user's last observed transaction, if that transaction results in the user having a bitcoin balance of zero. If the user's last transaction results in a bitcoin balance above zero, the user is regarded as active until the end of our sample in April 2017.	I & D
Darknet Sites	A transaction-weighted average of the number of operational illegal darknet marketplaces at the time a user transacts (the sum of number of operational darknet marketplaces at every transaction, divided by *Transaction Count*). The logic is that if a user transacts at a time when there is a lot of illegal darknet marketplace activity, they are more likely to be involved in illegal activity than if they are active when there is little or no illegal darknet activity.	I & D

Table 2.1 (continued)
Variable definitions

Variable	Definition	DCE equation
Tumbling	Tumbling refers to techniques or services used to obscure a user's holdings or transaction history. Wash transactions, in which a user is both the sender and receiver of bitcoin, are also sometimes used for such purpose. Illegal users are likely to have greater incentives to obscure their activity than legal users. We classify tumbling transactions using the following three approaches. Approach 1: transactions with known tumbling service providers (such as Coin Fog). Approach 2: transactions in which a user sends bitcoin to another user (potential tumbler) and that user sends the bitcoin back (in one or several transactions), less a tumbling fee of between 0 to 10 % within 10 blocks. Approach 3: transactions with users that display the characteristics of tumbling service providers (a *Transaction Count* of 10 or above and displays the two tumbling characteristics above in at least 8% of transactions). For each user, we compute their percentage of tumbling and wash transactions out of their total number of transactions.	I
Darknet Shock Volume	The percentage of the user's transaction value that occurs immediately after shocks to darknet marketplaces, including one week after each seizure or "exit scam" of a darknet marketplace. Seizures by law enforcement officials and "exit scams" in which darknet sites close without warning are likely to result in increased activity from illegal users as they turn to alternative marketplaces or relocate their holdings. At the same time, shocks to darknet marketplaces are unlikely to materially affect the activity of legal users.	I
Bitcoin Hype	The transaction-weighted average of the Google Trends value for "bitcoin" (calculated from Jan 1, 2009 to May 1, 2017). For each user, we record the intensity of Google searches for the term "bitcoin" (scaled from 0-100) in the months of their transactions and then compute the average for each user across all of their transactions. The logic is that the more intensive the search activity for bitcoin on Google is, the more likely the user is transacting for speculative (as opposed to illegal) purposes.	I
Bitcoin Market Cap	The transaction-weighted average log market capitalization of bitcoin at the time of each user's transactions. For each user, we calculate the log market capitalization of bitcoin at the time of each user's transactions. We then compute the average across all of the user's transactions. The logic is that as the value of bitcoin increases, the likelihood of illegal activity is lower as more speculators are present.	I
Shadow Coins	The transaction-weighted average log market capitalization of major opaque cryptocurrencies ("shadow coins": Dash, Monero, and ZCash) at the time of each user's transactions. The logic is that if illegal users make use of shadow coins, the likelihood of illegal activity in bitcoin will be lower when shadow coins are more prevalent.	I
Alt Coins	The transaction-weighted average log market capitalization of other non-privacy coins ("alt-coins": all cryptocurrencies excluding bitcoin and "shadow coins") at the time of each user's transactions. The logic is that when there is a lot of interest in alternative non-privacy cryptocurrencies, which cannot be used in darknet markets, all else equal, it is likely there is proportionally more legal/speculative trade in cryptocurrencies and thus a lower fraction of illegal activity.	I
Pre-Silk-Road User	Dummy variable that is equal to one if the user commenced transacting in bitcoin prior to the seizure of Silk Road 1 on October 1, 2013. The logic is that an illegal user that was using bitcoin prior to the first major darknet seizure by law enforcement authorities has a higher probability of having been detected than a user that started transacting in bitcoin after that seizure because such a user could not have been "detected" in the first seizure.	D

Table 2.2 reports descriptive statistics about the user-level variables. Focusing on the variables that characterize a user's bitcoin transaction activity (Panel A), we see that a typical (median) user engages in three bitcoin transactions (mean *Transaction Count* is 5.7 transactions) with three different counterparties (mean of *Counterparties* is around 4.2). Thus, a typical user has a low degree of concentration in counterparties, in that they do not repeatedly transact with the same counterparty (our measure of *Concentration*, which is a normalized Herfindahl-Hirschman Index, has a median of zero). There are a small number of highly active entities, with the most active having almost 11.4 million transactions and 4.4 million counterparties.

The average transaction size is around $5,000, but a typical transaction (the median *Transaction Size*) is much smaller at $112. Some transactions are very large, with the largest exceeding $90 million. For most users, their first and last bitcoin transaction occurs within the same month (the median *Existence Time* is one month), although some users are present for many years (the maximum *Existence Time* is 101 months, or just over eight years). The other variables (Panel B) are more specific indicators of the nature of the activity in which a user is likely to be engaged and are thus important in our empirical models. We therefore define and discuss these variables when we turn to the empirical models.

Table 2.2
Descriptive statistics for all users
This table reports descriptive statistics about bitcoin users. *Transaction Count* is the total number of bitcoin transactions involving a given user. *Transaction Size* (in USD) is the user's average transaction value. *Transaction Frequency* is the average rate at which the user transacts between their first and last transactions, annualized to transactions per year. *Counterparties* is the total number of other users with which the given user has transacted. *Holding Value* is the average value of the user's bitcoin holdings (in USD), where holdings are measured after each transaction. *Concentration* takes values between zero and one, with higher values indicating a tendency to repeatedly trade with a smaller number of counterparties. *Existence Time* is the number of months between the date of the user's first and last transaction. *Darknet Sites* is the average number of operational darknet sites at the time of each of the user's transactions. *Tumbling* is the percentage of the user's transactions that attempt to obscure the user's holdings (wash or tumbling trades). *Darknet Shock Volume* is the percentage of the user's total dollar volume that is transacted during the week after marketplace seizures or "exit scams". *Bitcoin Hype* is a measure of the intensity of Google searches for the term "bitcoin" around the time of the user's trades. *Pre-Silk-Road User* is a dummy variable taking the value one if the user's first bitcoin transaction is before the seizure of the Silk Road on October 2013. *Bitcoin Market Cap*, *Shadow Coins*, and *Alt Coins* are transaction-weighted average log market capitalizations of bitcoin, major opaque cryptocurrencies, and non-privacy cryptocurrencies excluding bitcoin, respectively, at the time of each user's transactions.

Variable	Mean	StdDev	Min	P25	Median	P75	Max
Panel A: Transactional characteristics							
Transaction Count	5.70	1,622.74	1.00	2.00	3.00	3.00	11,410,691.00
Transaction Size	5,207.61	56,939.00	1.00	22.06	111.91	668.44	92,504,688.00
Transaction Frequency	29.88	659.27	0.12	7.20	24.00	36.00	3,077,978.00
Counterparties	4.18	553.71	0.00	2.00	3.00	3.00	4,385,500.00
Holding Value	3,974.05	55,011.00	0.00	15.91	83.96	551.37	115,529,839.00
Concentration	0.10	0.28	0.00	0.00	0.00	0.00	1.00
Existence Time	6.61	11.91	1.00	1.00	1.00	5.00	101.00
Panel B: Characteristics associated with particular types of activity							
Darknet Sites	17.14	5.10	0.00	15.00	18.00	20.00	27.00
Tumbling	0.45	0.05	0.00	0.00	0.00	0.00	181.82
Darknet Shock Volume	16.51	0.36	0.00	0.00	0.00	0.00	100.00
Bitcoin Hype	28.29	15.44	0.00	19.00	24.00	38.00	100.00
Pre-Silk-Road User	0.07	0.26	0.00	0.00	0.00	0.00	1.00
Bitcoin Market Cap	9.82	0.49	5.14	9.71	9.94	10.09	10.40
Shadow Coins	7.07	2.52	0.00	7.28	7.78	8.32	9.10
Alt Coins	8.68	2.04	0.00	8.76	9.21	9.34	10.26

2.4 Identifying a sample of illegal users

We identify a sample of addresses (and therefore users) involved in illegal activity using three approaches described below.

2.4.1 First approach: Bitcoin seizures by law enforcement agencies

Our first approach exploits bitcoin seizures by law enforcement agencies such as the US FBI. We manually identify bitcoin seizures from news articles (via searches using Factiva) and US court records (via searches of the digital PACER records). Table A2 in Appendix A reports the list of seizures that we use. For each seizure, we extract information from court records and law enforcement agency disclosures about any identified bitcoin addresses or transactions (amounts and dates). From these details we uniquely identify the users involved in the illegal activity, by matching up the bitcoin address or transaction identifier with our user-level data constructed from the bitcoin blockchain.

In some cases (e.g., the US FBI's seizure of Silk Road and Ross Ulbricht's holdings, and the Australian law enforcement's seizure of Richard Pollard's holdings) the law enforcement agency auctioned the seized bitcoin to the public. Selling seized assets that are not themselves illegal is common practice among law enforcement agencies. Given the public nature of the auctions, we are able to identify the auction transactions on the bitcoin blockchain and work backwards to identify the seized bitcoin addresses, which in turn identify those individuals that were involved in illegal activity and had some or all of their bitcoin holdings seized by law enforcement agencies. Using this approach we identify 1,016 known illegal users, which we refer to as "*Seized Users*".

2.4.2 Second approach: Illegal darknet marketplaces and their users

Our second approach exploits the known "hot wallets" of major illegal darknet marketplaces. These are central accounts, many of which operate like escrow accounts, into which users of darknet marketplaces deposit or withdraw funds. We are able to identify 17 such marketplaces using data from the Wallet Explorer service, which in turn identifies these marketplaces using an approach similar to Meiklejohn et al. (2013), i.e., on the basis of small "probing" transactions undertaken with a given entity.

From these hot wallets, we identify slightly over six million darknet marketplace users as individuals that send to and/or receive bitcoin from a known darknet marketplace. We refer to the darknet marketplace hot wallets and their contributors/recipients as "*Black Market Users*".

An underlying assumption is that the trade that occurs in darknet marketplaces is illegal. This assumption is supported by ample anecdotal evidence, objective empirical evidence in the form of darknet market scrapes that show the goods and services traded there (e.g., Christin, 2013; Aldridge Décary-Hétu, 2014; Van Buskirk et al., 2014; Soska and Christin, 2015), as well as actions by law enforcement agencies, including indiscriminate seizures of *all* bitcoin from such markets.

2.4.3 Third approach: Users identified in darknet forums

Our third approach exploits information contained in the darknet, in particular the bitcoin addresses of users identified in darknet forums as selling goods/services. We use systematic scrapes of darknet forums from 2013 to 2017.[31] This allows us to identify users that might never have been caught by authorities and might not be otherwise identified in the data through transactions with known darknet marketplaces. Users often post bitcoin addresses in cases such as fraud (they did not receive their

[31] A list of known darknet markets is in Table A1 of Appendix. An archive of darknet forums during 2013-2015 is available at https://www.gwern.net/index. We scrape information from active darknet sites during 2016-2017.

goods), quality checking, and for the purposes of advertising the address to which funds should be sent, including in privately negotiated trade. While other studies have also scraped darknet marketplaces for certain types of information (e.g., Soska and Christin, 2015; Van Buskirk et al., 2016), as far as we know no other study has used scrapes to identify the bitcoin addresses of illegal users.

Using this approach, we identify an additional 448 users that were not already identified in either of the previous two approaches. We refer to these as "*Forum Users*".

2.4.4 The sample of illegal users

Table 2.3 shows the number of illegal users identified using the three approaches above and various measures of their activity. Together, there are 6,223,359 "observed" illegal users, representing 5.86% of all bitcoin participants. They account for an even larger share of transactions—a total of 196 million transactions, or around one-third of all transactions (32.38%). They also account for an even larger share of bitcoin holdings—throughout the sample period, the average dollar value of the bitcoin holdings of observed illegal users is around $1.3 billion, which is close to half (45.28%) of the average dollar value of holdings for all users.[32] Observed illegal users control around one-quarter (26.33%) of all bitcoin addresses and account for approximately 12.96% of the total dollar value of all bitcoin transactions.

Within the three subgroups of illegal users, the largest group in terms of number of users is the "*Black market users*", followed by "*Seized users*" and then "*Forum users*". *Seized users* and *Forum users* are nevertheless meaningful subgroups in terms of their share of total transactions.

The results in Table 2.3 indicate that the sample of "observed" illegal users is already a substantial proportion of users and bitcoin transaction activity, without yet having applied methods to estimate the population of illegal users/activity. Capturing a relatively large sample of illegal activity is important because it provides rich information to our empirical methods that estimate the totality of illegal activity. The fact that the sample of illegal activity is drawn from three different approaches is also likely to help the subsequent empirical models by providing a more diverse sample.

[32] The average holdings numbers are considerably lower than current holdings because for the first few years of bitcoin's existence, its market capitalization was much lower than it is currently.

Table 2.3
Size and activity of observed user groups
This table reports the size and activity of (1) all users, (2) observed illegal users, and (3) other users. The observed illegal user group has three subgroups: users that had bitcoin seized by law enforcement agencies ("*Seized Users*"), illegal darknet marketplace escrow accounts (hot wallets), users that have interacted (sent or received bitcoin) with those accounts ("*Black Market Users*"), and users whose bitcoin address(es) are mentioned in darknet forums ("*Forum Users*"). The measures of group size and activity are: the number of users (*Users*), the number of transactions (*Transaction Count*), the average dollar value of monthly bitcoin holdings (*Holding Value*), the number of bitcoin addresses (*Number Of Addresses*), and the dollar volume of transactions (*Volume*). The percentage of total users/activity is reported in parentheses below each value.

Group / Subgroup	Users	Transaction Count (Mil)	Holding Value ($Mil)	Number Of Addresses (Mil)	Volume ($Bil)
1. All Users	106,244,432	605.69	2,964.66	221.71	1,862.51
	(100.00%)	(100.00%)	(100.00%)	(100.00%)	(100.00%)
2. Observed Illegal Users	6,223,359	196.11	1,342.43	58.38	241.46
	(5.86%)	(32.38%)	(45.28%)	(26.33%)	(12.96%)
2A. Seized Users	1,041	23.83	9.39	8.30	17.51
	(0.00%)	(3.93%)	(0.32%)	(3.74%)	(0.94%)
2B. Black Market Users (not in 2A)	6,221,870	157.30	1,324.32	49.71	220.91
	(5.86%)	(25.97%)	(44.67%)	(22.42%)	(11.86%)
2C. Forum Users (not in 2A or 2B)	448	14.98	8.72	0.38	3.03
	(0.00%)	(2.47%)	(0.29%)	(0.17%)	(0.16%)
3. Other Users	100,021,073	409.58	1,622.23	163.33	1,621.05
	(94.14%)	(67.62%)	(54.72%)	(73.67%)	(87.04%)

A limitation of the sample of observed illegal users is that it predominantly contains users that are involved in buying and selling illegal goods and services online in darknet marketplaces. There are other forms of illegal activity that involve bitcoin, such as money laundering, evasion of capital controls, payments in ransomware attacks, and bitcoin thefts. Without an initial sample of these forms of illegal activity, our empirical models are likely to underestimate their prevalence. Illegal activity that is similar in characteristics to illegal activity in darknet markets, and illegal activity that involves transacting with darknet market participants will be captured by our estimates of illegal activity, even if it is not in our sample of directly observed illegal activity. However, illegal activity that is dissimilar to darknet market activity and does not interact with such participants is unlikely to be captured by our empirical models. Thus, our estimates are likely to underestimate some forms of illegal activity involving bitcoin.

Given the nature of illegal activity could change through time, it is also important that our sample of observed illegal users spans different time periods and is not completely concentrated at one point in time. Figure 2.2 indicates that this is the case for our sample of observed illegal users and their activity.

These time-series show that the observed illegal users are present during all points in time throughout our sample period. Their share of activity is highest at the start of the sample in 2009, and then again during a period from 2012 to the end of 2015. The first of these periods (the year 2009) is not particularly economically meaningful as the first year or two of bitcoin's existence involves a very small number of users and transactions compared to subsequent years. In contrast, the activity in the second period, 2012-2015, is meaningful. This period corresponds to the time when illegal darknet marketplaces grew rapidly in number and popularity. Silk Road 1 was established in January 2011 and soon became a popular venue in which to buy and sell illegal goods and services (e.g., Soska and Christin, 2015). After Silk Road 1 was shut down by the US FBI in October of 2013, a large number of other illegal darknet marketplaces commenced operating throughout 2013-2015 (see Table A1 of Appendix A). Thus, perhaps somewhat unsurprisingly, the peak activity of our sample of observed illegal users coincides with substantial darknet marketplace activity. However, we also observe a reasonable number of illegal users and illegal activity outside of this peak window.

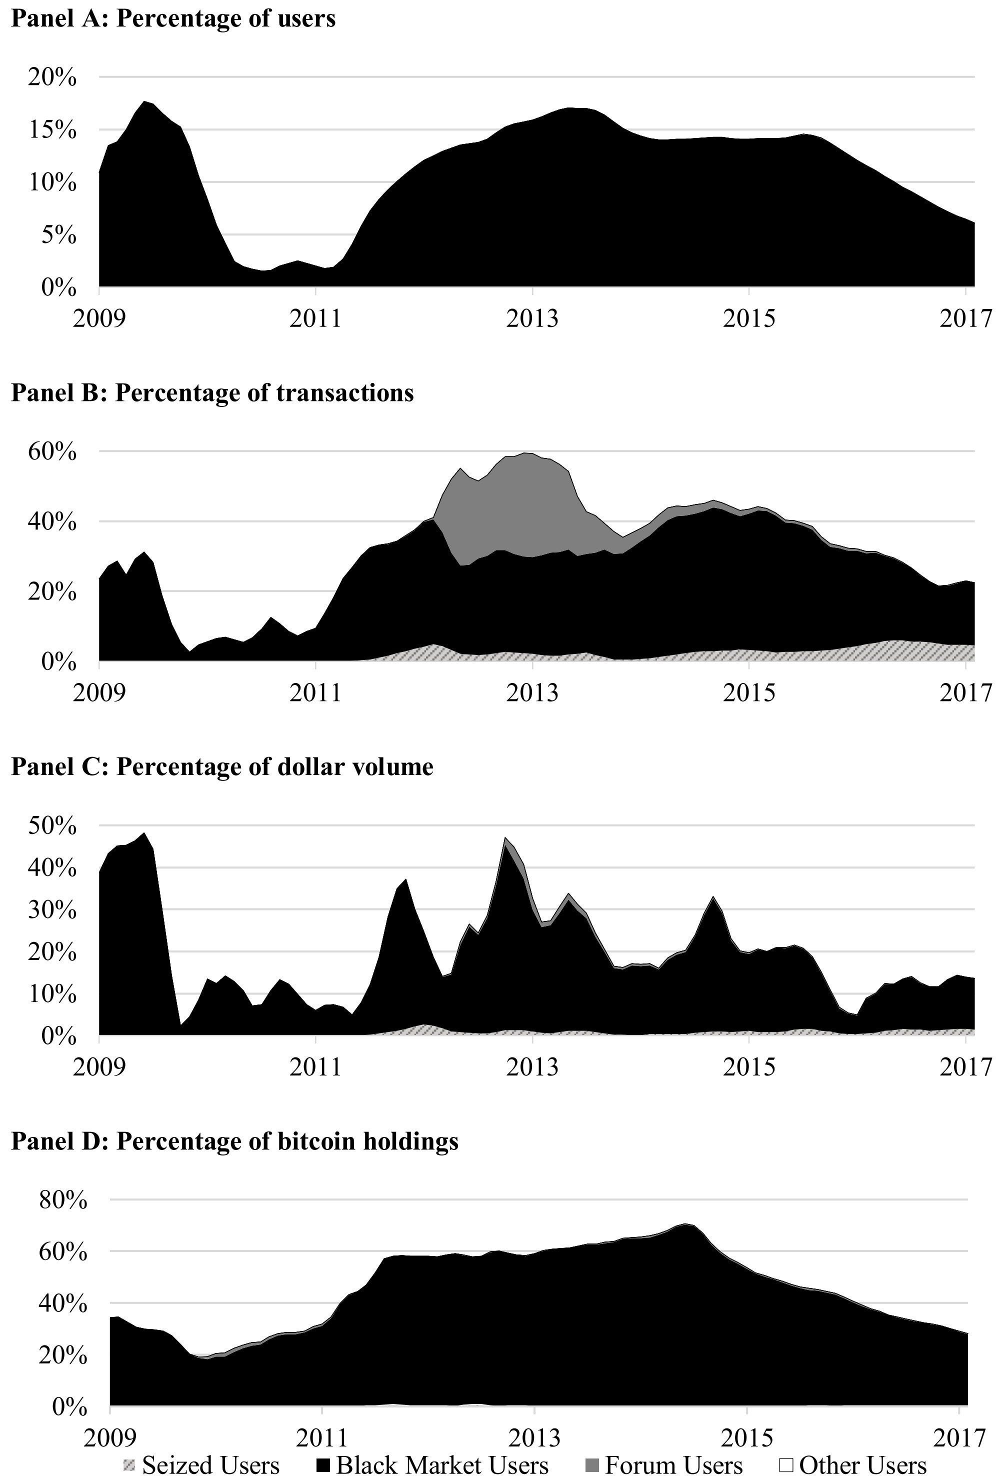

Panel A: Percentage of users

Panel B: Percentage of transactions

Panel C: Percentage of dollar volume

Panel D: Percentage of bitcoin holdings

Figure 2.2
Size and activity of the sample of "observed" illegal bitcoin users
This figure illustrates the time-series of the three subgroups of observed illegal users as a percentage of total users (Panel A), their number of transactions as a percentage of all transaction (Panel B), the dollar value of their transactions as a percentage of the dollar value of all transactions (Panel C), and the dollar value of their bitcoin holdings as a percentage of the dollar value of all bitcoin holdings (Panel D). The observed illegal user group includes three subgroups: users that had bitcoin seized by law enforcement agencies ("*Seized Users*"), illegal darknet marketplace escrow accounts (hot wallets), and users that have sent or received bitcoin from those accounts ("*Black Market Users*"), and users whose bitcoin address(es) are mentioned in darknet forums ("*Forum Users*"). "*Other Users*" corresponds to all bitcoin users other than those in the sample of observed illegal users. The values are smoothed using a three-month moving average.

2.5 Quantifying and characterizing all illegal activity

Having identified a substantial sample of bitcoin users that are involved in illegal activity, our next step is to use the information in this sample to estimate the totality of illegal activity that uses bitcoin. We use two different methods to classify users into those that are primarily involved in illegal activity ("illegal users") and those that are primarily involved in legal activity ("legal users"). Subsequently, we measure the size and activity of the two groups.

At an intuitive level, the first method exploits the network topology—the information about who trades with whom. Trade networks reveal "communities" of users and can thereby identify other illegal users that were not part of our initial sample. In contrast, the second method exploits characteristics that distinguish illegal users from legal users (controlling for non-random detection).

Both methods allow a user that was initially classified as an "observed" illegal user to be reclassified as a user that is predominantly engaged in legal activity (a "legal user"). Averaging across the three categories of observed illegal users, 0.17% of all users (2.84% of observed illegal users) are reclassified as legal users by the models that exploit the network topology and 0.60% of all users (10.17% of observed illegal users) are reclassified as legal users by the models that exploit characteristics of users. The reclassified users reflect both (i) errors in the classification models, and (ii) users that predominantly engage in legal activity but have some involvement in illegal activity. The relatively low reclassification rates suggest that most of the "observed" illegal users predominantly use bitcoin for illegal activity.

The two methods provide independent estimates of the illegal activity and its characteristics. Given that the methods rely on completely different assumptions and exploit different information, their concurrent use provides robustness and the ability to cross-validate results. The methods are described below in separate subsections. We then report the results of how many users and how much trade is estimated to be associated with illegal activity, after which we characterize the nature of the illegal users and their trading activity compared to legal users.

2.5.1 Method 1: Network cluster analysis

The first method exploits network topology to identify "communities" of users based on the transactions between users. In simple terms, the method works as follows. If users A, B, and C are known to be involved in illegal activity (e.g., their bitcoin was seized by law enforcement agencies), a user X that trades exclusively or predominantly with users A, B, or C is likely to also be involved in illegal activity. Similarly, a user Y that trades predominantly with users that are not identified as illegal is likely to be a legal user. This intuition drives the classification of users into legal and illegal on the basis of their transaction partners.

More formally, the method we apply is a network cluster analysis algorithm that takes as inputs the set of users ("nodes" in network terminology) and the trades between users ("edges" or "links" in network terminology). The output of the algorithm is an assignment of users to communities such that the "modularity" of the communities (density of links within communities and sparsity of links between communities) is maximized. The method labels a user as illegal (legal) if the disproportionate share of their transactions is with members of the illegal (legal) community. The method does not assume that users only engage in either legal or illegal activity—users can do both. Therefore, there will be some trades between the legal and illegal communities.

We apply a variant of the Smart Local Moving (SLM) algorithm developed by Waltman and van Eck (2013), adapted to our specific application. The algorithm's name ("smart moving") comes from the fact that the algorithm finds the underlying community structure in the network by moving nodes from one community to another, if such a move improves the model fit. The SLM algorithm is among the leading network cluster analysis algorithms.[33] Applied to our data, the algorithm is as follows:

- Step 1: Assign all the observed illegal users to the illegal community and all of the remaining users to the legal community.
- Step 2: Loop through each user, performing the following action on each:
    - If the user disproportionately transacts with members of the user's currently assigned community, then leave the user in that community[34];
    - Otherwise, move the user to the other community (if the user is assigned to the illegal community, move the user to legal community, and vice versa).
- Step 3: Repeat Step 2 until, in a complete loop through all users, no user switches between communities. At that point the assignment to communities is stable and ensures that each member trades disproportionately with other members of the same community.

Due to the iterative nature of the algorithm, not all of the "observed" illegal users will necessarily remain in the illegal community. For example, it is possible that some of the users that had bitcoin seized by authorities were involved in some illegal activity (hence getting bitcoin seized) but were mainly using bitcoin for legal purposes. This will be recognized by the algorithm in Step 2 and the user will be moved to the legal community.

[33] For example, Emmons et al. (2016) in their comparison of multiple methods find that the SLM algorithm performs the best in terms of maximizing cluster quality metrics.

[34] "Disproportionately" is if the proportion of transactions the user makes with other members of the same community is greater than or equal to the community's proportion of total transactions. In robustness tests we consider the proportion of volume transacted rather than transactions and find consistent results.

2.5.2 Method 2: Detection-controlled estimation (DCE)

The second method we use to estimate the population of users involved in illegal activity ("illegal users") is detection-controlled estimation (DCE). Intuitively, this method exploits the differences in the characteristics of legal and illegal users of bitcoin to probabilistically identify the population of illegal users. If we had a random sample of illegal users and a set of characteristics that differ between legal and illegal users (e.g., measures of the extent to which a user has employed tools to conceal their activity), this task would be relatively simple and could be achieved with standard techniques (regression, discriminant analysis, and so on). A complication is that detection (as in most settings where violators attempt to conceal their illegal activity from authorities) is not random, and this non-randomness must be accounted for to obtain unbiased estimators.[35] We use "detection" in the broad sense of an illegal user having been identified by any of the three approaches to detecting illegal users described in Section 4.

Fortunately, this econometric challenge is not unique to illegal activity in bitcoin and methods to overcome it exist. The same challenge occurs in quantifying other forms of misconduct such as tax evasion, fraud, insider trading, and market manipulation, as well as contexts such as nuclear power plant safety regulation breaches, cancer detection by mammograms, and so on. The standard tool for these settings is DCE. Since its development by Feinstein (1989, 1990), DCE models have been applied to various financial misconduct settings including tax evasion (Feinstein, 1991), corporate fraud (Wang et al., 2010), and market manipulation (Comerton-Forde and Putniņš, 2014). By explicitly modelling both underlying processes (violation and detection) simultaneously, one can obtain unbiased estimates of the illegal activity, which is otherwise only partially observed.

Figure 2.3 illustrates the two-stage DCE model that we estimate. On the left is the starting point, the data, which in our case is the set of all bitcoin users. In the middle we have the two processes, violation (undertaking illegal activity) and detection (e.g., bitcoin seizures). On the right-hand side are the joint outcomes of those processes: the observable classifications of users into detected illegal users (the set A) and other users (the complement set A^C, comprising legal users and undetected illegal users).

[35] A further complication is that the determinants of this non-randomness are not separately observed (unlike, for example, non-respondents in a survey, or people that choose not to participate in the labor force) and therefore the classic tools to deal with sample selection bias (e.g., Heckman models) cannot be applied.

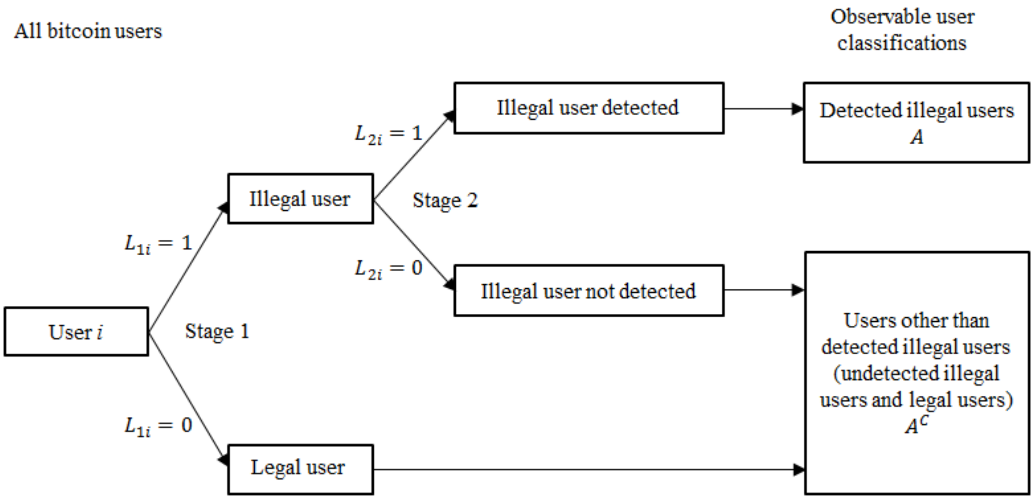
All bitcoin users
Observable user classifications
User i
Stage 1
$L_{1i}=1$
$L_{1i}=0$
Illegal user
Legal user
Stage 2
$L_{2i}=1$
$L_{2i}=0$
Illegal user detected
Illegal user not detected
Detected illegal users A
Users other than detected illegal users (undetected illegal users and legal users) A^C

Figure 2.3
Two-stage detection-controlled estimation (DCE) model
The figure illustrates the structure of the two-stage DCE model. Stage 1 models how legal and illegal users of bitcoin differ in characteristics. Stage 2 models the determinants of the probability that an illegal user was "detected" (had bitcoin seized by a law enforcement agency, was identified in darknet forums, or was observed in the blockchain data as having transacted with a known illegal darknet marketplace). Both stages are estimated simultaneously using maximum likelihood to select parameter values that maximize the likelihood of the observable user classifications, A and A^C.

The first branch models whether a bitcoin user, i, is predominantly involved in illegal or legal activity. This branch is modelled as an unobservable binary process (L_{1i}) driven by a continuous latent function (Y_{1i}) of a vector of characteristics, x_{1i}, that can distinguish between legal and illegal users:

$$Y_{1i} = \beta_1 x_{1i} + \epsilon_{1i} \tag{2.1}$$

$$L_{1i} = \begin{cases} 1 & (Illegal\ user) \\ 0 & (Legal\ user) \end{cases} \text{ if } \begin{Bmatrix} Y_{1i} > 0 \\ Y_{1i} \leq 0 \end{Bmatrix} \tag{2.2}$$

The second branch models whether or not an illegal user is "detected" (they enter our sample of observed illegal users). This detection process is modelled as another unobservable binary process (L_{2i}) driven by a different continuous latent function (Y_{2i}) of a vector of characteristics, x_{2i}, that affect the probability that an illegal user is detected:

$$Y_{2i} = x_{2i}\beta_2 + \epsilon_{2i} \tag{2.3}$$

$$L_{2i} = \begin{cases} 1 & (Detected) \\ 0 & (Not\ detected) \end{cases} if \begin{Bmatrix} Y_{2i} > 0 \\ Y_{2i} \leq 0 \end{Bmatrix} \tag{2.4}$$

Both stages of the model are estimated simultaneously using maximum likelihood. The likelihood function for the model is derived in Appendix B. Intuitively, this process finds estimates for the vectors of model parameters, β_1 and β_2, that maximize the likelihood of the observed data (the classification of users into sets A and A^C). From the estimates of β_1 and β_2, we compute each user's probability of being involved in illegal activity and construct a binary classification of legal and illegal users.

Similar to the SLM approach, the DCE model does not assume that detected illegal users were engaged solely or predominantly in illegal activity. Once the DCE model is estimated, the classification of users into legal and illegal categories can result in some detected illegal users being reclassified as predominantly legal users.[36]

Similar to Heckman models, identification in a DCE model without instruments is possible, relying on functional form and distributional assumptions. However, more robust identification is achieved through instrumental variables that affect one process but not the other. We take the more robust route of using instrumental variables. The next subsection describes the instrumental variables and their descriptive statistics.

2.5.3 Variables used in the DCE model and their descriptive statistics

One of the instrumental variables associated with illegal activity is the extent to which the user employs methods to conceal their identity or obfuscate their transaction history. For example, to partially conceal their identities from an observer of the bitcoin blockchain, users can use “tumbling” and “wash trades” to alter the addresses of their bitcoin holdings, increasing the difficulty of tracing their activity. Tumbling, in its simplest form, involves a user sending bitcoin to a tumbling provider who (in return for a small fee) returns the balance to a different address controlled by the user. Wash trades involve a user sending bitcoin from one address to another (new) address that they also control. Legal users have little reason to take such actions to conceal their actions (and incur associated costs). In contrast, users involved in illegal activity are likely to use these concealment techniques. As such, the use of tumbling services and wash trades is likely to be a predictor of whether a user is involved in illegal activity. Importantly (for this to be an instrumental variable), using wash trades and tumbling does not alter the probability of “detection” by law enforcement agencies via the seizures of bitcoin from darknet sites. The seizures confiscated all bitcoin held in darknet marketplace escrow accounts (“hot wallets”) irrespective of whether the user employed tumbling or wash trades. For each user, we measure the percentage of their transactions that are tumbling or wash trades and call this variable *Tumbling*.

Another set of instruments for the likelihood that a user is involved in illegal activity involves time-series variables that are likely to correlate with the type of activity in which bitcoin users are engaged. For example, for each user we construct a measure of the average number of operational illegal darknet marketplaces at the time the user transacts (we label the variable *Darknet Sites*). All else equal,

[36] For example, suppose a user was involved in some illegal activity and had bitcoin seized by authorities but was mainly using bitcoin for legal purposes. Such a user will have characteristics that are similar to those of legal users and not very similar to illegal users, which would lead to a classification by the DCE model into the legal user category. In contrast, a predominantly illegal user, even if not detected or observed, is likely to have characteristics similar to other illegal users and therefore (after controlling for the differences in characteristics due to non-random detection) the user is likely to be classified as illegal by the DCE model.

illegal transactions (and thus users involved in illegal activity) are more likely when there is a lot of illegal darknet marketplace activity than when there is little or no illegal darknet activity.

In a similar spirit, we construct a measure of the popularity of opaque cryptocurrencies (Dash, Monero, and ZCash). This measure, which we label *Shadow Coins*, is the average log market capitalization of the opaque cryptocurrencies at the time of a user's transaction. These "shadow coins" were developed to provide more privacy than bitcoin. If criminals are drawn to these shadow coins and start using them instead of bitcoin, the probability that a bitcoin user is involved in illegal activity will be inversely related to the market capitalization of such coins.

We also measure the popularity of bitcoin using its log market capitalization and the Google Trends search intensity for the keyword "bitcoin". We label these variables *Bitcoin Market Cap* and *Bitcoin Hype*, respectively. We also measure the popularity of other cryptocurrencies using the total log market capitalization of cryptocurrencies excluding bitcoin and the shadow coins (we label the variable *Alt Coins*). We measure these three variables at the time of each user's transaction and then for each user we average each variable across the user's transactions. To the extent that these variables correlate with speculative trading in bitcoin and mainstream (legal) use of cryptocurrencies, they will have an inverse association with the likelihood that a given user is involved in illegal activity. To avoid issues with co-linearity, we do not concurrently include *Bitcoin Market Cap* and *Bitcoin Hype* in the DCE models.

Our final instrument for involvement in illegal activity exploits the anecdotal evidence that significant darknet marketplace shocks such as seizures of darknet marketplaces by law enforcement agencies or closures of such marketplaces due to scams or hacks result in a brief spike of transaction activity by illegal users as they turn to alternative marketplaces or relocate their holdings in response to the shock. At the same time, shocks to darknet marketplaces are unlikely to materially affect the activity of legal users. Therefore, for each user, we measure the fraction of the user's transaction value that occurs in the one week period after each major darknet marketplace shock (marketplace "raids", "scams", and "hacks" in Table A1 of Appendix A). We label this variable *Darknet Shock Volume*.

As determinants of the probability of detection, we include a binary variable for whether the user started using bitcoin (date of first bitcoin transaction) before the first bitcoin seizure by law enforcement agencies from Silk Road 1 (we label the variable *Pre-Silk-Road User*). Because users that enter the bitcoin network after the first seizure can only be detected in subsequent seizures, post-Silk-Road-seizure users are likely to have a lower detection probability.

A few things are worth noting about the variables used in the DCE model. First, while the instrumental variables help identify the model, they are not the only characteristics that help separate legal and illegal users—the full set of characteristics used in the model serve that purpose, including variables common to both detection and violation equations (they have different coefficients in each equation). The full list of variables is presented in Table 2.1. Second, identification of the model requires only one variable that is associated with either the probability of being involved in illegal activity or the probability of detection, but not both. We have more candidate instrumental variables than this minimum of one, and in robustness tests we examine how sensitive the results are to the assumptions about these instruments. We do so by relaxing the assumed exclusion restrictions on a subset of the instruments one at a time, from which we conclude that the results are not particularly sensitive to any individual instrumental variable's exclusion restriction.

Table 2.2 Panel B reports descriptive statistics about the variables that serve as instruments. *Darknet Sites* indicates that for the average bitcoin participant, there are on average 17 operational darknet marketplaces around the time of their transactions. This number ranges from a minimum of zero to a maximum of 27. *Tumbling* indicates that only a relatively small proportion of users (less than 25%) engage in "tumbling" and/or "wash trades" to obscure the user's holdings. Thus, while techniques exist to help a bitcoin user conceal their activity, it appears that few bitcoin users adopt such techniques.

The variable *Darknet Shock Volume* indicates that while most users do not trade in the period immediately following darknet shocks (median of zero), some users conduct a large fraction of their trading during these periods, with the average bitcoin user undertaking around 17% of their trading following darknet shocks. The variable *Bitcoin Hype* indicates that for the average user, the intensity of Google searches for "bitcoin" is around 28% of its maximum of 100%. The *Pre-Silk-Road User* dummy indicates that only around 7% of all bitcoin participants started transacting before October 2013, when the first darknet marketplace seizure by law enforcement agencies occurred (the seizure of Silk Road 1 by the FBI). The mean/median and maximum of *Bitcoin Market Cap* are close in value indicating that the majority of users make the majority of their transactions in bitcoin around bitcoin's peak market value.

2.5.4 How much illegal activity involves bitcoin?

Both methods—network cluster analysis (SLM) and detection-controlled estimation (DCE)—arrive at probabilistic classifications of bitcoin users into those primarily involved in legal activity and those primarily involved in illegal activity. Once the users have been partitioned into the legal and illegal "communities", we use those categorizations to quantify the size and activity of the two groups.

Table 2.4 presents the main results at the aggregate level, across the sample period. Panel A reports the estimated size of the groups and their level of activity, while Panel B re-expresses these values as percentages for each group. First, the percentage of bitcoin users estimated to be predominantly involved in illegal activity is 29.12% using the SLM and 23.23% using the DCE, giving a midpoint estimate of about one-quarter of bitcoin users (26.17%, the average of the estimates from the two models). The 99 percent confidence interval around this estimate is 20.13% to 32.21%.[37] The

[37] We use a form of bootstrapped standard errors to form the confidence interval. First, we obtain standard errors from the DCE model using a bootstrap of 200 samples in which, for computational reasons, we are forced to reduce the sample size by taking a random sample (this is a conservative step as it inflates the estimated standard errors relative to the standard errors for the full sample size). We add to these standard errors the estimation

midpoint estimate suggests around 27.81 million bitcoin users are predominantly involved in illegal activity, versus 78.44 million legal users.

The estimated number of illegal users is around four times larger than our sample of observed illegal users. Given our sample of observed illegal users is based on a comprehensive approach and includes all users that can be observed transacting with one of the known darknet marketplaces, the results suggest that without empirical methods such as the SLM or DCE, illegal activity that can be inferred from involvement with known darknet marketplaces represents only a small (and likely non-random) fraction of all illegal activity. Thus, our results suggest that studies of known/identifiable darknet markets (e.g., Soska and Christin, 2015; Meiklejohn et al., 2013) only scratch the surface of all illegal activity involving bitcoin.

Table 2.4 also indicates that illegal users account for an even larger share of all transactions—around 46.17% (45.67% using the SLM and 46.67% using the DCE) or approximately 280 million transactions. Thus, the average illegal user is involved in more transactions than the average legal user. This result is consistent with the notion that illegal users are likely to use bitcoin as a payment system (which involves actively transacting), whereas legal users may hold bitcoin for reasons such as speculation. A similar proportion is observed for holding values—illegal users on average hold around one-half (49.22%) of all outstanding bitcoins. One reason for the large share of illegal user holdings (relative to their share of the number of users) is related to the calculation of this variable as a time-series average. A high fraction of illegal users early in the sample (when there are fewer bitcoin users) can generate such a result even if the holdings *per user* are lower among illegal users compared to legal users.

Illegal users are estimated to control around 39.31% of bitcoin addresses and account for about one-fifth (23.06%) of the dollar volume of bitcoin transactions. In dollar terms, illegal users conduct approximately $429 billion worth of bitcoin transactions. Because illegal users account for a larger share of transactions than their share of dollar volume, they tend to make smaller value transactions than legal users. This result is consistent with illegal users primarily using bitcoin as a payment system rather than holding it as an investment or speculative asset.

Three general conclusions can be drawn from the results in Table 2.4. First, illegal users account for a sizeable proportion of both users and trading activity in bitcoin, with the exact proportion varying across different measures of activity and the two estimation models. Second, the estimates from both the SLM and DCE are fairly similar across the various activity measures, despite relying on completely different assumptions and information. Third, even a fairly comprehensive approach to identifying illegal activity directly (such as the approach used in the previous section and that used in other darknet

uncertainty from the SLM model, which is captured by the users that cannot be uniquely assigned to the legal or illegal categories. We then apply the conservative bootstrapped standard errors to approximate the error in the midpoint estimate.

market studies) only captures a small fraction of the total illegal activity, highlighting the importance of extrapolation beyond a directly observed sample.

Table 2.4
Estimated size and activity of legal and illegal user groups
This table reports the size and activity of legal and illegal user groups. The measures of group size and activity are: the number of users (*Users*), the number of transactions (*Transaction Count*), the average dollar value of bitcoin holdings (*Holding Value*), the number of bitcoin addresses (*Number Of Addresses*), and the dollar volume of transactions (*Volume*). Panel A reports the values of these measures for the two user groups, while Panel B expresses the measures for each group as a percentage of the total. Different rows report different approaches to classifying the legal and illegal user groups. *SLM* provides estimates from the network cluster analysis approach to classification (a variant of the "Smart Local Moving" algorithm). *DCE* provides estimates from the detection-controlled estimation (DCE) approach to classification, which exploits the characteristics of legal and illegal users. *Midpoint* is the average of the estimates from the SLM and DCE models. *Upper bound* and *Lower bound* provide a 99 percent confidence interval around the *Midpoint*, using a form of bootstrapped standard errors.

Group	Classification	Users (Mil)	Transaction Count (Mil)	Holding Value ($Mil)	Number Of Addresses (Mil)	Volume ($Bil)
Panel A: Values						
Illegal	SLM	30.94	276.63	1,394.76	87.95	436.78
	DCE	24.68	282.70	1,523.87	86.35	422.05
	Upper bound	34.22	308.72	1,782.44	99.67	558.23
	Midpoint	27.81	279.67	1,459.32	87.15	429.41
	Lower bound	21.39	250.62	1,136.20	74.63	300.60
Legal	SLM	75.31	329.06	1,569.90	133.76	1,425.73
	DCE	81.57	322.99	1,440.79	135.37	1,440.45
	Upper bound	84.86	355.07	1,828.47	147.09	1,561.91
	Midpoint	78.44	326.02	1,505.35	134.56	1,433.09
	Lower bound	72.02	296.97	1,182.22	122.04	1,304.28
Panel B: Percentages						
Illegal	SLM	29.12%	45.67%	47.05%	39.67%	23.45%
	DCE	23.23%	46.67%	51.40%	38.95%	22.66%
	Upper bound	32.21%	50.97%	60.12%	44.96%	29.97%
	Midpoint	26.17%	46.17%	49.22%	39.31%	23.06%
	Lower bound	20.13%	41.38%	38.32%	33.66%	16.14%
Legal	SLM	70.88%	54.33%	52.95%	60.33%	76.55%
	DCE	76.77%	53.33%	48.60%	61.05%	77.34%
	Upper bound	79.87%	58.62%	61.68%	66.34%	83.86%
	Midpoint	73.83%	53.83%	50.78%	60.69%	76.94%
	Lower bound	67.79%	49.03%	39.88%	55.04%	70.03%

2.5.5 How does the illegal activity vary over time?

There is interesting time-series variation in the amount of illegal activity and its share of all bitcoin activity. Figures 2.4 to 2.7 plot the estimated amount of illegal activity that uses bitcoin through time from the first block in 2009 to 2017. The figures show the estimated number of illegal users, the number and dollar value of their transactions, and the value of their bitcoin holdings. Panel B of each of the figures shows these activity measures as a percentage of the total across all bitcoin participants.[38]

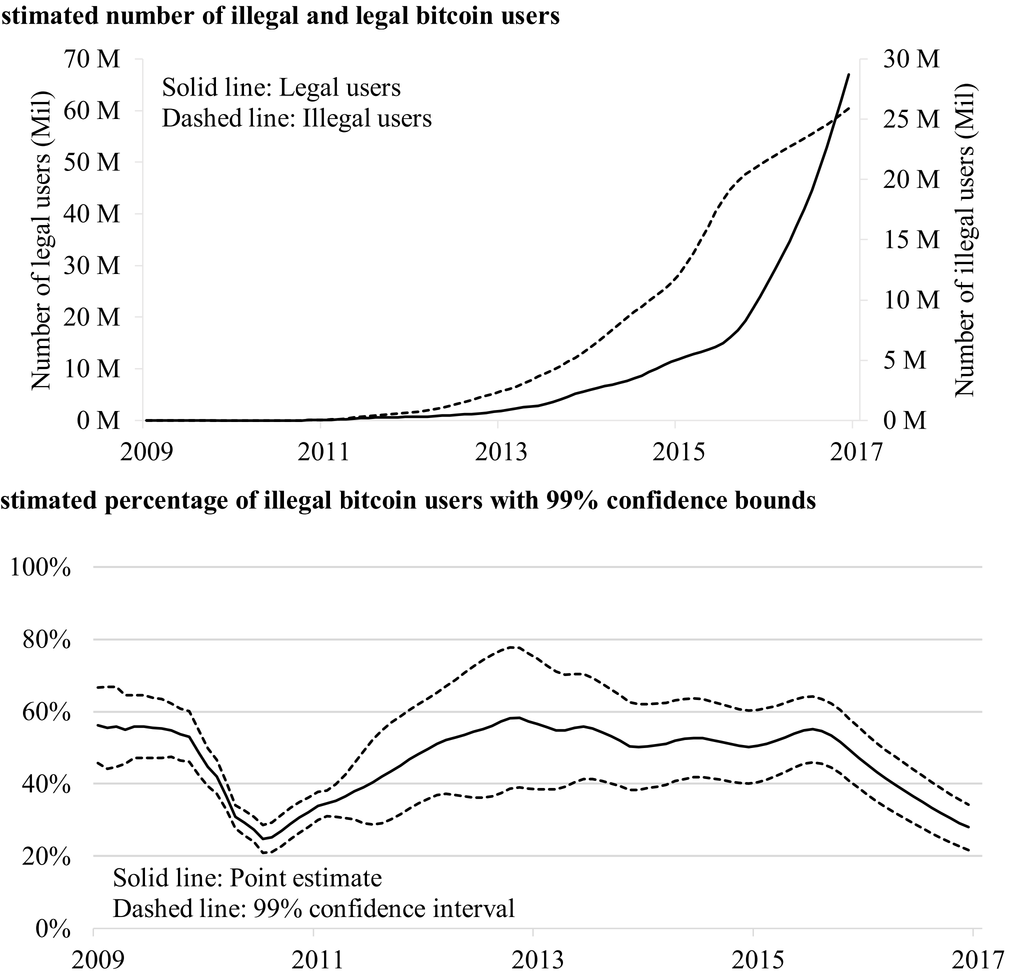
Panel A: Estimated number of illegal and legal bitcoin users

Panel B: Estimated percentage of illegal bitcoin users with 99% confidence bounds

Figure 2.4
Estimated number and percentage of bitcoin users involved in illegal activity
This figure illustrates the time-series of the estimated number of illegal and legal bitcoin users (Panel A) and the percentage of illegal users (Panel B). In Panel A, the number of legal users is plotted with the solid line using the left-hand-side axis and the number of illegal users is plotted with the dashed line using the right-hand-side axis. In Panel B, the solid line is the point estimate of the percentage of illegal users and the dashed lines provide a 99 percent confidence interval using bootstrapped standard errors. The estimates come from a combination of two empirical models (the average of the estimates produced by the SLM and DCE models). Values are smoothed using a moving average.

[38] Figures 2.4-2.7 use the average of the SLM and DCE model estimates. The SLM and DCE time-series estimates are separately reported in Figures C1-C8 in Appendix C.

Panel A: Estimated number of illegal and legal bitcoin user transactions per month

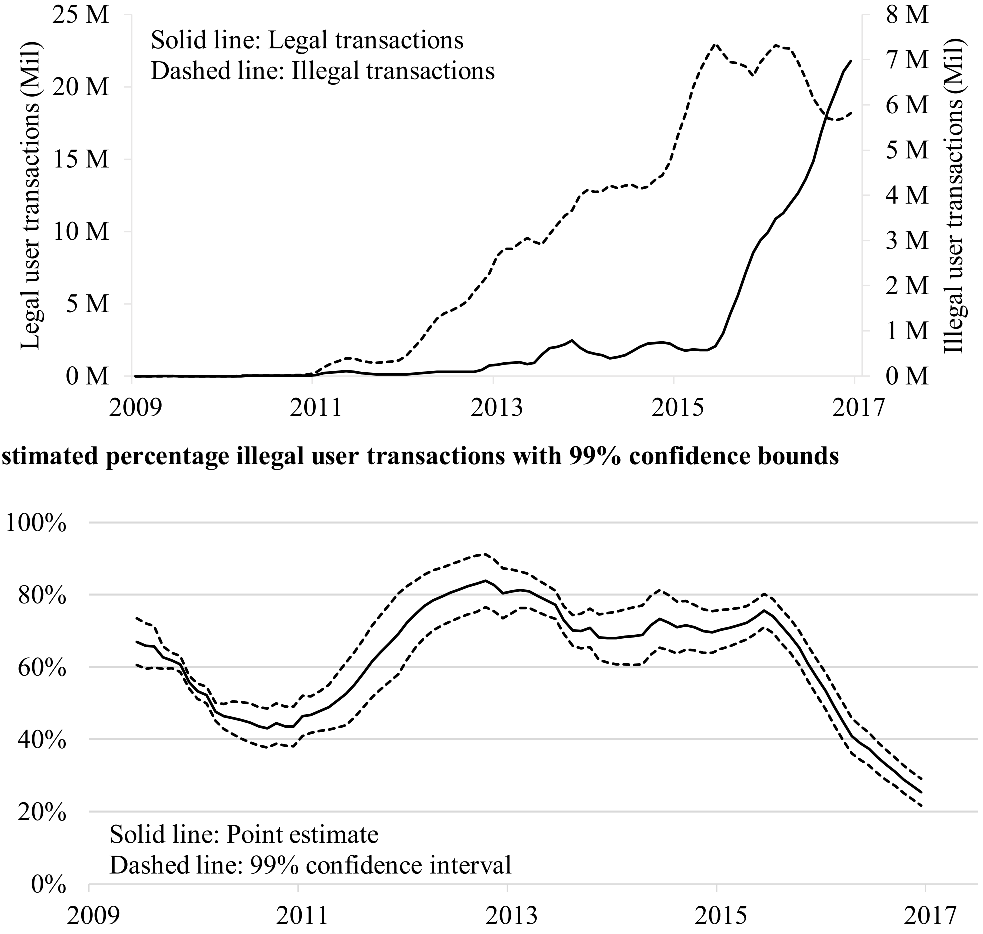

Panel B: Estimated percentage illegal user transactions with 99% confidence bounds

Figure 2.5
Estimated number and percentage of illegal bitcoin users transactions
This figure illustrates the time-series of the estimated number of illegal and legal bitcoin user transactions per month (Panel A) and the percentage of illegal user transactions (Panel B). In Panel A, the number of legal user transactions is plotted with the solid line using the left-hand-side axis and the number of illegal user transactions is plotted with the dashed line using the right-hand-side axis. In Panel B, the solid line is the point estimate of the percentage of illegal user transactions and the dashed lines provide a 99 percent confidence interval using bootstrapped standard errors. The estimates come from a combination of two empirical models (the average of the estimates produced by the SLM and DCE models). Values are smoothed using a moving average.

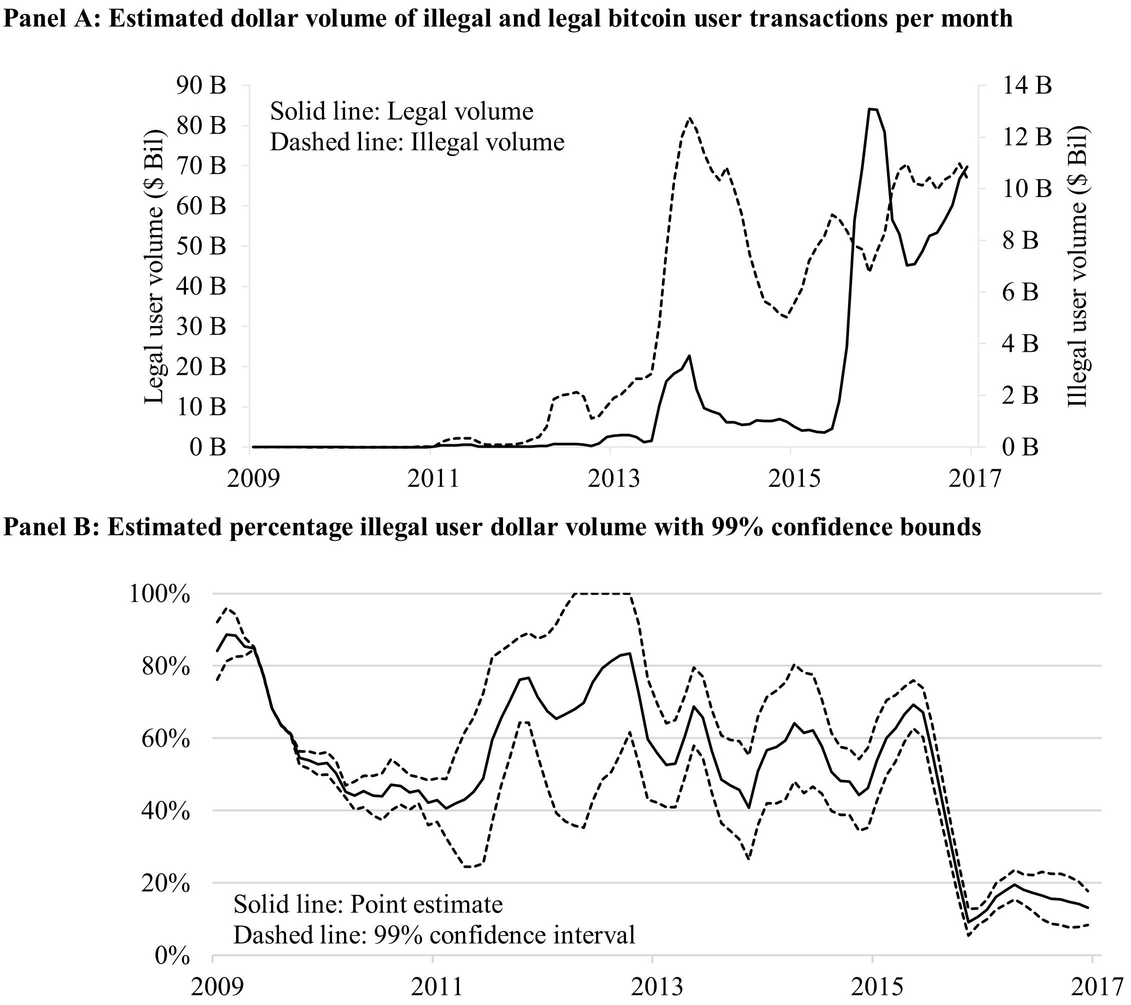

Panel A: Estimated dollar volume of illegal and legal bitcoin user transactions per month

Panel B: Estimated percentage illegal user dollar volume with 99% confidence bounds

Figure 2.6

Estimated dollar volume and percentage dollar volume of illegal bitcoin user transactions

This figure illustrates the time-series of the estimated dollar volume of illegal and legal bitcoin user transactions per month (Panel A) and illegal user dollar volume as a percentage of total dollar volume of bitcoin transactions (Panel B). In Panel A, the dollar volume of legal user transactions is plotted with the solid line using the left-hand-side axis and the dollar volume of illegal user transactions is plotted with the dashed line using the right-hand-side axis. In Panel B, the solid line is the point estimate of the illegal dollar volume as a percentage of total dollar volume and the dashed lines provide a 99 percent confidence interval using bootstrapped standard errors. The estimates come from a combination of two empirical models (the average of the estimates produced by the SLM and DCE models). Values are smoothed using a moving average.

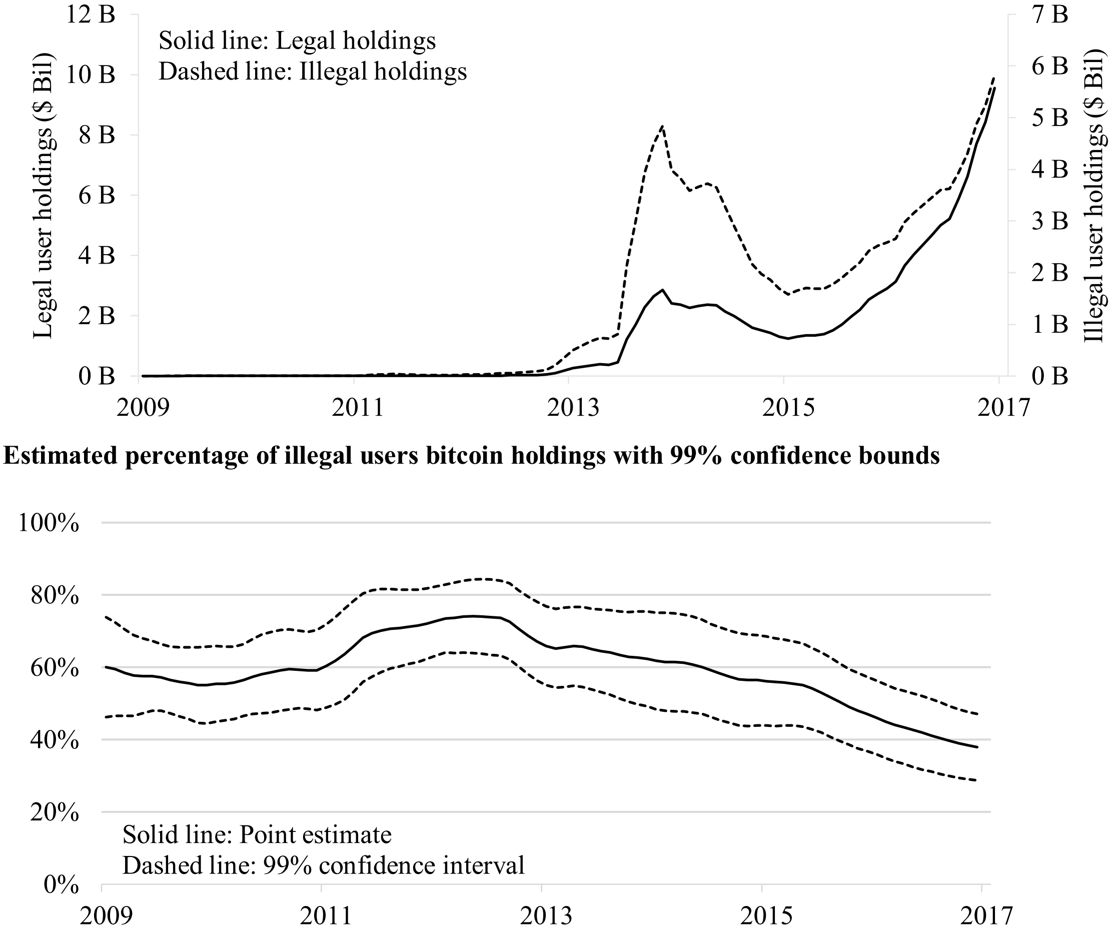

Panel A: Estimated dollar value of illegal and legal user bitcoin holdings

Panel B: Estimated percentage of illegal users bitcoin holdings with 99% confidence bounds

Figure 2.7
Estimated dollar value and percentage of illegal user bitcoin holdings
This figure illustrates the time-series of the estimated dollar value of illegal and legal user bitcoin holdings (Panel A) and illegal user holdings as a percentage of total bitcoin holdings (Panel B). In Panel A, the dollar value of legal user bitcoin holdings is plotted with the solid line using the left-hand-side axis and the dollar value of illegal user holdings is plotted with the dashed line using the right-hand-side axis. In Panel B, the solid line is the point estimate of the illegal user holdings as a percentage of total bitcoin holdings and the dashed lines provide a 99 percent confidence interval using bootstrapped standard errors. The estimates come from a combination of two empirical models (the average of the estimates produced by the SLM and DCE models). Values are smoothed using a moving average.

A pattern that is observed across all activity measures is that illegal activity, as a percentage of total bitcoin activity, tends to be high at the start of the sample in 2009, and then again from 2011 to the beginning of 2016, after which it steadily declines through to 2017. The activity levels indicate that there is only a very small (negligible) level of activity in bitcoin until about the middle of 2011, so the activity at the start of the sample is not economically meaningful. In contrast, the high relative level of illegal activity between 2012 and 2016 is noteworthy and coincides with the growth in the number of illegal darknet marketplaces, starting with the Silk Road in 2011. After the Silk Road was shut down in October 2013, a large number of other illegal darknet marketplaces commenced operating between 2013 and 2015 (Table A1 of Appendix A).

What could drive the decline in the relative level of illegal activity from beginning of 2016 onwards? The first thing to note is that the decline is observed in relative terms (that is, illegal activity as a fraction of total bitcoin activity), but *not* in absolute terms. Thus, it is not the case that the level of illegal activity in bitcoin has declined in recent years; rather, there has been a disproportionate increase in the legal use of bitcoin since the beginning of 2016. For example, from the beginning of 2016 to April 2017, the estimated number of illegal bitcoin users increases from around 21 million to around 27 million, reflecting growth of almost 30%, whereas the estimated number of legal bitcoin users increases from around 21 million to around 77 million, reflecting growth of around 250%. The rapid growth of legal use is likely driven by factors such as increased interest from investors and speculators (e.g., the emergence of "cryptofunds", and more recently bitcoin futures) and increased mainstream adoption as a payment system (e.g., cafes and internet merchants accepting bitcoin).

The emergence of new cryptocurrencies that are better at concealing a user's activity might also have contributed to the decline in the share of illegal activity in bitcoin as criminals migrate to these more opaque alternatives. We shed some light on this issue by examining how the estimated illegal activity in bitcoin was impacted by one of the major darknet marketplaces, Alphabay, beginning to accept an opaque alternative cryptocurrency, Monero, on its platform from August 22, 2016. Given that Alphabay's adoption of Monero is only expected to impact illegal activity, we isolate the impact using a difference-in-differences model of illegal and legal transaction activity in bitcoin during the eight weeks either side of August 22, 2016. The results (reported in Table D1 of Appendix D) show a significant decrease in the illegal activity in bitcoin after the event. Illegal users are estimated to make around 90 thousand fewer transactions in bitcoin per day after Alphabay's adoption of Monero (relative to legal users). This is an economically meaningful change given that illegal users made around 217 thousand transactions per day before the change. Figure D1 in Appendix D illustrates that the change in illegal activity occurs quickly around Alphabay's adoption of Monero. It also illustrates that the parallel trends assumption appears valid (further testing of this assumption could use a distributed lags approach).

While the effect of Alphabay's adoption of Monero appears quite large, there are three further considerations when interpreting the effect size. First, Alphabay is only one of many darknet sites, and it is likely other sites began accepting Monero at a similar time. Second, it is possible that the transactions that migrate to Monero are smaller, leaving a proportionally larger dollar value of darknet activity in bitcoin. Finally, it is possible that some of the darknet participants that initially switched to Monero returned to bitcoin some time later, potentially due to the continued widespread use of bitcoin in darknet marketplaces.

The time-series of legal and illegal activity levels show strong growth in both illegal and legal activity throughout the sample period, in particular since 2012. Interestingly, the strong growth in illegal activity precedes the strong growth in legal activity—by about three or four years. Thus, it seems illegal users were relatively early adopters of bitcoin as a payment system. Because of the rapid growth in the legal use of bitcoin in the final two years of the sample, the aggregate proportion of illegal bitcoin activity reported in the previous subsection understates the proportion that exists throughout most of the sample period. For example, for most of the period from 2009 to 2017, the estimated proportion of illegal users is closer to one-half than one-quarter (the aggregate estimate). The aggregate estimate is heavily influenced by the large number of legal users that enter in the last two years of the sample. Similarly, for much of the sample period, the estimated proportion of bitcoin transactions involved in illegal activity is between 60% and 80%, contrasting with the aggregate estimate of around 46%.

The most recent estimates of illegal activity (at the end of our sample in April 2017) suggest there are around 27 million illegal users of bitcoin. These users conduct around 37 million bitcoin transactions annually, valued at around $76 billion, and collectively hold around $7 billion in bitcoin.[39]

2.5.6 Robustness tests

We conduct a number of different robustness tests. Perhaps the most rigorous robustness test of an empirical model is to compare its results with results from a completely different model/approach that makes different assumptions and draws on different information. Throughout the chapter we put our two empirical models through this test. The two models, one based on a network cluster analysis algorithm and the other on a structural latent variables model drawing on observable characteristics, provide highly consistent results. The two models tend to agree, within a reasonable margin of error, on the overall levels of illegal activity, as well as the differences between legal and illegal users in terms of characteristics and network structure.

We also subject each of the models to specific tests that vary key assumptions or modelling choices. Table 2.5 reports the estimated amount of illegal activity for the most notable of these tests. For the SLM, we re-estimate the model using transaction volumes as the measure of interaction between users rather than transaction counts (*SLM Alternative 1*). We also consider a modification of the SLM

[39] For these estimates, we have halved the double-counted volumes so that the estimates can be interpreted as the volume/value of goods/services bought/sold by the illegal users.

algorithm in which we impose a constraint that does not allow the sample of "observed" illegal users to be moved to the legal community (*SLM Alternative 2*).

For the DCE model, one set of robustness tests involves examining the sensitivity to relaxing key exclusion restrictions. For example, in the baseline model, *Darknet Sites* (the number of operational darknet marketplaces at the time a user transacts) is included only as a determinant of illegal activity. As a robustness test (*DCE Alternative 1*), we include it in both equations, allowing it to also affect the probability of detection. *Darknet Sites* could affect detection if the existence of many darknet marketplaces is a catalyst for increased surveillance and enforcement by law enforcement authorities. We also test sensitivity to the key exclusion restriction in the detection equation by including *Pre-Silk-Road User* in both equations (*DCE Alternative 2*), thereby allowing it to also affect the probability of illegal activity. Finally, we relax the restriction that tumbling does not impact the probability of detection (*DCE Alternative 3*).

Table 2.5 shows that the estimated overall levels of illegal activity across the various activity measures are not overly sensitive to modifications of the baseline model, although there is some variation in individual estimates of illegal activity. For example, across the various alternative model specifications, the estimated proportion of illegal users varies from a minimum of 22.29% to a maximum of 29.12%. Similarly, the estimated characteristics of illegal users are not overly sensitive to these modifications (results not reported for conciseness). The Appendix E Table E1 reports the coefficient estimates of the three DCE models described above in which we relax key exclusion restrictions, showing that the key results are also not particularly sensitive to these modifications.

We also examine the robustness of the DCE model to the initial parameter values used in estimating the model. We initialize the model with different starting values (-1, 0, +1, and randomly drawn starting values), and find that our results are not sensitive to the choice of starting values, suggesting convergence to a global rather than local maximum of the likelihood function.

We re-estimate the standard errors used in confidence bounds around the estimated illegal activity and significance tests. Instead of the bootstrapped standard errors that we use in the main results, we instead compute standard errors using analytic expressions. We find that the analytic standard errors are considerably smaller than the bootstrapped standard errors. This finding suggests that using bootstrapped standard errors in the main results is the more conservative of the two approaches.

Finally, the characteristics of illegal users could change through time (for example, in response to seizures by law enforcement agencies), which could lead to model mis-specification. To examine this possibility, we estimate difference-in-differences models of how illegal user characteristics change after the Silk Road seizure relative to the changes in legal user characteristics. Controlling for the changes in legal user characteristics removes potentially confounding time-series variation that is due to the evolution of the bitcoin ecosystem.

Table 2.5
Robustness tests
This table reports robustness tests for the sensitivity of the overall estimated amount of illegal activity in bitcoin to variations in the specification of the underlying empirical models. The rows reflect estimates from different models. *SLM Baseline* and *DCE Baseline* are the SLM and DCE models used to produce the main results, and are included for comparison. The models labelled "Alternative" are variations on the corresponding baseline model. *SLM Alternative 1* is an SLM model that considers the transaction volume (in bitcoins) rather than the transaction count as a measure of trading activity when applying the network cluster analysis algorithm. *SLM Alternative 2* is a variation of the baseline SLM model in which observed (known) illegal users are constrained from leaving the illegal community. *DCE Alternative 1, 2*, and *3* are variations of the baseline DCE model in which exclusion restrictions for the instrumental variables are relaxed one at a time (these models correspond to Models 1-3 of Table E1 in Appendix E) respectively. The measures of group size and activity are: the number of users (*Users*), the number of transactions (*Transaction Count*), the average dollar value of bitcoin holdings (*Holding Value*), the number of bitcoin addresses (*Number Of Addresses*), and the dollar volume of transactions (*Volume*). Panel A reports the values of these measures for the two user groups, while Panel B expresses the measures for each group as a percentage of the total.

Group	Model	Users (Mil)	Transaction Count (Mil)	Holding Value ($Mil)	Number Of Addresses (Mil)	Volume ($Bil)
Panel A: Values						
Illegal	SLM Baseline	30.94	276.63	1,394.76	87.95	436.78
	SLM Alternative 1	28.95	270.69	1,418.42	85.10	400.29
	SLM Alternative 2	23.68	287.42	1,866.47	89.23	441.94
	DCE Baseline	24.68	282.70	1,523.87	86.35	422.05
	DCE Alternative 1	27.12	317.69	2,349.53	98.51	479.13
	DCE Alternative 2	24.59	276.56	1,474.28	82.47	420.22
	DCE Alternative 3	25.73	284.45	1,444.59	87.52	414.52
Panel B: Percentages						
Illegal	SLM Baseline	29.12%	45.67%	47.05%	39.67%	23.45%
	SLM Alternative 1	27.25%	44.69%	47.84%	38.38%	21.49%
	SLM Alternative 2	22.29%	47.45%	62.96%	40.25%	23.73%
	DCE Baseline	23.23%	46.67%	51.40%	38.95%	22.66%
	DCE Alternative 1	25.53%	52.45%	79.25%	44.43%	25.72%
	DCE Alternative 2	23.14%	45.66%	49.73%	37.20%	22.56%
	DCE Alternative 3	24.22%	46.96%	48.73%	39.47%	22.26%

Table F1 in Appendix F reports the difference-in-differences results using three different definitions of illegal users: illegal users identified by the SLM model, illegal users identified by the DCE model, and the directly observed sample of known illegal users that exist before and after the Silk Road seizure (corresponding to 2B and 2C in Table 2.3). The changes in most of the characteristics are not statistically distinguishable from zero. The statistically significant changes, using the directly observed illegal user group, suggest that after the Silk Road seizure illegal users tend to make fewer transactions, use smaller transactions, trade at a lower frequency, and hold smaller bitcoin balances. Such changes could impact the DCE model estimates and given the direction of the changes they could bias against the DCE identifying users as illegal. However, all of the estimated changes are relative to legal users and therefore some of the differences might be driven by the increase in speculative and mainstream interest in bitcoin in the later years of the sample. Most of the estimated changes in characteristics are small relative to the overall means of the characteristics. Therefore, there do not appear to be major changes in illegal user characteristics following the Silk Road seizure. Simpler models of pre-post changes in the illegal user characteristics provide qualitatively similar results, also suggesting there are no major changes in illegal user characteristics.[40]

2.6 Discussion

It is important to consider the differences between cryptocurrencies and cash. After all, cash is also largely anonymous (traceable only through serial numbers) and has therefore traditionally played an important role in facilitating crime and illegal trade (e.g., Rogoff, 2016). The key difference is that cryptocurrencies (similar to PayPal and credit cards) enable digital transactions and thus e-commerce. Arguably, the ability to make digital payments revolutionized retail and wholesale trade. Online shopping substantially impacted the structure of retailing, consumption patterns, choice, marketing, competition, and ultimately supply and demand. Until cryptocurrencies, such impacts were largely limited to legal goods and services due to the traceability of digital payments. Cryptocurrencies may have changed this, by combining the anonymity of cash with digitization, which enables efficient anonymous online and cross-border commerce. Cryptocurrencies therefore have the potential to cause an important structural shift in how the black market operates.

While the emergence of illegal darknet marketplaces illustrates that this shift may have commenced, it is not obvious to what extent the black market will adopt the opportunities for e-commerce and digital payments via cryptocurrencies. This is an important empirical question. Our findings illustrate the dynamics of this adoption process and suggest that eight years after the introduction of the first cryptocurrency, the black market has indeed adopted this form of electronic payment on a meaningful scale. Thus, our results suggest that cryptocurrencies are having a material impact on the way the black market for illegal goods and services operates.

Our findings have a number of further implications. Blockchain technology and the systems/protocols that can be implemented on a blockchain have the potential to revolutionize numerous industries. In shedding light on the dark side of cryptocurrencies, we hope this research will reduce some of the regulatory uncertainty about the negative consequences and risks of this innovation, facilitating more informed policy decisions that assess both the costs and benefits. In turn, we hope this enables these technologies to reach their potential. Second, the chapter contributes to our understanding the intrinsic value of bitcoin, highlighting that a significant component of its value as a payment system derives from its use in facilitating illegal trade. This has ethical implications for bitcoin as an

[40] Future work could estimate DCE models that use only the characteristics that do not change through time, examine in more detail how user characteristics respond to seizures, or estimate models using only the observed illegal users that were not part of the Silk Road seizure.

investment, as well as valuation implications. Third, the chapter moves the literature closer to understanding the welfare consequences of the growth in illegal online trade. A crucial piece of this puzzle is understanding whether illegal online trade simply reflects a migration of activity that would have otherwise occurred on the street, versus the alternative that by making illegal goods more accessible, convenient to buy, and less risky to buy due to anonymity, "black e-commerce" encourages growth in the aggregate black market. Our estimates contribute to understanding this issue, but further research is required to relate these estimates to trends in the offline black market to further understand the welfare consequences.

2.7 Conclusion

As an emerging FinTech innovation, cryptocurrencies and the blockchain technology on which they are based could revolutionize many aspects of the financial system, ranging from smart contracts to settlement, interbank transfers to venture capital funds, as well as applications beyond the financial system. Like many innovations, cryptocurrencies also have their dark side. We shed light on that dark side by quantifying and characterizing their use in illegal activity.

We find that illegal activity accounts for a sizable proportion of the users and trading activity in bitcoin, as well as an economically meaningful amount in dollar terms. For example, approximately one-quarter of all users and close to one-half of transactions are associated with illegal activity, equating to around 27 million market participants with illegal turnover of around $76 billion per year in recent times. Much of this illegal activity involves trading in darknet marketplaces. There are likely to be other forms of illegal activity such as evasion of capital controls that are not fully captured by our estimates.

Our results have a number of implications. First, by shedding light on the dark side of cryptocurrencies, we hope this research will reduce some of the regulatory uncertainty about the negative consequences and risks of this innovation, thereby allowing more informed policy decisions that weigh up the benefits and costs. In turn, we hope this contributes to these technologies reaching their full potential.

Second, the techniques developed in this chapter can be used in cryptocurrency surveillance in a number of ways. The methods can be applied going forward as new blocks are added to the blockchain, allowing authorities to keep their finger on the pulse of illegal activity and monitor its trends, its responses to regulatory interventions, and how its characteristics change through time. Such information could help make more effective use of scarce regulatory and enforcement resources. The methods can also be used to identify individuals of strategic importance in illegal networks.

Third, our chapter suggests that a significant component of the intrinsic value of bitcoin as a payment system derives from its use in facilitating illegal trade. This has ethical implications for those that view bitcoin as an investment, as well as valuation implications. For example, changes in the demand to use bitcoin in illegal trade (e.g., due to law enforcement crackdowns or increased adoption of more opaque cryptocurrencies in illegal trade) are likely to impact its fundamental value.

Finally, our chapter moves the literature closer to answering the important question of the welfare consequences of the growth in illegal online trade. A crucial piece of this puzzle is understanding whether online illegal trade simply reflects migration of activity that would have otherwise occurred on the street, versus the alternative that by making illegal goods more accessible, convenient to buy, and less risky due to anonymity, the move online encourages growth in the aggregate black market. Our estimates of the amount of illegal trade facilitated via bitcoin through time contribute to understanding this issue, but further research is required to relate these estimates to trends in the offline black market.

The DCE approach developed in this chapter is used in Chapter 3 to characterize the illegal activity in bitcoin and the determinants of its detection. Chapter 3 complements this chapter by providing deeper insights into how law enforcement may use their resources efficiently when combatting illegal activity in cryptocurrencies.

Chapter 3

Characteristics of illegal activity in bitcoin

3.1 Introduction

The detection controlled estimation (DCE) model developed in Chapter 2 quantifies unobserved illegal activity and its development over time. However, DCE has applications beyond measuring the scale and development of illegal bitcoin activity—the model can estimate the characteristics of illegal bitcoin users and the determinants of their detection. In this chapter, we use the DCE model to estimate the effects and magnitudes of illegal user characteristics and detection variables. We also analyze the topology of the illegal bitcoin network with a suite of network metrics. These results are particularly useful to market agents such as regulators, whose mandate is to reduce criminal activity in the markets they oversee. They provide insights into how illegal user activity differs from legal activity, thereby enhancing the ability of market agencies to identify misconduct.

The rapid growth in cryptocurrencies and the anonymity that they provide users has created considerable regulatory challenges. An application for a $100 million cryptocurrency Exchange Traded Fund (ETF) was rejected by the US Securities and Exchange Commission (SEC) in March 2017 (and several more rejected in 2018) amid concerns including the lack of regulation. The Chinese government banned residents from trading cryptocurrencies and made initial coin offerings (ICOs) illegal in September 2017. Central bank heads, such as the Bank of England's Mark Carney, have publicly expressed concerns about cryptocurrencies. While cryptocurrencies have many potential benefits including faster and more efficient settlement of payments, regulatory concerns center around their use in illegal trade (drugs, hacks and thefts, illegal pornography, even murder-for-hire), potential to fund terrorism, launder money, and avoid capital controls. There is little doubt that by providing a digital and anonymous payment mechanism, cryptocurrencies such as bitcoin have facilitated the growth of online "darknet" marketplaces in which illegal goods and services are traded. The recent FBI seizure of over $4 million worth of bitcoin from one such marketplace, the "Silk Road," provides some idea of the scale of the problem faced by regulators.

Given the incredible growth of illegal online activity facilitated by cryptocurrencies, it is now especially important for regulators and law enforcement to gain an understanding into the characteristics of illegal cryptocurrency activity. Insights into the characteristics illegal users and their network may inform law enforcement on how to allocate resources efficiently by pinpointing key individuals (or hubs) in the network. Commercial businesses such a cryptocurrency exchanges may also use known characteristics of illegal users to flag unwanted market participants. This chapter aims to rectify the situation by estimating the characteristics of the illegal users, the determinants of their detection, and the topology of their illegal network.

The analysis in this chapter shows that bitcoin users that are involved in illegal activity differ from other users in several characteristics. Illegal users tend to transact more, but in smaller transactions. They are also more likely to repeatedly transact with a given counterparty. These differences in transactional characteristics are generally consistent with the notion that while illegal users predominantly (or solely) use bitcoin as a payment system to facilitate trade in illegal goods/services, some legal users treat bitcoin as an investment or speculative asset. Despite transacting more, illegal users tend to hold less bitcoin, consistent with them facing risks of having bitcoin holdings seized by authorities.

We find several other robust predictors of involvement in illegal activity. A user is more likely to be involved in illegal activity if they trade when there are more darknet marketplaces in operation, lower combined market value of shadow coins, less mainstream interest in bitcoin as measured by Google search intensity, and immediately following darknet marketplaces seizures or scams. A user is also more likely to be involved in illegal activity if they use “tumbling” and/or “wash trades”—two trading techniques that can help conceal one’s activity.

We find that the network of bitcoin transactions between illegal users is three to four times denser than the legal user network, with users much more connected with one another through transactions. The higher density is consistent with illegal users transacting more and using bitcoin primarily as a payment system for buying/selling goods.

The next section develops hypothesis on the characteristics of illegal users and their network community. Section 3.3 tests these hypotheses using univariate statistics, DCE model coefficients, and network metrics. In Section 3.4, we discuss the implications of chapters 2 and 3 as well as their contribution to the literature. Section 3.5 concludes.

3.2 Hypotheses

In this section, we develop hypotheses about the characteristics of illegal users, the determinants of their detection, and the topology of the illegal user network. To create these hypotheses, we rely on the literature on crime in cash-based black markets and the more recent literature on cryptocurrencies.

3.2.1 Characteristics of illegal activity in cryptocurrencies

As a means of payment, bitcoin has challenges; among the issues that the cryptocurrency faces is the enormous amounts of electricity needed to maintain the network, the slow transaction processing, and the high transaction fees (De Vries, 2018). For example, as of 2021, bitcoin uses almost 120 TWh per year (more than half of Denmark’s yearly energy consumption), while a transaction can cost anywhere between $10 and $60 and takes around ten minutes to complete.[41] The blockchain community responds with a new blockchain protocol called "Proof-of-Stake" that consumes less energy by limiting individuals' mining power to their cryptocurrency holdings (their "Stake") and avoids a computational arms race (Saleh, 2021). While Ethereum is in the process of switching to "Proof-of-Stake", the bitcoin blockchain still uses a "Proof-of-Work" protocol that requires enormous amounts of energy. The type of user who choses bitcoin as a means of transaction accept these downsides because of the anonymity it provides—generally, these users tend to be the illegal users. The majority of the remaining users are likely (legal) investors, who invest and hold the currency, which involves much fewer transactions. Illegal transactions may also be smaller in amount than legal investment transactions because illegal users buy small quantities of illegal items to limit the cost of their losing purchased items to law enforcement.[42] This leads to the following hypothesis:

H1. Illegal users use bitcoin as a means of transaction rather than an investment, so they transact in small, frequent transactions.

Illegal users who view bitcoin as an essential part of their business model (and not an investment), likely limit their holdings to what is needed to purchase and sell products and services on the darknet. Extending their holdings beyond the necessary amounts increases the cost of capital because the funds could be used on alternative investments. Second, illegal users also face the risk of losing their bitcoin, either in a seizure or scam, so they attempt to limit their potential losses by keeping their bitcoin balances low. Users associated with legal activity in bitcoin likely use the cryptocurrency as an investment and are not concerned with darknet scams or seizures—therefore, their holdings are larger.

H2. Illegal bitcoin users hold fewer bitcoin to limit their losses in bitcoin seizures and to lower the opportunity cost of capital.

The number of darknet sites on the darkweb is limited to around one hundred,[43] which vary depending on popularity, product offering, and accepted cryptocurrency. Each market has its own address (or "hot wallet") that customers send bitcoin to when purchasing items on the website. Second, the interface on darknet marketplaces allows for comments and feedback on products sold, which likely instills a sense of community and loyalty between consumers and sellers. The community also must

[41] For an estimate on bitcoin energy consumption see May 5, 2021 Harvard Business Review article "How much energy does bitcoin actually consume" and see https://ourworldindata.org/energy/country/denmark for Danish energy consumption.

[42] Law enforcement work with postal services to examine suspicious packages sent via mail. For more information on law enforcement strategies see Kruithof et al. (2016).

[43] See May 6, 2019 *CYBERSCOOP* article "How may dark web markets actually exist? About 100".

rely much more on trust than legal communities because there are no laws or contracts protecting the market participants (see, e.g., Kinsella, 2006). On darknet sites, trust is so important that sellers with a good reputation can charge a premium for their products (Hardy and Norgaard, 2016). The size and "tightly-knit" nature of the illegal community implies that illegal users often repeatedly transact with the same user. We would therefore expect the following:

H3. Illegal activity is concentrated to a small community, and illegal bitcoin users transact repeatedly with the same counterparties.

Most of the early adopters of bitcoin were the illegal users—the first darknet marketplace (the "Silk Road") started its operations already in 2011, only two years after the inception of bitcoin. Bitcoin provided early vendors with a lucrative opportunity because the payment system expanded their previously cash-based business to internet commerce while keeping their business activity anonymous. Therefore, the longer a user has been active in the bitcoin network, the more likely they are to be involved in illegal activity:

H4. Illegal users have been active in the bitcoin network for longer than legal users because they were among the first adopters of bitcoin.

An increased number of darknet sites increases competition among the markets and lowers the transactions costs for market participants. In financial markets, competition among exchanges (market fragmentation) increases liquidity, decreases transaction costs, and encourages market participation (O'Hara and Ye, 2011). Similarly, illegal activity in cryptocurrencies likely increases when there are more darknet markets that, through competition, drive down the transaction costs. We therefore hypothesize:

H5. Illegal bitcoin users become more active when more darknet sites are available.

As with other types of illegal activity, illegal bitcoin users are less likely to use bitcoin for illegal purposes if the costs outweigh the rewards. This is similar to the rational expectation models in the tax evasion literature, where tax evaders base their decision to evade on the benefits (tax savings) and costs probability of detection and penalties if detected (see, e.g., Allingham and Sandmo, 1972; Yitzhaki, 1974). In the bitcoin network, illegal users can lower the probability of detection by using services that further obfuscate their transactions—so-called "tumbling services." Tumbling services (or "mixing services") allow the user to mix their bitcoin with other bitcoin to obscure their original source.

We assume bitcoin users engaged in illegal activity use tumbling services more frequently than legal users to lower the probability of detection, so we hypothesize the following:

H6. Illegal bitcoin users frequently use tumbling services to conceal their activity.

Darknet sites frequently shut down, either because law enforcement agencies seize the market or because the owners scam market participants by closing the site and stealing the cryptocurrencies held in escrow.[44] For example, in 2013 the US government seized the infamous darknet site, the "Silk Road," along with 69 thousand bitcoin (worth around $3 billion today). Similarly, in 2016, the administrator of the "Evolution" darknet market, shut down the site and stole all the bitcoin held in escrow. After such events, illegal users who were customers in the closed darknet site will likely fill their demand by reallocating their holdings to competing markets. They may also take precautionary measures to ensure the safety of their bitcoin holdings from law enforcement by using tumbling services and wash trades. We would thus expect the following:

H7. A darknet market seizure or scam will likely increase illegal user trading activity as they relocate their holdings to competing sites or take precautionary measures to protect their funds.

The increased demand for anonymous payments has prompted blockchain developers to create new cryptocurrencies that rival (and supersede) bitcoin in their ability to protect users from detection. Some of the features that set these privacy-enhancing cryptocurrencies apart from bitcoin include creating a new receiving address (called "stealth address") for every transaction[45], "ring signatures," which obfuscate the sender of transactions by mixing their signature with other non-genuine senders, or encrypting addresses and transacted amounts through "zero-knowledge proofs." The demand for bitcoin as a means of illegal transaction will likely drop as darknet markets and the illegal cryptocurrency community as a whole adopt these alternative means of covert payment, hence producing the following hypothesis:

H8. The increased number alternative cryptocurrencies, including shadow coins, specifically designed to conceal user activity lowers illegal activity in bitcoin.

A given user is less likely to be engaged in illegal activity as the mainstream popularity of bitcoin increases. Since the first bitcoin exchanges in 2010, bitcoin has risen from a fringe asset trading

[44] For example, the "silk road" darknet seizure in 2013 or the "Evolution" darknet market scam.

[45] A cryptocurrency "address" resembles a traditional bank account in that it contains (cryptocurrency) funds. Creating a new address for every transaction disguises the user's aggregate activity because the owner of individual addresses is unobservable on the blockchain.

for less than $1 dollar to almost $40 thousand today, partly because of the large speculative investments in the cryptocurrency.[46] All else equal, increasing the number of (legal) bitcoin investors in the bitcoin network lowers the proportion of illegal users and decreases the likelihood that a randomly picked user is illegal. Thus, we hypothesize the following:

H9. Increased mainstream attention measured by bitcoin market capitalization lowers the probability of a user being illegal.

3.2.2 Determinants of illegal user detection

The first darknet market seizure (the "Silk Road") occurred in October of 2013, sparking increased attention from law enforcement around the world on bitcoin activity in darknet markets. The Silk Road seizure (naturally) only caught illegal users who were active before October 2012 and, therefore, users who became active after the event have a lower probability of detection. This leads to the following hypothesis:

H10. Detection is more likely for illegal users who traded bitcoin before the first darknet site (the "Silk Road") seizure.

There have been many seizures since the Silk Road, including the recent (and largest) seizure of "DarkMarket" in January 2021, which hosted an astounding 500,000 users and facilitated 320,000 illegal transactions. The total duration that illegal users are active in the network increases their exposure to darknet seizures and their likelihood of detection, so expect the following:

H11. The likelihood of illegal user detection increases with the duration of their activity in the bitcoin network.

Law enforcement have limited resources, so they focus on the most egregious misconduct so that their efforts have the largest crime-reducing effect, creating headlines that deter other potential violators. The likelihood that a user is detected increases with the intensity of their activity in the illegal community; repeated small transactions to the same counterparties (such as a darknet market) characterizes illegal activity and increases the likelihood of detection.

[46] See https://www.coindesk.com/ for a timeseries of btc/usd quotes.

H12. Detection is more likely for illegal users who transact repeatedly with the same counterparty and use bitcoin as a means of transaction (as opposed to investment) by transacting small amounts frequently.

3.2.3 The illegal user network

In the illegal network, vendors sell their products to many different customers, and customers may buy from several vendors to meet their demand. Vendors may also act as customers in some transactions when purchasing products for their personal consumption (and vice versa). Therefore, the illegal network is much more "connected" than the legal network, where investors buy and hold bitcoin and rarely transact with other legal users in their network.

H13. The illegal network is very connected in that illegal users transact with many other illegal users and occasionally both send and receive bitcoin from the same user.

Kinsella (2006) views conventional arms trade as a network rather than a marketplace. A social network is usually heterogeneous in the number of links each participant has. For example, black markets for illegal firearms have a high level of centrality in that some market participants interact with considerably more counterparties than others (Kinsella, 2006)—the participants with many interactions are usually sellers and those with a few are typically buyers. There is likely also a high level of heterogeneity in the interactions of our illegal bitcoin sample because it consists of a few vendors (who interact with many buyers) and many buyers (who interact with a few vendors).

H14. The illegal users are more heterogeneous in the number of other users they interact with than legal users.

3.3 Illegal user characteristics

In this section, we characterize the legal and illegal users as well as the topology of their networks. First, we use univariate statistics to test the differences in means between legal and illegal users. Second, we re-estimate the detection controlled estimation model from Chapter 2 to observe the characteristics of illegal users and the determinants of their detection. Third, we use network characteristics to measure the differences between the legal and illegal networks.

In Chapter 2, we used two models, the SLM model and the DCE model, to classify users as either legal or illegal. The midpoint of the SLM and DCE estimates from the chapter indicate that around one-quarter (26.17%) of all users in the bitcoin network are illegal and half of all transactions (46.17%) equal to almost \$430 billion are associated with illegal activity. Both models use a sample of hand-collected "observed" illegal users to estimate the true (unobserved) illegal proportion. The "observed" sample consists of users caught by the authorities, users that directly transact with known darknet markets, and users that list their addresses on darknet forums.

The SLM model uses the transactional activity between "observed" illegal users and "other" users (users that are not observed as illegal) to estimate the true proportion of illegal bitcoin activity. The model classifies users as "legal" if they predominately transact with other "legal" users and "illegal" if they predominately transact with other "illegal" users.[47]

The DCE model also uses the "observed" illegal sample when estimating illegal activity but estimates two equations simultaneously—a violation equation (equations 2.1-2.2) and detection equation (equations 2.3-2.4). The violation equation uses a set of illegal user characteristics to model whether the user is predominantly engaged in illegal or legal activity. The detection equation, models whether or not a user is detected using a set of variables that affect the probability of detection. A user is detected in the second equation, if they are in the "observed" illegal sample.

3.3.1 What are the characteristics of the illegal users?

We assess the differences between legal and illegal user characteristics in two ways: univariate statistics that compare observed or estimated illegal users with their legal counterparts, and multivariate tests exploiting the coefficients of the estimated DCE model. The univariate tests compare the means of illegal and legal users. The illegal classification comes from our hand-collected sample of illegal participants (the "*Observed*" category), the SLM model estimates (the "*SLM*" category), and the DCE model (the "*DCE*" category). Finally, we examine the characteristics of the illegal users and the determinants of illegal user detection with two variants of the detection controlled estimation model—the first DCE variant ("*Model 1*"), is the baseline model used for the main results in Chapter 2, and the second variant (*Model 2*"), adds two additional controls.

Starting with a univariate difference in means, Table 3.1 compares the characteristics of the sample of "observed" illegal users with the characteristics of other users. The "other users" are not all legal users—they contain a mix of legal users and undetected illegal users. Therefore, the table also compares the characteristics of users classified by the SLM and DCE models as being involved in illegal activity with those of the users classified as legal. Interestingly, despite being based on completely different assumptions, the SLM and DCE models generally agree on how the characteristics of legal users differ from illegal users.

Consistent with hypothesis *H1*, the SLM and DCE models agree that illegal users tend to transact more (have a two to three times higher *Transaction Count*), but use smaller sized transactions (about half the average size of legal transactions). This result could be a reflection of illegal users predominantly using bitcoin to buy and sell goods and services, whereas some legal users also use

[47] Section 2.5.1 explains the SLM model in detail.

bitcoin for investment and speculation.[48] With the average size of an illegal transaction being around $3,000, bitcoin transaction fees even at their peak of around $150 (see Basu et al., 2019) are small relative to the average illegal transaction.

[48] While the result could also reflect illegal users engaging in techniques to conceal their trading, this is less likely to be an explanation because a similar result holds in multivariate (DCE) tests that control for tumbling and wash trades.

Table 3.1
Differences in characteristics between illegal and legal users

This table reports differences in mean characteristics for illegal vs legal bitcoin users. The first three columns (“*Observed*”) compare observed illegal users (those identified through seizures, darknet marketplaces, and darknet forums) and other users (including both legal and undetected illegal users). The second three columns (“*SLM*”) compare illegal and legal users, classified by a network cluster analysis algorithm (SLM). The final three columns (“*DCE*”) compare illegal and legal users, classified by a detection-controlled estimation (DCE) model. The characteristics are as follows. *Transaction Count* is the total number of bitcoin transactions involving the given user. *Transaction Size* (in USD) is the user’s average transaction value. *Transaction Frequency* is the average rate at which the user transacts between their first and last transactions, annualized to transactions per year. *Counterparties* is the total number of other users with which the given user has transacted. *Holding Value* is the average value of the user’s bitcoin holdings (in USD), where holdings are measured after each transaction. *Concentration* takes values between zero and one, with higher values indicating a tendency to repeatedly trade with a smaller number of counterparties. *Existence Time* is the number of months between the date of the user’s first and last transaction. *Darknet Sites* is the average number of operational darknet sites at the time of each of the user’s transactions. *Tumbling* is the percentage of the user’s transactions that attempt to obscure the user’s holdings (wash or tumbling trades). *Darknet Shock Volume* is the percentage of the user’s total dollar volume that is transacted during the week after marketplace seizures or “exit scams”. *Bitcoin Hype* is a measure of the intensity of Google searches for the term “bitcoin” around the time of the user’s trades. *Pre-Silk-Road User* is a dummy variable taking the value one if the user’s first bitcoin transaction is before the seizure of the Silk Road on October 2013. *Bitcoin Market Cap*, *Shadow Coins*, and *Alt Coins* are transaction-weighted average log market capitalizations of bitcoin, major opaque cryptocurrencies, and non-privacy cryptocurrencies excluding bitcoin, respectively, at the time of each user’s transactions. The significance of the difference in means is computed with t-tests. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels respectively.

Variable	Observed			SLM			DCE		
	Other	Illegal	Difference	Legal	Illegal	Difference	Legal	Illegal	Difference
	(1)	(2)	(2-1)	(1)	(2)	(2-1)	(1)	(2)	(2-1)
Transaction Count	4.09	31.51	27.42***	4.37	8.94	4.57***	3.96	11.46	7.50***
Transaction Size	5,346.87	2,969.38	-2,377.49***	6,225.51	2,729.66	-3,495.85***	5,791.25	3,278.30	-2,512.95***
Transaction Frequency	28.91	45.46	16.54***	29.77	30.16	0.39**	28.50	34.45	5.95***
Counterparties	3.53	14.61	11.08***	3.77	5.18	1.42***	3.57	6.19	2.62***
Holding Value	4,021.77	3,207.06	-814.71***	4,625.45	2,388.31	-2,237.14***	4,359.86	2,698.71	-1,661.16***
Concentration	0.09	0.20	0.11***	0.08	0.13	0.05***	0.09	0.12	0.04***
Existence Time	6.19	13.44	7.26***	5.91	8.31	2.40***	6.17	8.08	1.91***
Darknet Sites	17.17	16.67	-0.50***	17.13	17.17	0.04***	16.87	18.04	1.18***
Tumbling	0.40	1.18	0.78***	0.37	0.64	0.27***	0.31	0.89	0.58***
Darknet Shock Volume	15.84	27.25	11.40***	14.51	21.39	6.88***	10.57	36.14	25.56***
Bitcoin Hype	28.74	21.16	-7.58***	29.67	24.95	-4.72***	30.99	19.38	-11.60***
Pre-Silk-Road User	0.06	0.22	0.16***	0.06	0.12	0.07***	0.03	0.22	0.19***
Bitcoin Market Cap	9.85	9.45	-0.40***	9.88	9.68	-0.21***	9.96	9.36	-0.60***
Shadow Coins	7.18	5.34	-1.84***	7.30	6.51	-0.79***	7.67	5.11	-2.56***
Alt Coins	8.75	7.49	-1.26***	8.86	8.24	-0.62***	9.04	7.47	-1.57***

The models also agree that illegal users tend to hold less bitcoin (measured in dollar value) than legal users; their average *Holding Value* is about half that of legal users. This characteristic is consistent with the previous conjecture from *H2*—legal users might tend to hold larger bitcoin balances because some use bitcoin for investment/speculation purposes, whereas for an illegal user that buys/sells illegal goods and services using bitcoin, holding a large balance is costly due to (i) opportunity costs of capital, and (ii) risks associated with having holdings seized by authorities. For these reasons, illegal users are likely to prefer holding less bitcoin and this tendency is supported by the data.

Illegal users tend to have more counterparties in total, reflecting their larger number of transactions, but tend to have a higher counterparty concentration (consistent with hypothesis *H3*). This suggests that illegal users are more likely to repeatedly transact with a given counterparty. This characteristic might be a reflection of illegal users repeatedly transacting with a given illegal darknet marketplace or other illegal user once trust is established from a successful initial exchange. Illegal users have a longer *Existence Time* (time between their first and last transactions in bitcoin), consistent with *H4* and our observations from the time-series that illegal users tend to become involved in bitcoin earlier than legal users. Similarly, the differences in means also show that there is a higher proportion of Pre-Silk-Road users among the illegal users than the legal users (as indicated by the variable *Pre-Silk-Road User*).

The more specific indicators of illegal activity also show significant differences between the two groups. Consistent with *H5, i*llegal users tend to be more active during periods in which there are many illegal darknet marketplaces operating (a higher mean for the variable *Darknet Sites*). They make greater use of tumbling and wash trades to conceal their activity (two to three times more *Tumbling*), which provides support for *H6*. On average, a larger proportion of illegal volume, compared to legal volume, is transacted immediately following shocks to darknet marketplaces (*Darknet Shock Volume*). This finding is consistent with *H7* with anecdotal evidence that illegal users turn to alternative marketplaces in response to darknet marketplace seizures or scams.

Interestingly, illegal users are more likely to transact in bitcoin when there is lower combined market value of "shadow coins" consistent with such coins serving as alternatives to bitcoin in illegal transactions (hypothesis *H8*). This result matches anecdotal accounts of shadow coins attracting attention from the illegal community for their increased privacy and recent examples of hackers demanding ransom payments in shadow coins rather than bitcoin. The result also supports the evidence that illegal activity in bitcoin decreased after a major darknet marketplace, Alphabay, adopted one of the major shadow coins, Monero, as a form of payment in August 2016.

Another interesting result is that there tends to be relatively fewer illegal users when there is less *Bitcoin Hype*, measured by the Google search intensity for "bitcoin" (supporting *H9*). It therefore appears that Google searches for "bitcoin" are associated with mainstream (legal) adoption of bitcoin for payments, and/or speculative/investment interest in bitcoin. Similarly, there are relatively fewer illegal users when bitcoin market capitalization is higher and when other cryptocurrencies, "*Alt-Coins*" (excluding the opaque shadow coins), have higher value. This finding suggests that high valuations of bitcoin and other non-shadow cryptocurrencies correspond to periods of increased legal interest in cryptocurrencies.

In summary, the comparison of transactional characteristics (number and size of transactions, holdings, and counterparties) is consistent with the notion that illegal users predominantly use bitcoin for payments, whereas legal users are more likely to treat bitcoin as an investment asset. Furthermore, legal and illegal users differ with respect to when they are most active in bitcoin, with illegal users being most active when there are more darknet marketplaces, less bitcoin hype, lower bitcoin and other non-shadow cryptocurrency market capitalizations, and immediately following shocks to darknet marketplaces. The differences in characteristics for the instrumental variables are consistent with the hypothesized differences, lending support to their use as instruments.

The DCE model coefficients reported in Table 3.2 provide multivariate tests of how the characteristics relate to the likelihood that a user is involved in illegal activity. The results confirm most of the observations made in the simple comparison of means. The effects of all of the instrumental variables are consistent with their hypothesized effects. A user is more likely to be involved in illegal activity if they trade when: (i) there are many darknet marketplaces in operation, (ii) "shadow coins" such as Monero are not widespread (low market values), (iii) the market value of bitcoin is low, and (iv) darknet marketplaces have recently experienced seizures or scams. A user is also more likely to be involved in illegal activity if they use tumbling and/or wash trades, transact frequently in small sized transactions, and tend to repeatedly transact with a given counterparty. The value of other non-privacy cryptocurrencies (*Alt Coins*) at the time a user transacts is not statistically significant after controlling for the other variables, despite *Alt Coins* correlating with the likelihood of illegal activity in univariate tests. The results suggest that *Bitcoin Market Cap* is more closely related to the amount of mainstream and speculative interest in bitcoin and therefore *Alt Coins* is not a significant predictor of illegal activity after controlling for the value of bitcoin.

Table 3.2
DCE model estimates

This table reports the coefficient estimates and marginal effects of two detection-controlled estimation (DCE) models. Both models use the two-equation structure given by equations (2.1-2.4) of Chapter 2. Model 1 is the baseline model used for the main results in the chapter. Model 2 includes additional control variables. I() is the probability that a given user is predominantly using bitcoin for illegal activity. D() is the conditional probability of detection. Variables are defined in Table 2.1 in Chapter 2. Numbers not in brackets are the coefficient estimates. Numbers in brackets are the marginal effects (partial derivatives of the corresponding probability with respect to each of the variables, reported as a fraction of the estimated corresponding probability). Pseudo R^2 is McFadden's likelihood ratio index (one minus the ratio of the log-likelihood with all predictors and the log-likelihood with intercepts only). Significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively, using bootstrapped standard errors.

	Model 1		Model 2	
Variable	I()	D()	I()	D()
Intercept	-1.147***	0.265***	-1.054***	0.066
	(-0.755)	(0.126)	(-0.677)	(0.033)
Darknet Sites	1.005***		1.076***	
	(0.661)		(0.691)	
Tumbling	0.085***		0.103***	
	(0.056)		(0.066)	
Bitcoin Market Cap	-1.608***		-1.690***	
	(-1.059)		(-1.085)	
Shadow Coins	-0.649***		-0.679***	
	(-0.428)		(-0.436)	
Alt Coins	0.591		0.615	
	(0.389)		(0.395)	
Darknet Shock Volume	0.445***		0.496***	
	(0.293)		(0.319)	
Pre-Silk-Road User		0.430***		0.430**
		(0.204)		(0.213)
Transaction Frequency	0.438***	0.477***	0.230	0.474
	(0.288)	(0.227)	(0.148)	(0.235)
Transaction Size	0.005	-0.171***	-1.574***	-0.443**
	(0.003)	(-0.081)	(-1.011)	(-0.220)
Concentration	0.293***	0.542***	0.268***	0.500***
	(0.193)	(0.258)	(0.172)	(0.248)
Existence Time	0.098	1.744***	-0.058	1.405
	(0.064)	(0.829)	(-0.037)	(0.697)
Holding Value			3.602***	-0.537
			(2.312)	(-0.266)
Transaction Count			7.900	-0.593
			(5.071)	(-0.294)
Pseudo R^2	21.92%		22.08%	

The marginal effects in Table 3.2, reported in parentheses below the coefficient estimates, provide a sense of the magnitudes of the effects and their relative importance.[49] For example, the marginal effects indicate that a one standard deviation increase in the number of illegal darknet marketplaces at the time a user transacts in bitcoin increases the probability of that user being involved in illegal activity by a factor of 0.661, or 66.1% of what their probability would otherwise be.[50] The magnitudes generally show that most of the determinants of involvement in illegal activity and determinants of the detection probability are economically meaningful. In particular, the instrumental variables *Darknet Sites*, *Shadow Coins*, *Bitcoin Market Cap,* and *Darknet Shock Volume* all have strong relations with the probability that a user is involved in illegal activity.

The DCE model also sheds light on the determinants of the likelihood that an illegal user is "detected" by either of our three approaches and the results support hypotheses *H10*, *H11*, and *H12*. The main instrument, *Pre-Silk-Road User* has a strong relation with detection, indicating that illegal users that commence transacting in bitcoin prior to the first darknet marketplace seizure in October 2013 have a higher probability of being detected. Similarly, those users that transact in bitcoin for a longer period of time (higher *Existence Time*), trade more frequently (higher *Transaction Frequency*), or tend to trade repeatedly with a given counterparty such as a darknet marketplace (higher *Concentration*), have a significantly higher detection probability.

Model 2 in Table 3.2 adds further control variables, including *Holding Value* and *Transaction Count*, and finds that the main results do not change much in response to additional control variables. A risk of adding too many transactional control variables is co-linearity between such variables. In unreported results, we also find that the main results are robust to including a measure of bitcoin volatility. Somewhat unexpectedly, bitcoin volatility around the time a user transacts in bitcoin has a positive association with the likelihood that user is involved in illegal activity, all else equal.

3.3.2 What are the characteristics of the illegal user network?

Exploiting the fact that the bitcoin blockchain provides us with a complete record of every transaction between every pair of counterparties, we briefly explore how the trade network of illegal users differs from that of legal users. Our approach is to compute a few descriptive network metrics that capture different aspects of network topology and structure for each of the two groups or "communities"

[49] To make the comparisons and interpretation easier, before estimating the DCE models, we standardize all variables to have mean zero and standard deviation of one. We also log transform the right skewed variables (*Transaction Frequency*, *Size*, and *Count*, and *Holding Value*) and winsorize the variables at +/- three standard deviations to reduce the influence of extreme values.

[50] As an example of how to interpret the marginal effect of 0.661, if a user's illegal probability is say 20%, the predicted effect of a one standard deviation increase in *Darknet Sites*, holding all else constant, is to increase the user's probability to 20% × 1.661 = 33.2%, an increase of 66.1% of what their probability would otherwise be.

separately and then compare the values between the two communities. In mapping the networks, users form the "nodes", and transactions between users form the "edges" or "links" between nodes.

Table 3.3
Network characteristics of legal and illegal bitcoin user networks
This table reports metrics that characterize the trade networks of estimated legal and illegal bitcoin users. In the columns labelled "SLM" user classifications into legal and illegal communities are based on a network cluster analysis algorithm (SLM) and in the columns labelled "DCE" the classifications are from a detection-controlled estimation (DCE) model. *Density* takes the range [0,1] and indicates how highly connected the users are within a community (versus how sparse the connections are between users); it is the actual number of links between users within the given community (a "link" between two users means that they have transacted with one another) divided by the total potential number of links. *Reciprocity* takes the range [0,1] and indicates the tendency for users to engage in two-way interactions (both sending and receiving bitcoin to and from one another); it is the number of two-way links between users within the given community divided by the total number of links within the given community (two-way and one-way). *Entropy* measures the amount of heterogeneity among users in their number of links. It takes its minimum value of zero when all users have the same number of links (same degree).

Metric	SLM		DCE	
	Legal	Illegal	Legal	Illegal
Density (10^{-6})	0.04	0.13	0.04	0.17
Reciprocity	0.01	0.03	0.01	0.03
Entropy	1.50	1.75	1.53	1.73

Table 3.3 reports the results. The first metric, *Density*, takes the range [0,1] and indicates how highly connected the users are within a community (versus how sparse the connections are between users); it is the actual number of links between users within the given community (a "link" between two users means that they have transacted with one another) divided by the total potential number of links. It shows that the illegal trade network is three to four times denser in the sense that users are much more connected to one another through transactions. This observation is consistent with hypothesis *H13* and the fact that illegal users tend to transact more than legal users. It is also consistent with the notion that in the illegal community, bitcoin's dominant role is likely that of a payment system in buying/selling goods, whereas in the legal community, bitcoin is also used as an investment or for speculation.

Reciprocity takes the range [0,1] and indicates the tendency for users to engage in two-way interactions; it is the number of two-way links between users within the given community (a two-way link is when two users send and receive bitcoin to and from one another) divided by the total number of links within the given community (two-way and one-way). While *Reciprocity* is higher among illegal users than it is among legal users (supporting hypothesis *H13*), it is generally very low in both communities (1% among legal users and 3% among illegal users). Thus, interactions between bitcoin users are generally only one-way interactions with one counterparty receiving bitcoin from the other but not vice versa.

Entropy measures the amount of heterogeneity among users in their number of links to other members of the community. It takes its minimum value of zero when all users have the same number of links (same degree).[51] The results support *H14* and suggest that illegal users are a more heterogeneous group in terms of the number of links each user has with other members of the community. A driver of that heterogeneity could be that the illegal community at one end of the spectrum has darknet marketplaces that have hundreds of thousands of links to vendors and buyers, and at the other end has individual customers of a single marketplace, potential with only the one link.

A concluding observation is that both the SLM and DCE models provide a consistent picture of how legal and illegal users differ, this time in the context of their trade networks. Again, this suggests that the two different models tend to agree about the nature of the illegal activity in bitcoin.

3.4 Discussion

3.4.1 Implications

Blockchain technology and the systems/protocols that can be implemented on a blockchain have the potential to revolutionize numerous industries. Possible benefits to securities markets include reducing equities settlement times and costs (Malinova and Park, 2016; Khapko and Zoican, 2017), increasing ownership transparency leading to improved governance (Yermack, 2017), and providing a payments system with the network externality benefits of a monopoly but the cost discipline imposed by free market competition (Huberman et al., 2017). The technology has even broader applications beyond securities markets, from national land registries, to tracking the provenance of diamonds, decentralized decision making, peer-to-peer insurance, prediction markets, online voting, distributed cloud storage, internet domain name management, conveyancing, medical record management, supply chain, auditing, and many more.[52]

This technology, however, is encountering considerable resistance, especially from regulators. Regulators are cautious due to their limited ability to regulate cryptocurrencies and the many potential but poorly understood risks associated with these innovations. The negative exposure generated by anecdotal accounts and salient examples of illegal activity no doubt contributes to regulatory concerns and risks stunting the adoption of blockchain technology, limiting its realized benefits. In quantifying and characterizing this area of concern, we hope to reduce the uncertainty about the negative consequences of cryptocurrencies, allowing for more informed decisions by policymakers that assess both the costs and benefits. Hopefully, by shedding light on the dark side of cryptocurrencies, this research will help blockchain technologies reach their full potential.

[51] Formally, $Entropy = -\sum_d P(d)\log[P(d)]$, where $P(d)$ is the degree distribution (probability density of the degree for each user, where a user's degree, d, is the number of links the user has with other members of the same community).

[52] Iyengar, Saleh, Sethuraman, and Wang (2021) and Chod, Trichakis, Tsoukalas, Aspegren, and Weber (2021) study the application of blockchains in supply chains and Cao, Cong, and Yang study their application in auditing.

A second contribution of this chapter is the development of new approaches to identifying illegal activity in bitcoin, drawing on network cluster analysis and detection-controlled estimation techniques. These methods can be used by law enforcement authorities in surveillance activities. For example, our methods can be applied to blockchain data going forward as new blocks are created, allowing authorities to keep their finger on the pulse of illegal activity in bitcoin. Applied in this way, one could monitor trends in illegal activity such as its growth or decline, its response to various regulatory interventions such as seizures, and how its characteristics change through time. Such information could help make more effective use of scarce regulatory and enforcement resources.

During our sample period, a number of opaque cryptocurrencies such as Monero, Dash, and ZCash, also known as "privacy coins", emerged and gained some degree of adoption among illegal users. For example, some darknet marketplaces started accepting Monero for payments and our estimates suggest that such events negatively impacted the amount of illegal activity in bitcoin. While it is possible that further development of privacy coins could render our approach to detecting illegal activity less useful going forward, to date the major privacy coins have been shown to fall short of offering their users complete privacy. Using various heuristics and clustering algorithms, computer science researchers have been able to recreate user-level records and transaction activity in popular privacy coins such as Monero (Möser et al., 2018; Wijaya et al., 2018) and ZCash (Kappos et al., 2018). On the basis of such findings, privacy coins are perhaps not as private as they are intended to be. Therefore, even if illegal activity continues to migrate to popular privacy coins such as Monero and ZCash, law enforcement agencies and researchers could still use our approach applied across several cryptocurrencies, including privacy coins and non-privacy coins such as Bitcoin Cash, Litecoin, and Ethereum. It is possible that at some stage a truly private coin will be created for which it is not possible to undertake the type of analysis that is in this chapter.

Another surveillance application is in identifying individuals/entities of strategic importance, for example, major suppliers of illegal goods. Combining these empirical methods with other sources of information can "de-anonymize" the nameless entities identified in the data. This might be done, for example, by tracing the activity of particular individuals to the interface of bitcoin with either fiat currency or the regulated financial sector (many exchanges and brokers that convert cryptocurrencies to fiat currencies require the personal identification of clients). The methods that we develop can also be used in analyzing many other blockchains, though at present this might be more challenging for privacy coins.

Third, our finding that a substantial amount of illegal activity is facilitated by bitcoin suggests that bitcoin has contributed to the emergence of an online black market, which raises several welfare considerations. Should policymakers be concerned that people are buying and selling illegal goods such as drugs online and using the anonymity of cryptocurrencies to make the payments? This is an important question and the answer is not obvious. If the online market for illegal goods and services merely reflects a migration of activity that would have otherwise occurred "on the street" to the digital world of e-commerce, the illegal online activity facilitated via bitcoin might not be bad from a welfare perspective. In fact, there are many potential benefits to having illegal drugs and other goods bought and sold online rather than on the street. For example, it might be safer and lead to reduced violence (e.g., Barratt et al., 2016a). It could also increase the quality and safety of the drugs because darknet marketplaces rely heavily on user feedback and vendor online reputation, which can give a buyer access to more information about a seller's track record and product quality than when buying drugs on the street (e.g., Soska and Christin, 2015). There is also more choice in the goods offered, which has the potential to increase consumer welfare.

However, by making illegal goods more accessible, convenient, and reducing risk (due to anonymity), the darknet might encourage *more* consumption of illegal goods and increase reach, rather than simply migrating existing activity from the street to the online environment (Barratt et al., 2016b). Presuming illegal goods and services have negative net welfare consequences, then bitcoin and other cryptocurrencies could decrease welfare by enabling the online black market. Such negative consequences would have to be weighed up against welfare gains that also accompany cryptocurrencies.

Therefore, while the chapter does not provide a definitive answer to the question of welfare effects, it does get closer to an answer by having estimated both the trends and scale of illegal activity involving bitcoin (the most widely used cryptocurrency in darknet marketplaces). Future research might quantify the relation between drug trafficking on the street vs online (drawing on our methods or estimates) to understand to what extent we are experiencing a simple migration vs an expansion in the overall market. It might also quantify the benefits of moving to an online market and contrast them with the negative consequences of any expansion in the market as a result of it being more accessible / convenient / safe.

Our results also have implications for the intrinsic value of bitcoin. The rapid increase in the price of bitcoin in recent times has prompted much debate and divided opinions among market participants and even policymakers / central banks about whether cryptocurrency valuations are disconnected from fundamentals and whether their prices reflect a bubble. In part, the debate reflects the uncertainty about how to value cryptocurrencies and how to estimate a fundamental or intrinsic value. While we do not propose a valuation model, our results provide an input to an assessment of fundamental value in the following sense. One of the intrinsic uses of cryptocurrencies, giving them some fundamental value, is as a payment system. To make payments with bitcoin, one has to hold some bitcoin; the more widespread its use as a payment system, the greater the aggregate demand for holding bitcoin to make payments, which, given the fixed supply, implies a higher price. Our results suggest that currently, as a payment system, bitcoin is relatively widely used to facilitate trade in illegal goods and services and thus the illegal use of bitcoin is likely to be a meaningful contributor to bitcoin's fundamental value.

This observation—that a component of bitcoin's fundamental value derives from its use in illegal trade—raises a few issues. First, an ethical investor might not be comfortable investing in a security for which a meaningful component of the fundamental value derives from illegal use. Second, changes in the demand to use bitcoin in illegal trade are likely to impact its fundamental value. For example, increased attention from law enforcement agencies or increased adoption/substitution to more opaque alternative cryptocurrencies could materially decrease the fundamental value of bitcoin. Conversely, continued growth in the online black market with continued use of bitcoin, could further increase bitcoin's fundamental value. Third, the recent price appreciation of bitcoin greatly exceeds the growth in its use in illegal activity, suggesting either a substantial change in other components of bitcoin's fundamental value or a dislocation of the bitcoin price from its fundamental value.

3.4.2 Relation to other literature

This chapter contributes to three branches of literature. First, several recent chapter analyze the economics of cryptocurrencies and applications of blockchain technology to securities markets (e.g., Malinova and Park, 2016; Khapko and Zoican, 2017; Yermack, 2017; Huberman et al., 2017; Basu et al., 2019). The chapter contributes to this literature by showing that one of the major uses of cryptocurrencies as a payment system is in settings in which anonymity is valued (e.g., illegal trade).

Another related, although small, branch of literature examines the degree of anonymity in bitcoin by quantifying the extent to which various algorithms can identify entities/users in bitcoin blockchain data and track their activity (e.g., Meiklejohn et al., 2013; Ron and Shamir, 2013; Androulaki et al., 2013; Tasca et al., 2018). In doing so, some of these papers also provide insights about the different types of activities that use bitcoin. Of these papers, one of the closest to ours is Meiklejohn et al. (2013), who explore the bitcoin blockchain up to April 2013, clustering addresses into entities/users and manually identifying some of those entities by physically transacting with them. They are able to identify the addresses of some miners, exchanges, gambling services, and vendors/marketplaces (including one darknet marketplace), suggesting bitcoin entities are not completely anonymous. Tasca et al. (2018) use a similar approach to explore the different types of activity in bitcoin, focusing only on the largest entities, so-called "super clusters", and within that set, only those with a known identity. Fanusie and Robinson (2018) show how a sample of known illegal entities, many of which are darknet marketplaces, exchange bitcoin for other currencies or "wash"/launder their bitcoin holdings. They find that among the known illegal entities that they consider, darknet marketplaces account for most of the bitcoin exchange/laundering and that bitcoins from these illegal entities are mainly exchanged/laundered through bitcoin exchanges, bitcoin mixers/tumblers, and gambling providers.

None of these papers attempt to categorize all of the activity in bitcoin, nor do they try and quantify or characterize the population of illegal bitcoin users, which is the focus of this chapter. We exploit the lack of perfect anonymity that is documented in these studies and draw on some of the techniques from this literature to construct an initial sample of known illegal users. We add new methods to this literature, extending the empirical toolkit from making direct observations about individuals, to identification of communities and estimation of populations of users.

Yin and Vatrapu (2017) compare the performance of various supervised machine learning algorithms in classifying a sample of bitcoin users. Their analysis uses a sample of known entities, which includes some darknet marketplaces and other illicit entities. The algorithms that perform the best within their sample give widely varying estimates of the proportion of illegal users in sample, from 10.95% to 29.81%. While the study by Yin and Vatrapu focuses on the comparison of supervised machine learning algorithms, our study aims to provide comprehensive estimates of the scale and nature of illegal activity in bitcoin. The chapter therefore differs in that it analyzes all bitcoin activity, attempts to identify as much of the observable illegal activity as possible, and characterizes the trends and characteristics of the illegal activity.[53]

Finally, another related branch of literature is the recent studies of darknet marketplaces and the online drug trade, including papers from computer science and drug policy. For example, Soska and Christin (2015), use a web-crawler to scrape information from darknet marketplaces during 2013-2015, collecting a variety of data. Their paper provides valuable insights into these markets, including information about the types of goods and services traded (largely drugs), the number of goods listed, a lower bound on darknet turnover using posted feedback as a proxy (they do not have data on actual transactions/sales), the number of vendors, and the qualitative aspects of how these marketplaces operate (reputation, trust, feedback). The related drug policy studies often draw on other sources of information such as surveys of drug users and contribute insights such as: (i) darknet marketplaces like the Silk Road facilitate initiation into drug use or a return to drug use after cessation (Barratt et al., 2016b) and can encourage drug use through the provision of drug samples (Ladegaard, 2018); (ii) darknet forums can promote harm minimization by providing inexperienced users with support and knowledge from vendors and more experienced users (Bancroft, 2017); (iii) darknet marketplaces tend to reduce systemic violence compared with in-person drug trading because no face-to-face contact is required (Barratt et al., 2016a; Martin, 2018; Morselli et al., 2017); (iv) about one-quarter of the drugs traded on the Silk Road occur at a wholesale scale, suggesting that such markets might also indirectly serve drug users "on the street" by impacting dealers (Aldridge and Décary-Hétu, 2016); and (v) there are interesting cross-country differences in the use of the darknet marketplace "Agora" (Van Buskirk et al., 2016).

We contribute to this literature by quantifying the amount of illegal activity undertaken using bitcoin. All of the illegal activity captured by the existing studies of one or several darknet marketplaces is also in our measures because one of the approaches we use to construct a sample of observed illegal

[53] The results in Yin and Vatrapu (2017) are difficult to compare to ours for several reasons: their paper uses a non-random sample of bitcoin activity, whereas we analyze *all* bitcoin activity, they do not specify how they filter the blockchain data, cluster addresses to form entities / users or how they identify a sample of known entities as all of these steps were performed and provided by a data provider and are not reported, and their sample period is not specified.

activity involves measuring transactions with known darknet marketplaces. However, our estimates include much more than this activity—we use direct measures of transactions rather than a lower-bound measure such as feedback, consider all known darknet marketplaces (rather than one or a few), include two other methods of obtaining a sample of illegal activity, and most importantly, we estimate models that extrapolate from the *sample* of observed illegal activity to the estimated *population*. This yields vastly different and more comprehensive estimates. Empirically, we confirm that studies of darknet marketplaces only scratch the surface of the illegal activity involving bitcoin—the estimated population of illegal activity is several times larger than what can be "observed" through studying known darknet marketplaces. Furthermore, the studies of darknet marketplaces do not analyze how the characteristics of illegal and legal bitcoin users differ, or how recent developments such as increased mainstream interest in bitcoin and the emergence of new, more opaque cryptocurrencies impacts the use of bitcoin in illegal activity. These are further contributions of our chapter.

3.5 Conclusion

This chapter characterizes illegal users in bitcoin, the determinants of the detection, and the topology of the illegal bitcoin network. This chapter uses methods that characterize illegal activity while simultaneously controlling for selection bias that result from non-random detection.

This chapter's findings shed light on what drives illegal activity in cryptocurrencies and how law enforcement can use their resources efficiently to combat it. Illegal users of bitcoin tend to transact more, in smaller sized transactions, often repeatedly transacting with a given counterparty, and they tend hold less bitcoin. These features are consistent with their use of bitcoin as a payment system rather than for investment or speculation. Illegal users also make greater use of transaction techniques that obscure their activity, and their activity spikes following shocks to darknet marketplaces. The proportion of bitcoin activity associated with illegal trade declines with increasing mainstream interest and hype (bitcoin market value and Google search intensity), the emergence of more opaque alternative cryptocurrencies, and with fewer operating darknet marketplaces.

Second, the illegal network is characterized by much more interaction among the users consistent with the notion the illegal users use bitcoin as a means of transaction (not investment). In the illegal network, the users transact more, frequently act as both sender and recipient in transactions, and concentrate their transactions to few counterparties.

By enhancing our understanding of the characteristics of illegal activity in bitcoin and the illegal bitcoin network, this chapter sheds light on how to detect illegal activity in cryptocurrencies.

Chapter 4

Measuring market integrity

4.1 Introduction

Market integrity is one of the most important issues facing financial markets today, yet there are currently few direct methods for measuring the level of integrity. Regulators around the world often quote "market efficiency" and "market integrity" as their chief mandates.[54] However, market integrity receives much less attention than market efficiency, perhaps because it is so hard to measure. Regulators also often quote insider trading and market manipulation as the main types of misconduct challenging market integrity[55] and devote significant resources to combat it. For example, in 2003, the Bush administration increased the SEC's regulatory budget by 40% "to protect investors, root out fraud, and instill corporate social responsibility."[56] Financial misconduct such as insider trading also leads to significant costs when firms try to comply with insider trading legislation and implement compliance regimes. Legislators often vaguely define insider trading; therefore, firms are forced to implement highly restrictive yet costly[57] compliance regimes to avoid sanctions for ineffective compliance programs. Ironically, these costs eventually pass down to shareholders—the market participants the legislation was meant to protect (Anderson, 2015). This chapter uses insider trading and market manipulation when measuring market integrity because of the detrimental effects they can have on financial markets.

Around the world, governments use whistleblower schemes to uncover misconduct and penalties such as jail time and fines to deter it. Further, the past three decades have seen significant changes in financial market structure; fragmentation of trading across a number of competing trading venues has become the norm in many countries, while modern market features such as dark pools and colocation have become commonplace. With no real way to measure the effect of these regulatory strategies and market design features on integrity, governments and regulators often focus on efficiency—that is, the impact on transaction costs. The dynamics of the recent COVID-19 pandemic may also have increased insider trading opportunities by increasing the value and rate of new private

[54] The largest regulatory bodies who safeguard more than 50% of the world's listed domestic equity, including the US SEC, Canadian OSC, Australian ASIC, and the FCA in the UK, all list either "integrity" or "fair" in their mission statements.

[55] See, e.g., December 2015 markets article from the Australian regulator (ASIC): "Market integrity matters! You can play an important role in keeping our markets clean."

[56] See the *New York Times* article "Bush Proposes Big Increase In S.E.C. Budget" from February 2003.

[57] The costs come in terms of corporate culture, cost of compensation, share liquidity, and cost of capital (Anderson, 2015).

material information. Insider trading opportunities may further increase if the pandemic increases the amount of nonpublic information through delays in SEC disclosure filings.[58] Yet in the world's largest market—the US—insider trading prosecutions remain at an all-time low.[59] Meanwhile, regulators and industry professionals are struggling to distinguish between stock price manipulation and speculation when stock pumping, such as that perpetrated by the WallStreetBets Reddit forum, has become the "new normal."

This study develops three indices of market integrity (or lack thereof) applicable to financial markets around the world: an insider trading index, a market manipulation index, and a combined market integrity index. Our indices include the most detrimental and frequent forms of misconduct: insider trading before M&A events and closing price manipulation. Insider trading makes up about half of all US SEC and Department of Justice (DOJ) insider trading cases, while closing prices are one of the most important reference in valuing major financial instruments and entities.[60] These findings are just as important for emerging markets, which rely heavily on fair financial markets to attract outside investments and grow their economies.

Regulators, exchanges, and academics can use the two misconduct indices and the market integrity index from this chapter to measure the level of insider trading, market manipulation, and market integrity in any global equity exchanges. Regulators can use the indices to keep a finger on the "pulse" of their financial markets and measure the direct effects of their regulatory efforts. Industry professionals, including companies seeking capital and investors seeking to invest, can ascertain if the markets they operate in meet their integrity requirements. Academics can use the indices in empirical studies to further our understanding of market integrity, which has so far received far less attention than market efficiency.

Our market integrity index is composed of six insider trading and market manipulation metrics that we validate using a sample of prosecuted insider trading and market manipulation cases. We then combine the metrics into three indices of market misconduct—one insider trading index (IT), a market manipulation index (MM), and a market integrity index (MI).

When measuring insider trading, we use M&A events, and for market manipulation, we use closing prices. Using M&A events to measure insider trading has two advantages. First, M&A events are unscheduled (as opposed to earnings announcements), which decreases uninformed investor participation and increases the proportion of informed trading activity (to total trading activity) before the announcement. Second, M&A announcements usually lead to positive target stock returns. The low

[58] See the SEC's public statement on March 23, 2020, "Statement from Stephanie Avakian and Steven Peikin, co-directors of the SEC's Division of Enforcement, regarding market integrity" (https://www.sec.gov/page/news).

[59] See the *NPR* article "Under Trump, SEC enforcement of insider trading dropped to lowest point in decades" (npr.org).

[60] Including derivative instruments, director options, broker performance, mutual fund net asset values, and stock indices (Comerton-Forde and Putniņš, 2011). Patel and Putniņš (2021) note that the majority of prosecuted insider trading is related to insider trading before M&A announcements and that one in five M&A announcements are insider traded.

uninformed investor participation coupled with the (usually) one-directional stock returns increases the accuracy of the measures we create to capture insider trading. Similar to M&A announcements, closing prices also usually lead to one-directional stock returns because manipulators tend to manipulate the stock price up (not down).[61] Therefore, using closing prices to measure market manipulation increases the accuracy of the measures we create. Closing price manipulation is also a particularly detrimental type of market manipulation because closing prices are an important component when valuing financial instruments and entities (Kahan, 1992).[62] Three characteristics are statistically significant predictors of insider trading: abnormal returns, abnormal volumes, and abnormal order imbalances in the five days before M&A events. Similarly, three intraday characteristics are statistically significant predictors of closing price manipulation: price volatility and positive order imbalance in the last two hours of trading, as well as abnormal price reversals. We improve the individual ability of these measures to pick up insider trading and market manipulation by combining them into indices.

This chapter also has implications for the academic literature; compared with market efficiency, there is far less research in market integrity because it is so difficult to measure. Our indices contribute to the academic literature by providing measures of insider trading, market manipulation, and market integrity.

This chapter relates to the literature on the costs and benefits of insider trading and market manipulation. From an efficiency standpoint, Bhattacharya and Daouk (2002) show that insider trading increases the cost of capital as liquidity providers protect themselves against informed trades by widening bid-ask spreads, with large stockholders shifting from monitoring management to trading on stock tips. Insider trading can also cause the market to demand a higher risk premium over the risk-free rate on newly issued equity (Gregoire and Huang, 2009), hence decreasing investor participation by decreasing the liquidity and increasing the transaction costs (Leland, 1992). Price accuracy also benefits from information collection by outside investors, who may leave the market if faced with higher likelihood of trading with an insider (Fishman and Hagerty, 1992). However, Manne (1966) and Carlton and Fischel (1983) argue that insider trading can lead to an increase in market efficiency through accurate and timely incorporation of information into prices, thereby improving stock price accuracy. The measures developed in this chapter move the literature further by providing measures of insider trading and market manipulation that can be used when measuring the costs and benefits of either type of misconduct.

The academic literature is also divided on the effects of market manipulation on financial markets. Market manipulation can discourage market participation, decrease market liquidity, increase trading costs, and increase the cost of capital, which lowers the number of stock listings. Low liquidity

[61] The large majority of our prosecutions sample indicates cases where the manipulator pushes the closing price up.

[62] Some examples are stock indices, net asset values (NAVs) in mutual funds, derivative instruments, broker performance, and incentive devices such as CEO bonuses and stock options (Comerton-Forde and Putniņš, 2011).

levels and price distortions induced by market manipulation inhibit the market's price discovery mechanism and reduce stock price accuracy (Comerton-Forde and Putniņš, 2011).[63] Finally, inaccurate prices lead to inefficient resources allocation and reduce economic efficiency (Goldstein and Guembel, 2008). However, Hanson and Oprea (2009) argue that market manipulators can act as a type of noise trader, potentially attracting informed investors who increase price accuracy through their informed trades.

The closest we currently have to a global benchmark of market integrity is indices of securities laws (e.g., Cumming and Johan, 2008). Although such measures are helpful in understanding the legal stance taken by countries toward such conduct, they are limited in their ability to measure the level of integrity. For example, Bhattacharya and Daouk (2002) show that it is not merely the existence of rules, but rather the enforcement of those rules that affects market integrity and—through it—the cost of capital. There is also relatively little cross-sectional variation in such rules, whether across countries or through time. If measured accurately, integrity is expected to vary both cross-sectionally and over time. Finally, rules cannot be used to gauge the efficacy of a regulator in ensuring compliance with rules. We take the first step in measuring both forms of integrity.

We also contribute to the literature on combating market misconduct. There is ample empirical evidence of the positive effects a stronger regulatory regime can have on market quality. Beny (2006) finds that countries with more stringent insider trading laws have (i) more dispersed equity ownership (ii) more liquid stock markets, and (iii) more informative stock prices. Daouk, Lee, and Ng (2005) find that improved capital market governance decreases the cost of equity capital, increases market liquidity, and increases market pricing efficiency. Bhattacharya and Daouk (2002) demonstrate that the enforcement of insider trading laws is necessary to observe a meaningful effect on the cost of capital. Fernandes and Ferreira (2008) show that price informativeness improves after the first insider trading prosecution in developed countries. Augustin, Brenner, and Subrahmanyam (2015) find that only 9% of the options trades that they identify as informed are litigated by the SEC. Comerton-Forde and Putniņš (2011) find that only 0.4% of all closing price manipulation is prosecuted. We contribute to the literature by showing what types of stocks are vulnerable to insider trading and closing price manipulation, along with how regulatory strategies and changes in market design may increase market integrity.

The structure of the rest of the chapter is as follows: Section 4.2 documents the data collection process, Section 4.3 describes the creation and validation of our market integrity measures, and Section 4.4 applies the measures to examine market integrity in US stock markets. Section 4.5 concludes this chapter.

[63] Comerton-Forde and Putniņš (2011) find that closing price manipulation causes abnormal end-of-day returns that are six times larger than normal (between 1.4% and 1.9%) and increases spreads by between 0.11% and 0.63%.

4.2 Data

We collect data from a variety of data sources, including historical data on the details of prosecuted financial misconduct, M&A events, and trade and quote data, to produce and validate our metrics. This section explains how we collect and filter the data.

4.2.1 Prosecuted insider trading and market manipulation cases

We use a sample of prosecuted insider trading and market manipulation cases (the "prosecuted samples") to calibrate and validate our metrics. Our sample comprises cases in which individuals were prosecuted for using private information to buy M&A target stocks prior to public announcement or for manipulating the closing price of a particular stock.

We extract insider trading prosecutions from the SEC litigations releases concerning civil lawsuits brought by the SEC in federal court (dating back to September 20, 1995).[64] We supplement this with information included in SEC annual reports and all relevant court complaints available on PACER (the online archive of all court records in US courts). PACER data are obtained by sourcing a list of all court records associated with each SEC case and then cross-checking that list against the court documents available directly from the SEC. Any documents missing from the SEC archive are included in the data collection process.[65]

We extract manipulation prosecutions data from litigation releases and filings from US and Canadian market regulators.[66] We then supplement this sample with prosecution cases from the legal databases Lexis, Quicklaw, and West-law, supplementing them with data from PACER when data are missing.

We record the following features of each case: the unique case number and date of the earliest SEC release about this case; the total number people prosecuted in the case; the number of people who traded on this announcement (in the case); the number of "ring members" involved in the spread of the nonpublic information; and the degree of separation from the person who traded to the original source of the information. The name, age, occupation, and location of the individual making the trade are also recorded, along with the resulting financial penalty (if known) imposed on them from the case.

The minimum criteria for a prosecution to be included in our sample is that it concerns insider trading in an M&A target stock prior to a public announcement or closing price manipulation in an identifiable public company over an identified time period. Note that this includes both cases that go to

[64] SEC, "Litigation Releases," https://www.sec.gov/litigation/litreleases.shtml.

[65] Although certain cases may also have been pursued by the Department of Justice (in cases that went through the criminal, rather than civil, legal system), criminal court records from the Department of Justice are not yet available to the same degree. As such, the SEC's civil cases are the only current source of insider trading data.

[66] These include the US Securities and Exchange Commission (USA), Ontario Securities Commission (Canada), Market Regulation Services Inc. (Canada), Investment Dealers Association (Canada), Mutual Funds Dealers Association (Canada), Investment Industry Regulatory Organization of Canada (Canada), NYSE Regulation Inc. (USA), and AMEX Division of Regulation and Compliance (USA).

court and those settled out of court. Prosecution cases are removed from our sample if (i) the date on which the offense was committed or the name of the stock subject to the offense is missing; (ii) the stock is not a common stock; (iii) the violation occurred in an over-the-counter market; 3) the violation did not involve trade-based techniques; or (iv) trade and quote data are unavailable for the given stock on the date of violation.[67] In total, our insider trading prosecution sample consists of 442 M&A events from April 1996 to April 2016 within which insiders traded in stocks listed on the Nasdaq, NYSE, or AMEX and for which we have the required data. Our prosecuted manipulation sample consists of 191 instances of manipulation from October 1998 to June 2013 on the Nasdaq and New York Stock Exchange (NYSE) in the US and Toronto Stock Exchange (TSX) and TSX Venture (TSXV) in Canada, for which we have the required data.

We obtain trading data for each insider trading and market manipulation case from the Refinitiv Tick History database, including intraday five-minute snapshots of quoted and traded prices and volumes.

4.2.2 M&A event and trading data

We collect 150,516 M&A announcements in the US from 1996 to 2020 from SDC platinum. The sample only considers M&A announcements where the acquirer achieves functional control by acquiring more than 50% of the target stock. The sample also excludes all prosecuted cases of insider trading, hence constituting our "nonprosecuted" insider trading sample.

Within this sample, we consider stocks that meet the following two criteria:

1. The Stock's ticker or target nation must not be missing. A total of 140,124 events fail this criterion reducing the sample to 10,092 announcements.
2. We only keep material M&A announcements, where the materiality requires that the M&A announcement has a significant impact on the target stock price. Our condition is that the cumulative abnormal return from 1 day before to 1 day after the announcement is at least 10% and from 5 days before to 1 day after the announcement is at least 5%. We remove 5,361 M&A announcements, reducing the sample to 4,737 announcements.

For the almost 5000 stocks subject to acquisition bids in the US from 1996 to 2020, we extract intraday five-minute frequency trading data from the Refinitiv Tick History database.

Using M&A events has two advantages. First, anecdotal evidence of M&A events shows that announcement day returns are strongly positive on average. Therefore, insider trading in M&A events usually involve buying (as opposed to selling) the target stock, simplifying our metric construction.

[67] By "trade-based," we mean manipulation by submitting buy or sell orders to the market. For example, Jonathan G. Lebed, who made 273,000 in illegal gains by manipulating the prices of six stocks by posting optimistic messages in investing chat rooms, is a form of non-trade-based manipulation (Morgenson, 2006).

Second, M&As are unannounced and unscheduled, as opposed to other price-sensitive events such as earnings announcements. As a result, noise traders do not increase their participation in anticipation of the upcoming event, and most abnormal trading activity prior can be attributed to informed trading.

Third, previous studies have found that M&A events are particularly prone to insider trading. Augustin et al. (2015) look at options trading before M&A announcements in the US and find that 25% of all M&A deals in their sample are subject to informed options trading in the preannouncement period. In terms of stock trading, Morgenson (2006) reports that in 41% of the largest US mergers, the companies receiving buyout bids exhibit suspicious and abnormal trading activity immediately before the announcement. Further, Keown and Pinkerton (1981) show that almost half of the price adjustments in listed securities occurs before the M&A announcement date. M&A-related deals also seem to be more susceptible to insider trading relative to other firm announcements; Meulbroek (1992) finds that 79% of the SEC cases they examine are takeover related. Similarly, in two of the largest insider trading cases (Levine-Boesky-Milken from the late 1980s[68] and the Galleon hedge fund case in 2009[69]), the majority of the charges were related to insider trading in takeover cases.

Finally, M&A events present a significant enforcement challenge for regulators when attempting to prevent corporate insiders from trading on insider information prior to M&A announcements of their own firms. Despite corporate insiders being required to file their trades in their own companies with the SEC, Agrawal and Nasser (2012) document a rise in corporate insider trading prior to M&A events, suggesting ineffective regulatory enforcement.

4.2.3 Closing price data

While our M&A metrics require an announcement to assess potential integrity violations, our manipulation metrics simply rely on the market closing. Each market close is an opportunity for a potential manipulator to manipulate the closing price of stocks listed on the market, hence providing a more frequent measure of misconduct. The sample also excludes any prosecuted cases of closing price manipulation, therefore constituting our “nonprosecuted” market manipulation sample.

We randomly select 200 US stocks each year from 1996 to 2020 subject to the following minimum data requirements:

1. The stock must trade at least 80% of days (not counting weekends and holidays).
2. The stock must have an average daily volume exceeding the median average daily volume of the stocks in the country.

[68] See Frantz (1987).

[69] See Bray (2010), Bray and Strasburg (2009), and Sharma and Pulliam (2009).

For the US stock sample, we extract intraday five-minute frequency trading data from the Refinitiv Tick History database from 1996 to 2020.

4.2.4 The characteristics of the prosecuted and nonprosecuted samples

This section shows the univariate differences between the prosecuted and nonprosecuted insider trading and market manipulation samples. The prosecuted sample includes known cases of prosecuted insider trading and market manipulation, and the nonprosecuted samples include M&A events and closing prices where no misconduct (insider trading or market manipulation) was prosecuted. Table 4.1 shows the difference in the means and medians between our prosecuted and nonprosecuted insider trading and market manipulation samples.

Table 4.1
Prosecuted and nonprosecuted sample univariate statistics
This table reports univariate statistics for our prosecuted and nonprosecuted insider trading and market manipulation samples. The first two columns ("*Prosecuted*") report the means and medians for the prosecuted portion of our sample, the second two columns ("*Nonprosecuted*") report the means and medians for the nonprosecuted portion of our sample, and the last two columns ("*Difference*") report the difference between the prosecuted and nonprosecuted means and medians. In Panels A and B, we report the sample characteristics *Market cap* and *Volume* for our insider trading sample (Panel A) and market manipulation sample (Panel B). *Market cap* is the company market value, and *Volume* is the daily traded volume in its stock. The significance of the difference in means is computed with two-sided t-tests, and the significance of the difference in medians is computed with the Wilcoxon signed-rank test. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. We report all index values in percentage.

Variable	Prosecuted (1)		Nonprosecuted (2)		Difference (1-2)	
	Mean	Median	Mean	Median	Mean	Median
Panel A: Insider trading sample characteristics						
Market cap	2,971.02	1,035.53	1,613.92	238.98	1,357.11***	796.55***
Volume	31.31	9.09	12.92	1.30	18.39***	7.80***
Panel B: Market manipulation sample characteristics						
Market cap	740.61	85.21	8,332.79	1,528.90	-7,592.18***	-1,443.69***
Volume	3.37	0.13	44.41	11.99	-41.05***	-11.86***

Starting with our insider trading sample (Panel A), the prosecuted cases of insider trading tends to occur in larger, more liquid target firms. Companies with prosecuted cases of insider trading have an average market capitalization almost double the size of other target firms (the nonprosecuted sample) and a typical (median) company's market capitalization that is three and a half times larger. The prosecuted sample's mean daily dollar volume is about twice that of the nonprosecuted sample, while a median trading day's dollar volume is more than six times larger. The sample statistics accord with our expectations, with insider trading being more likely in liquid stocks of large companies, where it is easier to hide informed order flow.

In contrast, in Panel B, we find that prosecuted cases of market manipulation skew toward smaller, less liquid stocks. The average (mean) market capitalization for stocks in the nonprosecuted sample is ten times larger than the market capitalization of nonprosecuted stocks, while a typical (median) stock has a market capitalization that is almost twenty times larger. In terms of liquidity, the mean daily dollar volume is ten times larger for the stocks in the nonprosecuted sample than the stocks in the prosecuted, while the median is one-hundred times larger. These differences in size and liquidity are consistent with the findings of Comerton-Forde and Putniņš (2011), who note that the closing prices of less liquid stocks are easier to control because the manipulator needs to compete with fewer other traders.

4.3 How do we measure market integrity?

We now produce a set of insider trading and market manipulation metrics and explain the intuition behind their construction. We then combine our metrics into an insider trading index, a market manipulation index, and, a market integrity index. Finally, we measure the accuracy of our metrics and indices by validating them against the sample of known manipulation cases.

4.3.1 Development and validation of insider trading measures

We construct a suite of metrics capturing abnormal trading in the target's stock prior to the announcement. To compute our metrics, we base our metrics on the M&A target's price movements, traded volumes, and order imbalance (number of buys minus sells) in the days before and after the M&A announcement.

The case from Section 1.1.1 illustrates a typical example of insider trading in DSC Communications Corp ("DSC") shares. In the example, the insider trader, Gary D. Force, purchases 50,000 of DSC shares valued $18.44 on June, 2 1998 (two days before the M&A announcement). On the day of his trading, the price return is 11% compared with the DSC's average daily return in June of 3%. On June 3 (one day before the announcement), the traded volume in DSC shares increased to almost 12 million (11,486,000)—66% higher than the average daily traded volume in DSC shares in June. Other insider may also have traded on the day and contributed to the high return. [70] From the case, likely candidates of insider trading would be abnormal volume, abnormal returns, and abnormal buying activity.

Our first insider trading metric uses abnormal target stock returns prior to the M&A announcement. Many studies show that insider trading causes information to be reflected in prices before the public announcement (e.g., Meulbroek, 1992; Fishe and Robe, 2004) and in the specific case of M&A announcements that it is unlikely that the preannouncement price run-up is because of reasons

[70] The US District Court case file (05 CV 5411) specifies that Force's trading was part of a large insider trading scheme involving more than 20 people.

apart from insider trading (e.g., Augustin et al., 2015). Furthermore, the extent of the preannouncement run-up is determined by how widely and freely people trade on private information; for example, Del Guercio, Odders-White, and Ready (2017) find that insider trading enforcement intensity reduces preannouncement price run-ups. The Financial Conduct Authority (FCA) and Australian Securities and Investment Commission (ASIC) also use price run-ups when measuring market cleanliness and information leakage.[71]

We measure the price run-up that occurs in the five days before the announcement as a proportion of the total abnormal returns around the M&A announcement (from five days before to two days after):[72]

$$IT\_RunUp = \frac{CAR(-5,-1)}{CAR(-5,+2)}, \tag{4.1}$$

where $CAR(-5,-1)$ is the cumulative abnormal return (the midquote return of the target stock minus the midquote return on the corresponding market index) from five days before the takeover announcement to the day before the announcement. Similarly, $CAR(-5,+2)$ is the cumulative abnormal return from five days before to two days after the announcement. If $CAR(-5,-1) > CAR(-5,+2)$, then the metric is set to one because all information has then been impounded into the market prior to the M&A announcement. Similarly, any metric value less than negative one is set to negative one. Hence, IT_RunUp can take values in the range $[-1,+1]$, with larger values indicating a higher likelihood of insider trading in the target's stock.

Our second insider trading metric uses the volume traded in the target's stock prior to the M&A announcement, here in the spirit of Baruch, Panayides, and Venkataraman (2017). The notion of increased volume in the target's stock is consistent with Meulbroek (1992), who finds that traded volume is significantly higher in the presence of insider trading compared with past volume and market-wide volume. Akey, Gregoire, and Martineau (2020) also find that volume-based measures robustly measure insider trading prior to public announcements. The FCA and ASIC also use abnormal trading volumes when measuring insider trading before prices sensitive announcements.[73]

To capture the level of volume likely induced by insider trading, we measure the volume in the 5 days before the M&A announcement compared with a benchmark period starting 30 days before and ending 11 days before the M&A announcement:

$$IT\_Volume = \frac{V(-5,-1)-V(-30,-11)}{V(-5,-1)+V(-30,-11)}, \tag{4.2}$$

where $V(-5,-1)$ is the average daily volume traded in the target stock from 5 to 1 days before the announcement and $V(-30,-11)$ is the average daily volume traded in the target stock from 30 to 11 days before the announcement. IT_Volume can take values in the range $[-1,+1]$, with larger values

[71] See FCA occasional paper No.4 "Why has the FCAs market cleanliness statistics for takeover announcements decreased since 2009?" and ASIC report 487 "Review of Australian equity market cleanliness."

[72] In our prosecutions sample, we find that 92% of all insider trading happens in the five days before the M&A announcement.

[73] See 2019/20 market cleanliness statistics published on the FCA website (https://www.fca.org.uk/data/market-cleanliness-statistics-2019-20) and ASIC report 487 "Review of Australian equity market cleanliness."

indicating elevated volume in the days leading up to the announcement, which may indicate more insider trading.

Our third insider trading metric considers the direction of traded volume in the target's stock prior to the M&A announcement, here in the spirit of Irvine, Lipson, and Puckett (2007), Christophe, Ferri, and Hsieh (2010), Bernile, Hu, and Tang (2016), Kacperczyk and Pagnotta (2019), and Ahern (2020). We subtract seller-initiated volume from buyer-initiated volume and scale the difference by the total volume traded in the target stock prior the event date. This process creates a measure of volume imbalance. To assign direction to the volume traded in each five-minute interval, we use the bulk volume classification (BVC) developed by Easley, Lopez de Prado, and O'Hara (2016):[74]

$$IT\_Imbalance = \frac{B(-5,-1)-S(-5,-1)}{B(-5,-1)+S(-5,-1)}, \tag{4.3}$$

where $B(-5,-1)$ and $S(-5,-1)$ are buyer-initiated and seller-initiated volumes, respectively, from five days to one day before the M&A announcement. $IT\_Imbalance$ takes values in the range $[-1,+1]$, much like the other metrics, with larger values indicating a proportionately larger amount of buyer-initiated volume than seller-initiated volume, which is likely to be indicative of (potential) insider trading.

We measure the accuracy of our insider trading metrics in correctly identifying the prosecuted cases of insider trading in our sample of US M&A events by computing an AUROC score for each metric separately. AUROC scores range between zero and one, where higher values indicate a higher classification accuracy and values above 0.5 indicate classification accuracy above chance. For our specification, we use the AUROC scores to ascertain the performance of our metrics and indices in correctly identifying insider trading and market manipulation using a set of prosecuted (insider trading and market manipulation) cases for which we know misconduct occurred.[75] With technological advancement regulators being able to detect less egregious conduct, the more recent part of our prosecution sample includes cases in which the misconduct is less obvious.[76] We account for this variability in regulatory efficacy by computing the AUROC scores separately for each year (comparing prosecution cases to the nonprosecuted sample of the same year) and taking the average. Note that the nonprosecuted part of our sample almost certainly contains unprosecuted or undetected cases of insider trading, which makes the AUROC scores reflect a conservative estimate of the accuracy of the metrics because they are "penalized" by the AUROC for flagging unprosecuted cases of insider trading.

[74] We use the formula $V_\tau^B = V_\tau \cdot Z\left(\frac{\log(Close_\tau)-\log(Close_{\tau-1})}{\sigma_{\Delta P}}\right)$ to compute the buyer-initiated volume and $V_\tau^S = V - V_\tau^B$ to compute the sell-initiated volume, where V_τ is the volume traded during the time period τ, Z is the CDF of the standard normal distribution, $\log(Close_\tau) - \log(Close_{\tau-1})$ is the log price change between the two time periods, and $\sigma_{\Delta P}$ is a measure of the volume-weighted standard deviation of the log price changes.

[75] AUROC scores are commonly presented graphically with the "true positive rate" (TPR) on the y-axis and "true negative rate" (TNR) on the x-axis. For more on the application of AUROC scores, see Stein (2005) and Tang and Chi (2005).

[76] One such advancement is SMARTS software, which actively screens through trading data and flags potential cases of misconduct.

Panel A of Figure 4.1 illustrates the insider trading AUROC scores, here showing the performance of each metric in correctly identifying insider trading. All three metrics significantly predict insider trading. The best predictor of insider trading is the imbalance of buy and sell orders in the target stock five days up to the announcement ($IT\_Imbalance$). This metric has an AUROC score of 0.631, which is statistically significantly greater than 0.5 at the 99% confidence level. The other two predictors also have classification accuracy significantly above chance, with IT_RunUp scoring 0.573 (statistically significantly greater than 0.5 at the 99% confidence level) and IT_Volume scoring 0.574 (statistically significantly greater than 0.5 at the 90% confidence level).

We assess the differences in our metric and index values between the prosecuted and nonprosecuted samples as an additional way of gauging whether they are useful proxies for insider trading. Panel A of Table 4.2 compares the prosecuted sample of insider trading in M&As to all other M&A events. The nonprosecuted M&A events are not all "free" of insider trading—they contain a mix of non-violations and undetected violations. Despite this, the difference in means indicates that the prosecuted sample displays higher scores on the insider trading metrics compared with the nonprosecuted sample. The price runup and pre-M&A buying activity in the prosecuted sample are more than double the levels in the nonprosecuted sample (difference in mean IT_RunUp and $IT\_Imbalance$ is 10.24 and 1.51 respectively) and the overall trading activity is six times larger (the difference in IT_Volume is 11.86). The typical (median) M&A in the prosecuted sample has a run-up that is almost 50% larger (difference in median IT_RunUp is 2.71) than in the nonprosecuted sample and buying imbalance and overall trading activity that is around twice as high as the nonprosecuted sample (difference in median IT_Volume is 3.74, and $IT\_Imbalance$ is 0.78). The insider trading index that combines all three of the insider trading metrics is about 4 index points larger in the prosecuted sample (mean IT_Index is 55.92% in the prosecuted sample and 52.01% in the nonprosecuted sample).

Panel A: AUROC scores for market integrity metrics

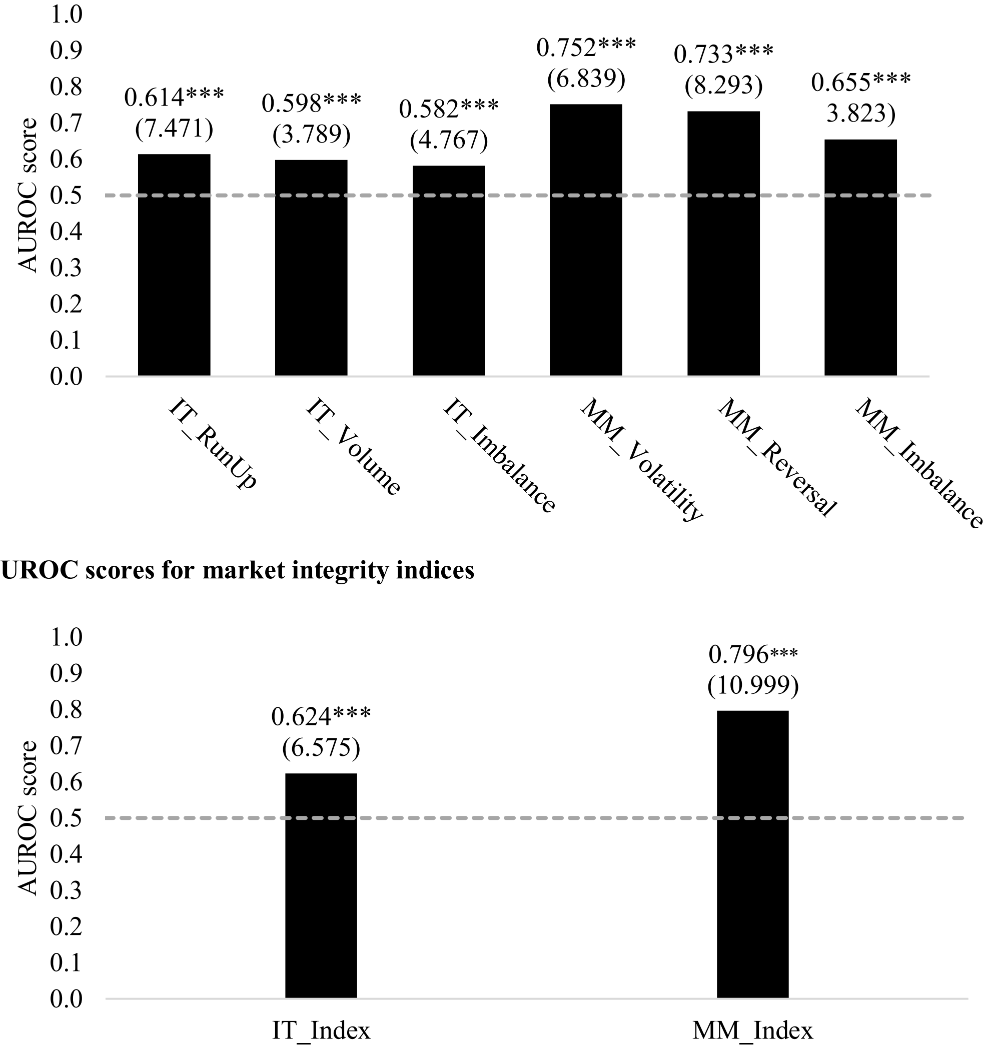

Panel B: AUROC scores for market integrity indices

Figure 4.1
Classification accuracy for market integrity metrics and indices
This figure reports the AUROC scores for the market integrity metrics (Panel A) and indices (Panel B). *IT_RunUp* measures the price run-up in the five days before an M&A announcement. *IT_Volume* measures the abnormal volume in the M&A target's stock in the five days before an M&A announcement. *IT_Imbalance* is the imbalance between buyers and sellers in the M&A target's stock in the five days before the M&A announcement. *MM_Volatility* measures the stock price volatility in the two hours before the market closes. *MM_Reversal* measures whether a day-end return reverses the following day. *MM_Imbalance* is the imbalance between buyers and sellers at the end of the trading day. Panel B reports the market integrity indices. The insider trading index (*IT_Index*) is the equally-weighted average of the insider trading metrics, and market manipulation (*MM_Index*) is the equally-weighted average of the market manipulation metrics. When computing the AUROC scores, we control for the change in legal enforcement effectiveness across time by computing an AUROC score for each year between 1996 and 2016 and then taking the average of the 20 values. The dashed gray line indicates an AUROC score of 0.5, which is equivalent to the classification accuracy of pure chance. Numbers reported above the bars are the AUROC scores. Numbers in brackets are the significance of the AUROC score's difference from 0.5 computed with two-sided t-tests. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 4.2
Insider trading and market manipulation metrics in prosecuted and nonprosecuted samples
This table reports univariate statistics for the prosecuted and nonprosecuted insider trading and market manipulation samples. The first two columns ("*Prosecuted*") report the means and medians for the sample of prosecuted misconduct cases, the second two columns ("*Nonprosecuted*") report the means and medians for the sample that does not have prosecuted misconduct, and the last two columns ("*Difference*") report the difference between the prosecuted and nonprosecuted means and medians. Panel A reports our insider trading metrics and index. *IT_RunUp* measures the price run-up in the five days before an M&A announcement. *IT_Volume* measures the abnormal volume in the M&A target's stock in the five days before an M&A announcement. *IT_Imbalance* is the imbalance between buyers and sellers in the M&A target's stock in the five days before the M&A announcement. The insider trading index (*IT_Index*) is the equally-weighted average of the insider trading metrics. Panel B reports our market manipulation metrics and index. *MM_Volatility* measures the stock price volatility in the two hours before the market closes. *MM_Reversal* measures whether a day-end return reverses the following day. *MM_Imbalance* is the imbalance between buyers and sellers at the end of the trading day. The market manipulation (*MM_Index*) is the equally-weighted average of the market manipulation metrics. The significance of the difference in means is computed with two-sided t-tests and the significance of the difference in medians is computed with the Wilcoxon signed-rank test. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. All index values are percentages.

Variable	Prosecuted (1)		Nonprosecuted (2)		Difference (1-2)	
	Mean	Median	Mean	Median	Mean	Median
Panel A: Insider trading metrics and index						
IT_RunUp	18.85	8.66	8.61	5.95	10.24***	2.71***
IT_Volume	14.23	6.52	2.36	2.78	11.86***	3.74*
IT_Imbalance	2.47	1.86	0.96	1.08	1.51***	0.78***
IT_Index	55.92	52.78	52.01	51.69	3.91***	1.10***
Panel B: Market manipulation metrics and index						
MM_Volatility	47.07	47.73	35.38	35.94	11.68***	11.79***
MM_Reversal	42.73	36.89	12.70	0.00	30.02***	36.89***
MM_Imbalance	7.83	5.95	0.12	0.00	7.71**	5.95***
MM_Index	49.41	47.98	37.19	34.95	12.22***	13.03***

4.3.2 Development and validation of market manipulation measures

We construct a suite of metrics designed to capture closing price manipulation. The metrics are based on return reversals, price volatility at the close, and traded volumes in the minutes before the close.

The case from Section 1.1.1 illustrates a typical example of closing price manipulation in TV Azteca ("TZA") shares. In the example, the manipulator, Moises Saba Masri, manipulates the closing price of DSC to above $5 with seven buy orders in the last minutes before the market closes. The buy orders account for 75% of all buy side activity in TZA on the day. The extreme price movement may also have caused significant volatility in TZA's share price. From the case, likely candidates of closing price manipulation are abnormal returns, abnormal buying activity, and abnormal price volatility.

Our first manipulation metric is day-end price volatility, here as measured by the standard deviation of stock prices in the hours before the market close. The use of price volatility to capture closing price manipulation follows the work of Hillion and Suominen (2004), who argue that market manipulation leads to increased volatility during the last minute of trading. We measure the standard deviation of the trade prices two hours before the close and benchmark this day-end volatility against the volatility during the entire day. This benchmarking reduces the tendency for the metrics to simply pick up volatile stocks, instead focusing on abnormal day-end volatility:

$$MM\_Volatility = \frac{\sigma(14{:}00,16{:}00)}{\sigma(14{:}00,16{:}00)+\sigma(9{:}30,16{:}00)}, \tag{4.4}$$

where $\sigma(14{:}00,16{:}00)$ and $\sigma(9{:}30,16{:}00)$ are the standard deviations of the trade prices in the last two hours and the entire trading session, respectively.[77] $MM\_Volatility$ can take values in the range of $[0,+1]$, with larger values indicating proportionately more volatility in the trade prices in the two hours before the market closes compared with the general level of volatility in the stock during the day, reflecting a higher likelihood of closing price manipulation.

Our second manipulation metric looks for day-end returns that reverse the next day. Comerton-Forde and Putniņš (2011) find that closing price manipulation causes abnormal end-of-day returns six times larger than normal and that these returns tend to reverse the following day, which is what distinguishes them from returns gained because of new information. We measure reversals by comparing day-end price movements on day d, which is denoted R_d, and price movements on the following day, R_{d+1}:

$$Reversal = \min(\max(0, R_d), \max(0, -R_{d+1})), \tag{4.5}$$

where R_d is the return from the midday midquote price to the closing price on day d:

$$R_d = \frac{CP_d - MQ_d}{MQ_d}, \tag{4.6}$$

and R_{d+1} is the return from the closing price on day d to the midday midquote on the following day:

$$R_{d+1} = \frac{MQ_{d+1} - CP_d}{MQ_{d+1}}, \tag{4.7}$$

and CP_d and MQ_d is the closing price and the midday midquote on day d.[78]

We control for stock volatility (more volatile stocks naturally produce larger reversals than less volatile stocks) by dividing the reversal by the largest reversal in the previous 20 days:[79]

$$MM\_Reversal = \frac{Reversal_d}{\max(Reversal_{d-20}, \dots, Reversal_d)}. \tag{4.8}$$

According to the prosecuted cases in our sample, closing price manipulation is almost always used to push prices up, not down. Therefore, we exclude negative reversals from the metric because these are more likely to reflect volatility as opposed to manipulation. We also constrain $MM\_Reversal$ to be less

[77] The markets in this study vary in design and timing. When required, we adjust $\sigma(9{:}30,16{:}00)$ to the opening and closing hours of the market studied.

[78] The time 12:45 is the exact middle of a trading day in the US market, which starts at 9:30 in the morning and ends at 16:00. Opening and closing times vary across markets, and individual markets do sometimes change them. When computing the metrics, we consider the opening and closing times of the specific market.

[79] We also try two other variations to control for stock price volatility, including i) dividing the reversal by the average reversal in the previous 20 days and ii) dividing the reversal by the standard deviation of stock returns in the previous 20 days. We settle on dividing by the largest reversal in the previous 20 days because it forces the measure to not exceed +1.

than or equal to one by setting values above one to one so that the range of the metric is $[0, +1]$. Larger values of $MM\_Reversal$ indicate a higher likelihood of closing price manipulation.

Our third manipulation metric considers the imbalance in the direction of traded volume in the stock in the minutes before the market close. Our use of order imbalance to capture manipulation is consistent with Comerton-Forde and Putniņš (2011), who find that manipulators create an order imbalance that persists for several minutes. To measure order imbalance, we subtract seller-initiated volume from buyer-initiated volume and scale the difference by total volume in the stock in the last two hours of trading. We again use bulk volume classification (BVC) developed by Easley et al. (2016) to determine the direction to traded volumes in five-minute buckets:

$$MM\_Imbalance = \frac{B(14{:}00,16{:}00) - S(14{:}00,16{:}00)}{B(14{:}00,16{:}00) + S(14{:}00,16{:}00)}, \quad (4.9)$$

where $B(14{:}00,16{:}00)$ and $S(14{:}00,16{:}00)$ are buyer-initiated and seller-initiated volumes, respectively, from 14:00 to 16:00. Thus, $MM\_Imbalance$ can take values in the range $[-1, +1]$, with higher values indicating an increased likelihood of (upward) closing price manipulation.

The last three bars in Panel A of Figure 4.1 illustrate the AUROC scores of our market manipulation metrics. Larger AUROC scores indicate that our metrics are more accurate in correctly classifying the prosecuted instances of closing price manipulation. Again, the AUROC measures are likely to understate the actual accuracy of the metrics because the metrics get penalized for flagging instances of nonprosecuted closing price manipulation.

In line with Comerton-Forde and Putniņš (2011), our best performing metric is $MM\_Reversal$, with an AUROC score of 0.751.[80] When manipulators set unnaturally high closing prices on a given day, those distorted prices tend to reverse the following day. Manipulation is also detectable using abnormal buying pressure because $MM\_Imbalance$ generates an AUROC score of 0.719. Finally, because manipulators dislocate closing prices, they also cause abnormal volatility in the last hours of trading, as captured by $MM\_Volatility$, which has an AUROC score of 0.631.

In Panel B of Table 4.2, we assess the differences in our metric and index values between the prosecuted and nonprosecuted samples as an additional way of gauging whether they are useful proxies for market manipulation. Both the means and medians of the market manipulation metrics are significantly larger in the prosecuted sample. The nonprosecuted manipulation sample also contains undetected cases of closing price manipulation, which will tend to bias the results against finding a difference between the two samples and understate the ability for these metrics to flag misconduct. The average day-end period is more volatile (mean $MM\_Volatility$ is around one-third larger) in the prosecuted sample compared to the nonprosecuted sample, shows around three times larger reversals (difference in mean $MM\_Reversal$ equals 30.02), and has more than 65 times the buy–sell imbalance (difference in mean $MM\_Imbalance$ is 7.71) in the two hours leading up to the market close. Similarly,

[80] Comerton-Forde and Putniņš (2011) find that closing price manipulation causes abnormal returns approximately six times larger than normal levels that reverse the following morning.

the medians of the metrics are larger in the prosecuted sample, with the typical (median) day-end period exhibiting 30% more volatility (difference in median $MM\_Volatility$ is 11.79), as well as larger reversals and imbalances (difference in median $MM\_Reversal$ and $MM\_Imbalance$ equals 36.89 and 5.95, respectively).[81] Finally, the mean and median of the market manipulation index in the prosecuted sample are both around 30% larger (the difference in mean and median MM_Index are 12.22 and 13.03, respectively) than in the nonprosecuted sample.

4.3.3 Creating a measure of aggregate market integrity

We combine the individual metrics that pick up insider trading and closing price manipulation into indices that we can use as benchmarks of market integrity. The indices are constructed in two steps, as follows:

In step one, we create an insider trading index (IT_Index) using our three insider trading metrics and a market manipulation index (MM_Index) using our three market manipulation metrics. We assess three different approaches for combining the metrics into indices, including (i) taking the simple average of the metrics, (ii) using the coefficients from a logistic regression with prosecution as a binary dependent variable and the metrics as independent variables, and (iii) using a principal component analysis (PCA). Given that all three approaches produce similar levels of accuracy, in the interest of simplicity, we report the results using the equal-weighting approach of constructing the indices.[82]

In step two, we subtract one from the simple average of the two indices (IT_Index and MM_Index) to produce one market integrity index (MI_Index). Figure 4.2 illustrates the two-step process.

Panel B of Figure 4.1 reports the accuracy of our insider trading and market manipulation indices using the AUROC scores. The AUROC scores of the indices are 0.591 and 0.789 and are significantly different from 0.5 at the 95% and 99% significance levels, respectively. Because there are almost certainly unprosecuted cases of insider trading and market manipulation in our nonprosecuted sample, these AUROC scores are conservative and should be viewed as a lower bound of the index's accuracy.

[81] The AUROC score of $MM\_Reversal$ consistently stays above 0.72 when changing the denominator to mean($Reversal_{d-20}, \ldots, Reversal_d$) or min(max($0, R_d$), max($0, -R_{d+1}$)).

[82] The AUROC scores of these alternative approaches are within +/- 2% of our equal weighting approach.

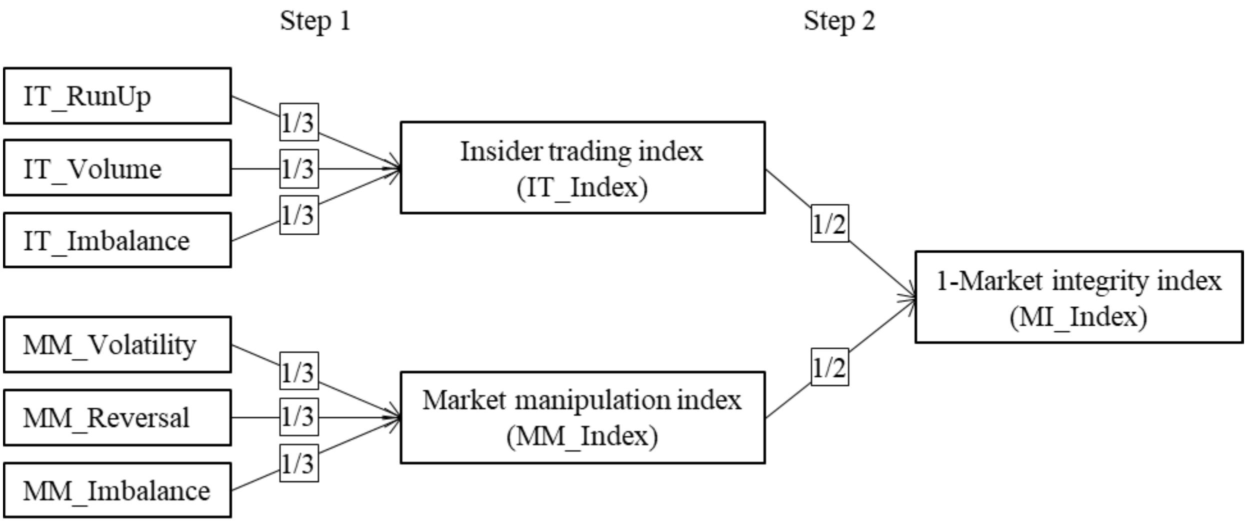

Figure 4.2
Index construction
This figure illustrates our index construction in two steps. In the first step, we combine insider trading and market manipulation metrics into an insider trading index (*IT_Index*) and a manipulation index (*MM_Index*). *IT_RunUp* measures the price run-up in the five days before an M&A announcement. *IT_Volume* measures the abnormal volume in the M&A target's stock in the five days before an M&A announcement. *IT_Imbalance* is the imbalance between buyers and sellers in the M&A target's stock in the five days before the M&A announcement. The insider trading index (*IT_Index*) is the equally-weighted average of the insider trading metrics. *MM_Volatility* measures the stock price volatility in the two hours before the market closes. *MM_Reversal* measures whether a day-end return reverses the following day. *MM_Imbalance* is the imbalance between buyers and sellers at the end of the trading day. The market manipulation (*MM_Index*) is the equally-weighted average of the market manipulation metrics. In the second step, we combine the indices into one market integrity index (*MI_Index*). The market integrity index (*MI_Index*) is the equally-weighted average of the insider trading index (*IT_Index*) and market manipulation (*MM_Index*).

4.4 Market integrity in US stock markets

In this section, we use the metrics and indices to analyze integrity in the US stock markets. This application serves two purposes. First, we gain a deeper understanding of the market integrity of one of the largest and most influential markets in the world. Second, applying our metrics to one of the most studied markets provides some further validation of the measures before extending them to the 25-country world sample in the next chapter.

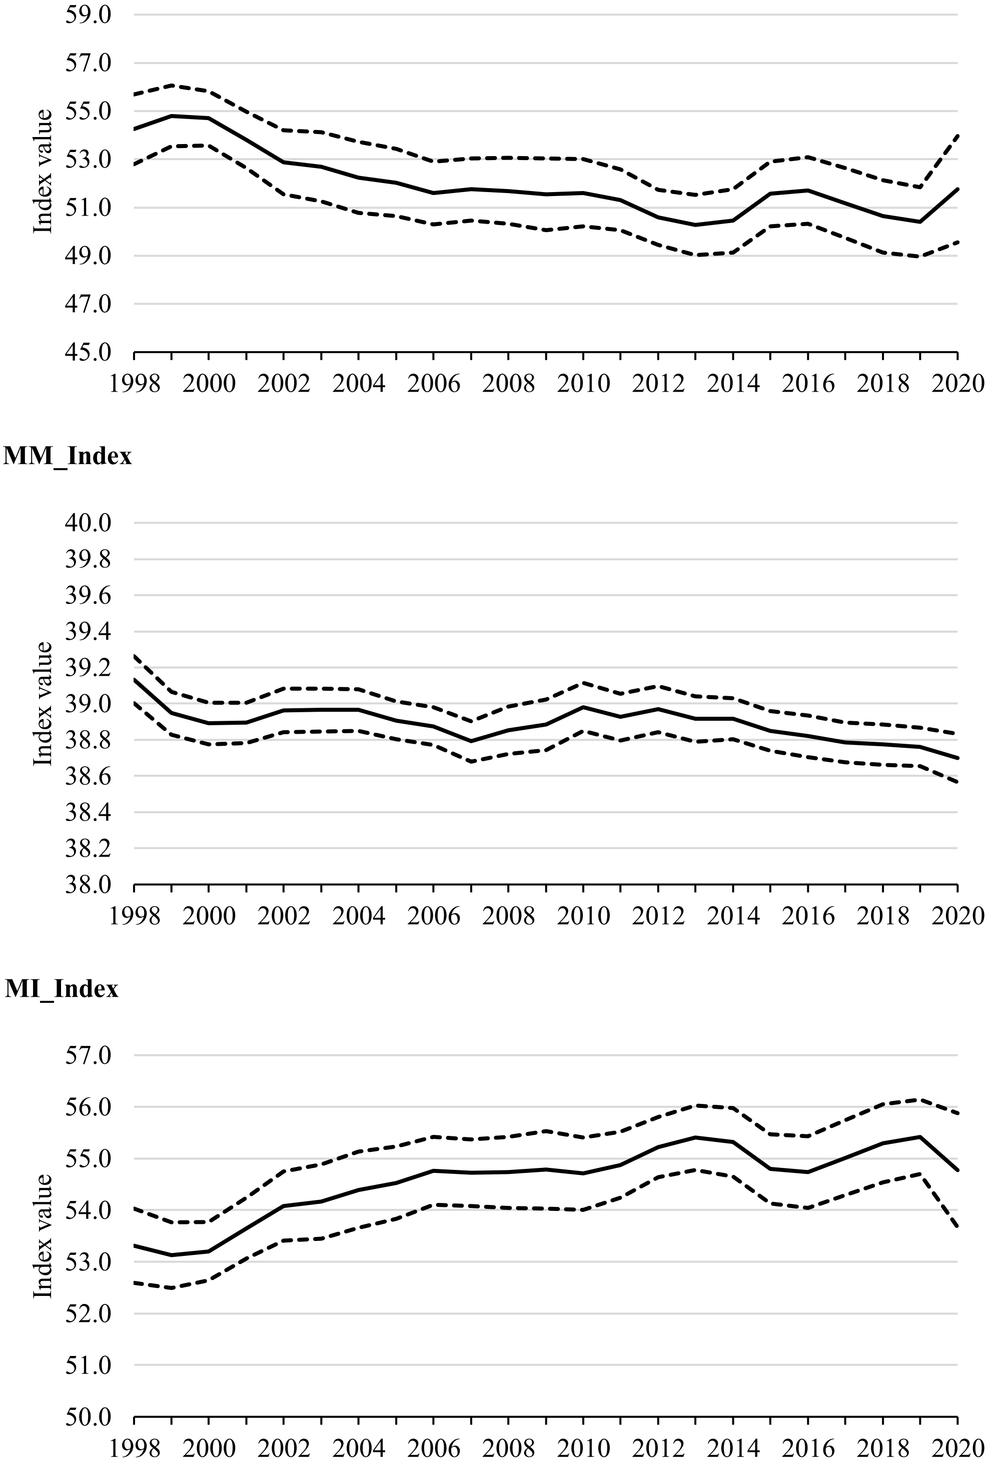

Panel A: IT_Index

Panel B: MM_Index

Panel C: MI_Index

Figure 4.3
US market integrity over time
This figureillustrates the values of our insider trading index (Panel A), market manipulation index (Panel B), and market integrity index (Panel C) in the US over time. We plot the index values with a solid black line and 95% confidence intervals with a dashed black line using double-clustered standard errors by stock and date for the market manipulation confidence intervals.

Figure 4.3 shows the level of market integrity in the US through time with 95% confidence bounds.[83] Both insider trading (in Panel A) and market manipulation (in Panel B) share similar time trends from 1998 to 2007. From 1998 to 2007, The insider trading index (IT_Index) falls by around 4.5%, the market manipulation index (MM_Index) falls by almost 1%, and the market integrity index (MI_Index) increases by around 2.7%.

The overall fluctuations in the insider trading and market manipulation indexes correspond to key legislative and regulatory/policy changes. The initial decrease in insider trading and market manipulation from 1998 to 2007 is likely the result of an increased focus on market misconduct by the US congress and strengthening of the SEC's enforcement powers. The *International Securities Enforcement Cooperation Act of 1990* enlarged the SEC's ability to address international securities issues. The act allowed the SEC to bar, sanction, or place conditions on the ability of market participants to engage in commission-regulated activities, if a foreign court had found the market participant guilty of illegal or improper conduct. The act also allowed for the confidential treatment of information from a foreign country if disclosure would violate the country's confidentiality requirements. Under the act, the SEC could also provide information to foreign and domestic authorities and accept reimbursement for any costs incurred from providing assistance. Our results are consistent with the legislation having been successful. We also see a drop in the insider trading index following the introduction of the 2011 US whistleblower program.[84] The increased risk of detection is likely to have discouraged some insider trading in the immediate years (2011 to 2014) following its implementation. There is an increase in the level of insider trading after the 2014 U.S. v. Newman ruling, which made prosecution less likely for individuals who are several links away from the original source of the private information. Ahern (2017) argues that these individuals are professional portfolio managers and our results are consistent with the notion that following the ruling they were able to trade more aggressively with a lower probability of prosecution. Finally, in 2020, insider trading spikes, which is likely because of the low conviction rate in 2019.[85] Insiders are more likely to break the law when the perceived likelihood of prosecution is lower.

[83] Because there are multiple days per stock and multiple stocks per day in our manipulation sample, we double cluster the standard errors used for the confidence intervals by stock and date. We combine the standard errors when computing the confidence intervals for the market integrity index using $SE_{GMC} = \sqrt{(0.5)^2 SE_{IT}^2 + (0.5)^2 SE_{MM}^2}$, where SE_{GMC} is the standard error of the combined market integrity index, SE_{IT}^2 is the squared standard error of the insider trading index, and SE_{MM}^2 is the squared double clustered standard error of the market manipulation index.

[84] The whistleblower program was part of the 2010 Dodd-Frank Act and promised whistleblowers an award between 10% and 30% of the monetary sanctions collected (See proposed rule release no. 34-63237 on SEC.gov).

[85] The number of people charged with insider trading was 46 in 2019 (the lowest since the Reagan administration from 1981 to 1989) down from around 100 in 2018. See *NPR* article from August 14, 2020: "Under Trump, SEC enforcement of insider trading dropped to lowest point in decades" by Tom Dreisbach.

4.5 Conclusion

Around the world, almost every financial market regulator’s mandate involves ensuring the country’s equity markets are fair and efficient. Compared with the extensive literature on market efficiency—how we measure it, what affects it, and how markets compare in efficiency—the academic literature has virtually ignored market integrity, which is much more difficult to measure. The current study develops and validates indices of insider trading and market manipulation which can be applied in equity markets around the world to measure market integrity.

We find that measures based on abnormal returns, large trading volumes, and positive order imbalances correlate with insider trading, thus serving as useful proxies, while manipulation exhibits intraday order imbalances, stock price volatility, and price reversals.

The indices developed in this chapter may be a useful tool for regulators to monitor and benchmark integrity in their markets including how it evolves through time or how it responds to various policies and regulatory activities. The indices could also be applied in real-time market surveillance, in particular if combined with other data.

This chapter takes one of the first steps in measuring market integrity, which has been largely ignored by the academic literature. We hope that our measures of integrity will facilitate further studies of market integrity, shedding more light on the drivers of market integrity and strategies to improve it.

Chapter 5

Market integrity around the world

5.1 Introduction

The mandate of regulators around the world is to maintain market efficiency and integrity in their jurisdictions.[86] The first step in achieving this is to understand their drivers and determinants. While the literature on market efficiency is extensive, market integrity is much less studied. In this chapter, we strive to fill this void and shed light on the following questions: which countries have the highest/lowest levels of market integrity? How does market integrity vary with stock characteristics? What country-level characteristics, such as wealth and culture, are associated higher integrity? How do changes in stock market structure impact market integrity? And finally, how do various regulatory actions, policies, and enforcement affect the level of market integrity?

We apply the new measures of financial market integrity from Chapter 4 on a sample of 25 countries using daily and intraday data spanning over two decades. We use the indices to examine in this global sample which stocks are the most vulnerable to insider trading and market manipulation and how economic development, income equality, and culture relate to market integrity. We then use the indices to measure how three major changes in stock market structure during the last three decades (market fragmentation, the growth in dark pools, and high-frequency trading) have affected market integrity. Finally, we measure the effect of whistleblower schemes, insider trading enforcement, penalty increases, and intercontinental cooperation on market integrity. In these tests, we use country fixed effects so that the countries that change market design or legislation are effectively the "treatment groups" while the rest of the countries serve as the control group.

Our analysis reveals a number of new findings. First, the market integrity index shows that the cleanest five countries, meaning countries with the highest levels of stock market integrity (integrity score given in parenthesis) are the US (56.12), Japan (55.85), Netherlands (55.76), Canada (55.63), and Australia (55.41),[87] and the least clean are Hong Kong (52.55), Singapore (52.43), Malaysia (51.09), Thailand (50.04), and China (43.85).

Our results also show that insider trading, measured by the insider trading index, is more likely to happen in stocks with larger market capitalization and high liquidity. One explanation is that these companies have more insiders employed who can trade on the firm's private information, and the stock's high liquidity helps insiders disguise their trades in the order flow. We also find that market

[86] See for example "The G20 Seoul Summit Leaders' Declaration" from November 2011 where G20 members commit to improving market integrity and efficiency.

[87] The UK comes in right after Australia at a close sixth with a market integrity score of (55.25).

manipulators are attracted to the opposite—they tend to seek out stocks with low market capitalization and low liquidity. First, the small size of the stocks may lower analyst coverage, lessen oversight, and lower the probability of detection. Second, manipulators may prefer low liquidity stocks where they compete with fewer trades when setting the closing price.

We also find that overall, market integrity around the world has increased during the past two decades, and although it appears the developed countries lead the improvements, emerging countries have made significant strides (especially in the late 1990s). Although our indices show that many markets have made significant progress in reducing the prevalence of insider trading, it seems only North America has significantly cut down on market manipulation levels.

The overall liquidity and low corruption levels of developed countries may be a contributing factor in their high market integrity levels: for every $1000 traded yearly per $1 of GDP, the level of market manipulation decreases by 4.14 index points and a 1 standard deviation decrease in corruption lowers insider trading by 1.42 index points (and increases market integrity by 0.81 index points). The effects are noteworthy considering the mean of the market manipulation index is 38.82 and its standard deviation is 1.07. Interestingly, when controlling for country development, corruption, and culture, the Asia-Pacific region still has a higher level of "unexplained" market manipulation than North America (0.66 index points higher). Note that when discussing changes in "index points" we mean a numerical (not a percentage) change in the level of the index.[88]

We also use our indices to measure the impact of major market microstructure events, including fragmentation of stock trading across competing exchanges or trading venues, dark trading, and high-frequency trading (HFT). We find that fragmenting stock markets decrease the market manipulation index by around 0.86 index points in North America and 1.39 in Australia, while dark trading restrictions and high frequency trading decrease the index by 1.16 and 0.42 respectively. The effects of these market changes are economically meaningful considering that the standard deviation changes of the market manipulation index is around 0.5 index points.

Finally, we regress the integrity indices on regulatory strategies, including whistleblower schemes, insider trading enforcement, penalties, and intercontinental cooperation. We find that whistleblower protection and increased fines for market misconduct tend to decrease insider trading by around 1 to 2 index points and increases the market integrity index by almost 1 index point. Our index indicates that by far the most effective insider trading deterrent is enforcement—we find that the first prosecution of insider trading in a country decreases the insider trading index by 4.71 index points and increases the market integrity index by 2.21 index points.[89] We also find that countries that implement The International Organization of Securities and Commissions (IOSCO) "best practice" directives for

[88] For example, a decrease in the market integrity index from 60% to 50% is a decrease of 10 (60-50) index points.
[89] This result is in line with Bhattacharya and Daouk (2002), who find that enforcing laws (and not just drawing them up) deters insider trading.

market integrity tend to have lower levels of market manipulation (by more than 0.5 index points) after implementing the directives.

Our results can help government institutions such as regulators, efficiently allocate scarce regulatory resources towards targeting the segments of the market that are more susceptible to insider trading and market manipulation. The results also provide some guidance as to what regulatory strategies appear to be effective around the world in improving market integrity. Finally, our results indicate that certain stock market designs may improve market integrity. Regulation can be incredibly expensive, and not all countries can afford a sophisticated regulatory body. Therefore, changes in secondary market design may be especially relevant for poorer countries.

The structure of the rest of the chapter is as follows: The next section develops testable hypotheses on the drivers of insider trading and market manipulation and the market designs and regulation that may deter it. Section 5.3 ranks countries by their integrity levels and tests the hypotheses, and Section 5.4 concludes this chapter.

5.2 Hypotheses

In the following section, we review the literature and develop hypotheses on what stocks are most vulnerable to misconduct, what country-level attributes are likely to explain cross-country differences in market integrity, how stock market design affects integrity, and how effective are various regulatory strategies in improving market integrity.

5.2.1 Stocks vulnerable to misconduct

The literature points toward three channels of stock vulnerability to market misconduct. In this section, we review the literature on the effects of firm size, liquidity, and asymmetric information on market integrity. Large (high market capitalization) companies have more analyst coverage, which discourages manipulation because detection is more likely (e.g., Comerton-Forde and Putniņš, 2014). Chung and Charoenwong (1998) argue that larger firms have more insiders and, therefore, more insider trading.

H1. Firm size increases insider trading levels, yet lowers market manipulation levels.

Insiders profit from making investments with a small price impact so that they reveal less information to other traders who may take a share in the profit by investing in the same direction or regulators who are more likely to detect the illicit activity. Therefore, insiders prefer more liquid stocks, where large trades cause small price movements.[90]

[90] See, e.g., Kyle (1985), Admati and Pfleiderer (1988), and Collin-Dufresne and Fos (2015, 2016).

Although insider traders appreciate liquidity because it acts as "camouflage" for their trading activity, closing price manipulators do not (see, e.g., Kyle, 1985; Admati and Pfleiderer, 1988; Collin-Dufresne and Fos 2015, 2016). Manipulating a stock's closing price to a desired level requires the manipulator to submit a sequence of trades in the last moments of the trading day. Therefore, closing price manipulators prefer illiquid stocks, where their trades have a higher price impact, making it easier (and less costly) to set the closing price (Comerton-Forde and Putniņš, 2011).

H2. Stock liquidity increases the level of insider trading and lowers the level of market manipulation.

Insider trading is contingent on the insider having an informational advantage, and the larger the informational gap (or level of asymmetric information) between the insider and the uninformed investors, the more profitable the information is. This is in line with Aboody and Lev (2000), who find that insider gains are higher in companies with high asymmetric information (using R&D activities as a proxy).

H3. Asymmetric information increases the level of insider trading.

5.2.2 Country development and market integrity

In this section, we develop hypotheses on the effects of country development, corruption, culture, and geographical location on market integrity. For the country's development overall and specifically stock market development, four standard development variables may be applicable: *Market capitalization to GDP*, *Stock market traded value to market capitalization*, *Stock market traded value*, and *GDP per capita*.

More developed countries have larger regulators with more resources at their disposal to detect, gather evidence, and prosecute financial crime. Bhattacharya and Daouk (2002) and La Porta, Lopez-de-Silanes, Shleifer, and Vishny (1998) note that it is not legislation but the effectiveness of enforcement in using legislation to prosecute misconduct that deters it. There is also a high correlation between the strength of the legal environment (in terms of legislation and enforcement) and size of the capital markets. Strong legal environments incentivize financiers to invest in the capital markets and help them grow (La Porta, Lopez-de-Silanes, Shleifer, and Vishny, 1997).

H4. More developed countries typically have larger regulatory bodies with more resources at their disposal to detect misconduct, enforce legislation, and prosecute misconduct. They therefore have higher levels of market integrity.

The literature shows a clear link between crime and poverty (see, e.g., Hsieh and Pugh, 1993; Fajnzylber, Lederman, and Loayza, 1998, 2002a, 2002b), especially in the type of crime studied in this thesis: nonviolent crimes (Kelly, 2000). Soares (2004), however, argues that the link between crime and development is because of the high detection rates in these countries, not the actual crimes committed. Amara and Khlif (2018) also note that corruption increases the extent of financial crime by lowering transparency.

H5. Countries with low poverty and corruption levels have higher levels of market integrity.

Country culture and market integrity

One of the key determinants discussed in the economics literature on crime is culture. The literature extensively uses Hofstede's cultural dimensions to study the effect between culture and ethical and criminal conduct, such as ethical decision making (Vitell, Nwachukwu, and Barnes, 1993), bribery (Sanyal, 2005), corruption (Husted, 1999), tax evasion (Tsakumis, Curatola, and Porcano, 2007), and money laundering (Yamen, Al Qudah, Badawi, and Bani-Mustafa, 2019). In this section, we consider three of Hofstede's culture dimensions: "power distance," "uncertainty avoidance," and "masculinity."

Power distance is the extent to which the less powerful individuals in a society accept their low standing and consider it as "normal." Therefore, we can think of the measure as an acceptance of inequality among the less fortunate. Countries with a high power distance tend to have high levels of bribery (Sanyal, 2005), corruption (Takyi-Asiedu, 1993; Husted, 1999), and tax evasion (Tsakumis, Curatola, and Porcano, 2007). These countries have paternal systems in place where superiors trade favors for loyalty from subordinates; subordinate loyalty may deter whistleblowing and lead to lower levels of market integrity.

H6. Power distance is associated with lower market integrity.

Uncertainty avoidance "indicates to what extent a culture programs its members to feel either uncomfortable or comfortable in unstructured situations" (Hofstede, 2011). The literature does not agree on the effect of uncertainty avoidance and crime. Uncertainty avoidance has a negative effect on crime, here as measured by money laundering (Yamen et al., 2019) and bribery (Sanyal, 2005), but a positive effect on tax avoidance (Tsakumis, Curatola, and Porcano, 2007) and corruption (Husted, 1999). Uncertainty-avoiding societies try to minimize the possibility of unstructured situations with strict behavioral codes and laws and rules (Hofstede, 2011). We argue that codes, laws, and rules discourage a society from engaging in criminal activity, hence leading to a higher level of market integrity.

H7. Uncertainty avoidance is associated with higher market integrity.

Hofstede (2011) defines societal masculine traits as more “assertive and competitive” than feminine traits, which are more “modest and caring.” Higher levels of masculinity are associated with financial crime (Yamen et al., 2019), bribery (Sanyal, 2005), and corruption (Takyi-Asiedu, 1993; Husted, 1999). Sanyal (2005) argues that masculine societies are more aggressive in the pursuit of success and achievement, which leads to corrupt conduct. We believe that this pursuit similarly affects our types of misconduct (insider trading and market manipulation) and lower market integrity.

H8. Masculinity is associated with lower market integrity.

Geographic region effects are common in world crime studies, such as the studies by Fajnzylber, Lederman, and Loayza (1998, 2002a, 2002b). Dissanaike and Lim (2015) argue insider trading and market manipulation is particularly prevalent in Asian markets where regulators lack necessary resources to enforce legislation, whistleblower programs are scarce, and insider trading is culturally more “acceptable”.

H9. There is a region-specific effect in Asia leading to a lower level of market integrity.

5.2.3 Stock market design and market integrity

The last three decades have seen three major market microstructure changes in stock markets—market fragmentation, high-frequency trading (HFT), and dark trading. Market fragmentation implies a market where several exchanges (as opposed to one or few) operate. High-frequency trading occurs when exchanges allow HFT firms to place their computer in the same premise as the exchange’s computers (also called colocation services) thereby lowering the HFT firms’ transaction latency. Market may also allow dark trading where the venue (the dark pool) avoids displaying any resting buy or sell orders before matching occurs, thereby avoiding pre-trade transparency.

Market fragmentation increases competition among exchanges, which increases liquidity (O’Hara and Ye, 2011). Liquidity decreases the ability for a market manipulator to influence prices[91] but the increased trading activity allows an insider to better camouflage their trades.[92]

H10. Market fragmentation increases market liquidity, which increases insider trading and decreases market manipulation.

[91] See Comerton-Forde and Putniņš (2014).

[92] See, e.g., Kyle (1985), Admati and Pfleiderer (1988), and Collin-Dufresne and Fos (2015, 2016).

Dark trading has attracted considerable attention from regulators and academics.[93] A major concern is that the process in which dark pools obtain prices from the lit market[94] facilitates manipulation strategies (Ye 2012; Klöck, Schied, and Sun, 2017). Mittal (2008) describes manipulation strategies where the manipulator (i) determines and order imbalance in the dark pool by submitting small orders (also called "fishing"), (ii) manipulates the lit market by trading in the direction of the dark pool order imbalance, and (iii) buys (or sells) the dark pool order imbalance at a discount (premium).[95]

H11. Dark trading increase market manipulation.

High-frequency traders are often accused by market participants of having negative impacts on markets,[96] even though the literature shows that high-frequency traders tend to increase liquidity on average, which may reduce market manipulation. For example, Aitken, Cumming, and Zhan (2015) note that as liquidity providers, high-frequency traders may discourage closing price manipulation.

H12. High-frequency trading increase market liquidity, which increases insider trading and decreases market manipulation.

5.2.4 Regulation and market integrity

Regulatory strategies, including whistleblowing schemes, insider trading enforcement, penalty increases, and cross-jurisdiction cooperation, are likely to impact market integrity. Countries around the world implement and change whistleblowing programs to help detect various forms of misconduct. To create incentives, whistleblowing schemes usually compensate the whistleblower with monetary awards or leniency if the whistleblower was involved in the misconduct themselves. Compensation is sometimes astounding in size—the US SEC awarded a record high $50 million to a whistleblower in April 2021.[97] As another example of incentives, the Australian financial market regulator, ASIC, announced a "get-out-of-jail free card,"[98] where guilty individuals gain legal immunity if they provide information on their accomplices.

Putniņš and Sauka (2015) measure the shadow economy in the Baltic countries, arguing that the decision to evade taxes is consistent with rational choice models—evaders base their decision on

[93] For more on the regulatory concerns regarding dark trading, see the IOSCO 2011 final report "Principals for Dark Liquidity," available at iosco.org.

[94] Lit markets are (as opposed to dark pools) pre-trade transparent as the orders are displayed in the order book.

[95] For a more elaborate explanation of dark pool manipulation strategies, see Mittal (2008).

[96] The SEC for example, blamed the May, 2010 flash crash on high-frequency trading (See September, 2010 report, prepared by the US CFTC and SEC "Findings Regarding the Market Events of May 6, 2010").

[97] See SEC press release 2021-62 "SEC awards over $50 million joint whistleblowers" from https://www.sec.gov/page/news.

[98] See February 24, 2021 *Australian Financial Review* article "'Rush to ASIC': White-collar whistleblowers to get legal immunity".

the benefits (tax savings) and costs (penalty and probability of detection). They argue that whistleblowing is effective in decreasing the Baltic shadow economy because it increases the evader's perceived probability of detection and deters tax evasion.

H13. Whistleblowing leads to higher levels of market integrity because it increases the probability of detection, thereby discouraging criminal activity among potential offenders.

Around the world, the laws against misconduct are plentiful; however, they are not always enforced. For example, Russia has only issued notices of seven insider trading investigations since laws were enacted and has yet to make any prosecutions (Anderson, 2021).[99] Bhattacharya and Daouk (2002) show that it is not laws but the enforcement of laws that deter misconduct. The enforcement of insider trading laws may be a prime contributor of cleaner financial markets.

H14. Enforcing insider trading laws sends a signal to potential offenders that enforcement agencies are able and willing to enact legislation. This lowers insider trading levels by increasing the perceived likelihood of detection among potential offenders.

Bris (2005) finds that it is not only the enforcement of insider trading laws but also the toughness that decreases insider trading profits. Putniņš and Sauka (2016) also find that higher penalties and perceived probability of detection lowers financial crime, such as tax evasion and misreporting.

H15. Increasing penalties increases market integrity by increasing the costs associated with criminal activity.

Finally, regulatory agencies vary in their efficiency and capability—resources that can be shared. For example, China accepted technical advice from the US SEC when developing their insider trading enforcement regime (Anderson, 2021). This ability to cooperate cross-jurisdictionally is especially relevant in today's globalized markets, where market misconduct occurs across jurisdictions. We use an event in 2011, in which G20 member countries agreed to cooperate on market integrity, to study the effect of cross-jurisdictional cooperation on market integrity.

[99] For information on Russian detected cases of insider trading and market manipulation, see http://old.cbr.ru/eng/finmarket/inside_detect/.

H16. Intercontinental cooperation lets enforcement agencies pool their resources and increases market integrity, especially among less developed countries, where the marginal effect of added resources is higher.

5.3 Data

In testing the hypotheses above, we expand our US sample from Chapter 4 to include other jurisdictions around the world. This section explains how we collect and filter the data.

5.3.1 M&A event and closing price data

In Chapter 4, we collected a sample of US M&A events and closing prices to measure insider trading and closing price manipulation in the US. We expand this sample to include 24 other jurisdictions in Asia-Pacific, Europe, and North America.[100] For each jurisdiction, we collect M&A events and closing prices from 1996 to 2020. For the M&A events, we include all M&A announcements in each country, but for the closing prices, we restrict our sample to include 200 randomly sampled stocks in each country-year. When sampling the stocks, we follow the same data requirements as described in Section 4.2.3. However, if for a given year a country has fewer than 200 listed stocks that meet our sampling criteria, then we include only those stocks that meet our criteria.

5.3.2 Stock and country characteristics data

Testing the hypotheses from Section 5.2 also requires additional stock- and country-level data, including data on country development, corruption, culture, and events data on changes in market design and regulatory legislation. We collect stock-level explanatory variables from the Refinitiv Tick History database, including stock market capitalization (*Market cap*), daily traded volume (*Volume*), daily stock price standard deviation (*Volatility*), and the stock's country of origin. We also collect country-level variables from the World Bank, including *Market capitalization to GDP*, *Trade value*, *Trade value to market cap*, *GDP per capita*, and *Control of corruption*. Hofstede (1997, 2001) uses country-level culture variables, which are available on his website[101], and from there, we collect *Power distance*, *Uncertainty avoidance*, and *Masculinity*. Finally, we collect the dates of market design changes and regulatory policies from regulatory reports, academic papers, and news articles. These include the dates of market fragmentation (*Frag NA*, *Frag EU*, and *Frag AU*), dark trading restrictions (*Dark*), colocation services (*HFT*), whistleblower schemes (*Whistleblowing*), first insider trading prosecutions (*IT enforcement*), changes in misconduct penalties (*Penalty increase*), and cooperative

[100] *Asia-Pacific* includes Australia, China, Hong Kong, India, Japan, Malaysia, Singapore, South Korea, Taiwan, and Thailand; *Europe* includes France, Germany, Israel, Italy, Netherlands, Norway, Poland, Sweden, Switzerland, and the UK; and *North America* includes the US and Canada.

[101] https://geerthofstede.com/research-and-vsm/dimension-data-matrix/.

efforts across jurisdictions (*Intercontinental cooperation*). Table 5.1 provides definitions for all of these variables and additional information about sources.

Table 5.1
Definitions of explanatory variables
This table defines the variables that we use in the analysis. The definitions for the insider trading and market manipulation metrics are in Section 4.3.1 and Section 4.3.2, respectively.

Variable	Definition
Panel A: Stock characteristics	
Market cap	The stock's average daily market capitalization in millions of dollars. Market capitalization is the share price multiplied by the number of outstanding shares. Sourced from Refinitiv Tick History database.
Volume	The stock's average daily traded volume in millions of dollars. Sourced from Refinitiv Tick History database.
Volatility	The annual average of the daily standard deviation of the five-minute intraday stock price. Sourced from Refinitiv Tick History database.
Panel B: Country-level drivers	
Market cap to GDP	The total market capitalization of all publicly listed companies in the country in thousands of dollars divided by its GDP in dollars, measured annually. Sourced from the World Bank's DataBank (databank.worldbank.org).
Trade value to market cap	The country's yearly *Trade value* in thousands of dollars divided by its *Market cap* in dollars. Sourced from the World Bank's DataBank (databank.worldbank.org).
Trade value	The yearly value of shares traded in the country in trillions of dollars. Sourced from the World Bank's DataBank (databank.worldbank.org).
GDP per capita	The country's GDP in thousands of dollars divided by the country's population. Sourced from the World Bank's DataBank (databank.worldbank.org).
Control of corruption	Measures the perception to which public power is used for private gain in units of a standard normal distribution, i.e., ranging from approximately -2.5 to 2.5. Sourced from the World Bank's DataBank (databank.worldbank.org).
Power distance	One of Hofstede's culture dimensions, measuring the extent to which the less powerful members of organizations and institutions accept and expect that power is distributed unequally. Hofstede's dimensions range from 0 to 100 and we divide the values by 10.
Uncertainty avoidance	One of Hofstede's culture dimensions, measuring society's tolerance for ambiguity and unstructured situations. Societies with high levels of uncertainty avoidance tend to have strict laws and rules and be less accepting of different opinions, religions, and philosophies. Hofstede's dimensions range from 0 to 100 and we divide the values by 10.
Masculinity	One of Hofstede's culture dimensions, measuring preference in society for achievement, heroism, assertiveness, and material rewards for success as opposed to cooperation, modesty, caring for the weak, and quality of life. Hofstede's dimensions range from 0 to 100 and we divide the values by 10.
Asia-Pacific	Dummy variable that is equal to one if the country is Australia, China, Hong Kong, India, Japan, Malaysia, Singapore, South Korea, Taiwan, or Thailand.
Europe	A dummy variable that is equal to one if the country is France, Germany, Israel, Italy, Netherlands, Norway, Poland, Sweden, Switzerland, or the UK.
North America	A dummy variable that is equal to one if the country is the US or Canada.
Other	A dummy variable that is equal to one if the country is Brazil, Russia, or South Africa.

Table 5.1 (continued)
Definitions of explanatory variables

Variable	Definition
Panel C: Market structure	
Frag_NA	Dummy variable that is equal to one if the country is the US and is the year after the US market was fragmented or the country is Canada and is the year after the Canadian market was fragmented. RegNMS fragmented the US market in 2005 and Chi-X entered (and fragmented) the Canadian market in 2008.
Frag_EU	Dummy variable that is equal to one if the country is a European Union (EU) member state and is the year after MiFID I fragmented the European market. MiFID I ended the "concentration rule" in 2007 that had until then prevented fragmentation.
Frag_AU	Dummy variable that is equal to one if the country is Australia and is the year is after Chi-X entered (and fragmented) the Australian market in 2011.
Dark trading restrictions	Dummy variable that is equal to one if the country is Australia or Canada and the year is after the country introduced restrictions on dark trading. Canada and Australia introduced minimum price improvement rules restricting dark trading in October 2012 and May 2013, respectively.
HFT	Dummy variable that is equal to one when a country's stock market allows colocation of trading servers with the exchange servers (considered an enabler of high-frequency trading, HFT) and zero otherwise. Sourced from Boehmer, Fong, and Wu (2020), found in the last column of Table G1 in Appendix G.
Panel D: Regulation and enforcement	
Whistleblowing	Dummy variable that is equal to one after the country has introduced substantial whistleblower protection regulation. The dates and regulations used are in Table G2 of Appendix G.
IT enforcement	Dummy variable that is equal to one after the first enforcement of insider trading laws. Sourced from Bhattacharya et al. (2002).
Penalty increase	Dummy variable that is equal to one in EU countries after the EU consolidated insider trading prison sentences in 2014 for all member states.
Intercontinental cooperation	Dummy variable that is equal to one in G20 countries after they started implementing IOSCO directives in October 2011. For more information on the directives, see IOSCO report "G20/FSB Recommendations related to Securities Markets."

5.4 Market integrity around the world and over time

Having validated the insider trading, market manipulation, and market integrity metrics in the previous chapter, we now apply the metrics to 25 countries. In Table 5.2 below, we group the indices by two periods of around ten years: one recent period spanning 2009 to 2020 (*Recent period: 2009–2020*) and one early period spanning 2008 to 1996 (*Early period: 1996–2008*).

Table 5.2
Market integrity in the cross-section of countries
This table ranks countries by market integrity from the highest to lowest integrity. The first five columns (“*2009–2020*”) rank market integrity in the latest part of our sample (2009–2020), and the last five columns (“*1996–2008*”) rank market integrity in the first part of our sample (1996–2008). For each period, we report the market integrity index (“*MI*”), insider trading index (“*IT*”), and market manipulation index (“*MM*”). All index values are in percentages.

2009–2020					1996–2008				
Rank	Country	MI	IT	MM	Rank	Country	MI	IT	MM
1	USA	56.12	51.05	36.72	1	Sweden	56.64	48.16	38.57
2	Japan	55.85	50.07	38.23	2	South Africa	54.76	51.81	38.67
3	Netherlands	55.76	50.96	37.51	3	USA	54.62	53.23	37.53
4	Canada	55.63	50.56	38.18	4	Japan	54.46	52.22	38.86
5	Australia	55.41	51.07	38.11	5	Norway	54.45	52.17	38.94
6	UK	55.25	51.60	37.90	6	Poland	54.02	53.48	38.48
7	Israel	55.25	50.58	38.92	7	Germany	53.97	54.29	37.77
8	Italy	55.20	51.53	38.07	8	Israel	53.96	52.98	39.10
9	South Africa	55.13	51.47	38.27	9	Australia	53.94	52.76	39.36
10	Norway	55.00	50.53	39.48	10	Canada	53.86	53.27	39.01
11	Switzerland	54.97	50.57	39.49	11	France	53.67	53.53	39.13
12	Sweden	54.92	50.99	39.16	12	Netherlands	53.59	54.45	38.36
13	Germany	54.90	51.87	38.33	13	UK	53.51	54.23	38.75
14	France	54.89	51.77	38.46	14	Brazil	53.00	55.22	38.79
15	Poland	54.16	52.55	39.13	15	Russia	52.97	55.95	38.11
16	South Korea	53.80	54.15	38.26	16	Italy	52.86	56.98	37.30
17	Russia	53.64	53.70	39.03	17	Switzerland	52.55	55.32	39.59
18	Brazil	52.80	56.16	38.24	18	South Korea	52.52	56.35	38.61
19	Taiwan	52.70	56.43	38.17	19	Taiwan	52.37	57.17	38.09
20	India	52.64	55.69	39.04	20	Hong Kong	52.36	54.30	40.97
21	Hong Kong	52.55	54.09	40.82	21	China	52.00	56.15	39.86
22	Singapore	52.43	54.85	40.29	22	Thailand	51.89	56.60	39.62
23	Malaysia	51.09	56.71	41.11	23	India	51.85	57.77	38.53
24	Thailand	50.04	59.83	40.08	24	Singapore	51.32	58.25	39.10
25	China	43.85	72.53	39.77	25	Malaysia	50.67	58.07	40.58

Wealthy, developed economies stand out as countries with high levels of integrity. These countries can afford to invest heavily in regulatory bodies but also in a sophisticated market design.[102] Market liquidity may also be an important factor; the US ranks first in the world in terms of market integrity, due to its low level of market manipulation. The US is by far the most liquid market in the world with the average yearly traded value per listed stock at $6 billion in the last 25 years. High liquidity makes markets more resilient to market manipulation because the manipulator must compete against more trades when attempting to set the closing price (Comerton-Forde and Putniņš, 2011).

[102] The resources available to the US SEC are among the highest with a regulator budget of more than $1 billion a year for the last ten years and $1.815 billion in 2020 (see www.sec.gov).

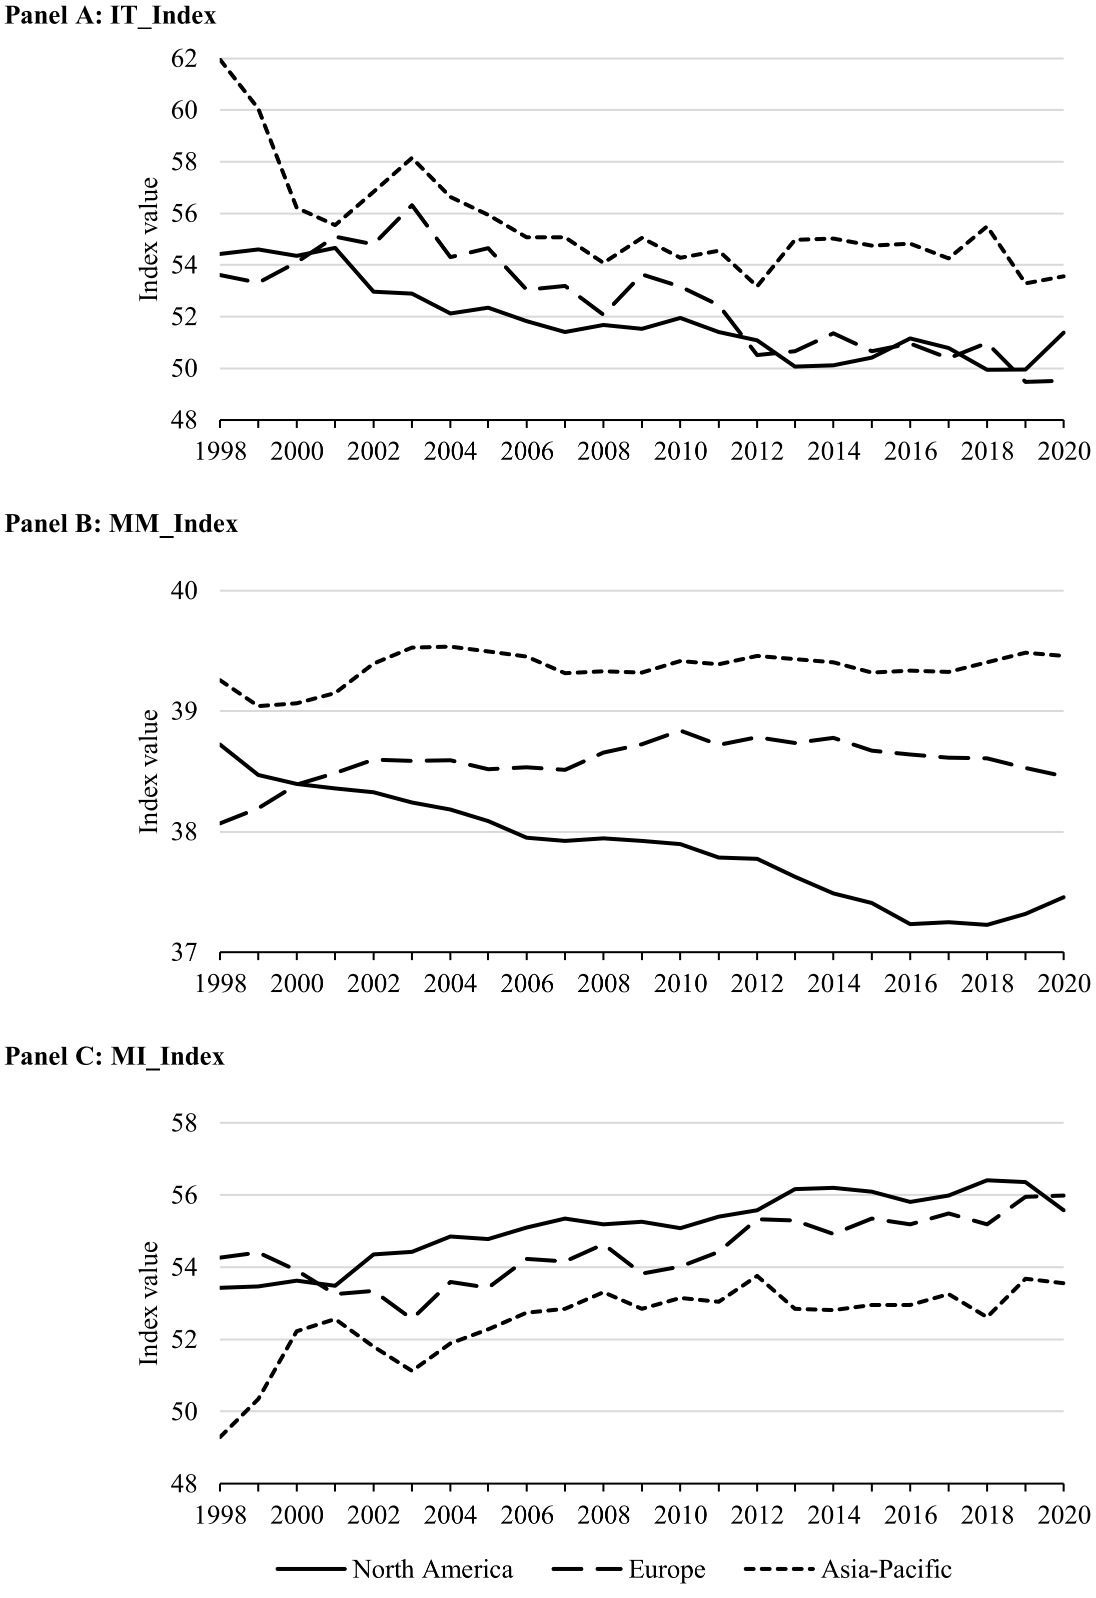

Panel A: IT_Index

Panel B: MM_Index

Panel C: MI_Index

Figure 5.1
Market integrity over time by location
This figure illustrates the values for our insider trading index (Panel A), market manipulation index (Panel B), and market integrity index (Panel C) for our world sample over time. *Asia-Pacific* includes Australia, China, Hong Kong, India, Japan, Malaysia, Singapore, South Korea, Taiwan, and Thailand; *Europe* includes France, Germany, Israel, Italy, Netherlands, Norway, Poland, Sweden, Switzerland, and the UK; and *North America* includes the US and Canada. We report all index values in percentages.

We group the countries into geographical locations and plot them over time to see how market integrity has evolved. Figure 5.1 shows that Asia-Pacific has higher insider trading levels than Europe and North America. Insider trading in Asia-Pacific drops substantially from 1998 to 2006 by about 7 index points (from 62% to 55%). Europe and North America also experience a decrease in insider trading, although the drop is not as substantial. The large drop in insider trading in the Asian countries is likely because of their rapid development from the mid-1990s to the late 2000s.[103] As these countries develop, they also become more sophisticated in cracking down on financial misconduct (Brownbridge and Kirkpatric, 2000). North America and Europe have also advanced their economies over the last twenty years but at a slower pace than Asia.

The results are similar when it comes to market manipulation: North America maintains the lowest market manipulation index from 2001 to 2020, Europe ranks second, and finally, Asia-Pacific has the highest level of market manipulation measured by the index. North America is the only region that appears to have lowered its levels of market manipulation throughout the sample period. The largest North American decrease in manipulation looks to happen after the Dodd–Frank Act of 2010. The law included new protections and rewards for whistleblowers of between 10% and 30% of collected monetary sanctions. The whistleblower program may have improved the SEC’s ability to investigate market manipulation.

Combining insider trading and market manipulation into our market integrity index, we see that although North America tends to have higher market integrity than Europe, they are very close. The market integrity index is the average of the insider trading index and the market manipulation index subtracted from 1. An increase (decrease) in the value of the market integrity index therefore corresponds to a(n) decrease (increase) in the average of the insider trading and market manipulation indices. Market integrity in Asia-Pacific is the lowest, but the region has seen the highest increase in market integrity, especially from 1997 to 2000, when the index increases by 3 index points (from a score of around 49% to 52%). This is likely because of the significant economic advancement among the Asian countries in 1997.

[103] Hong Kong, Singapore, South Korea, and Taiwan transitioned from “developing” countries to “developed” in 1997. China, India, Malaysia, and Thailand became newly industrialized countries (NICs) in the late part of the 2000s.

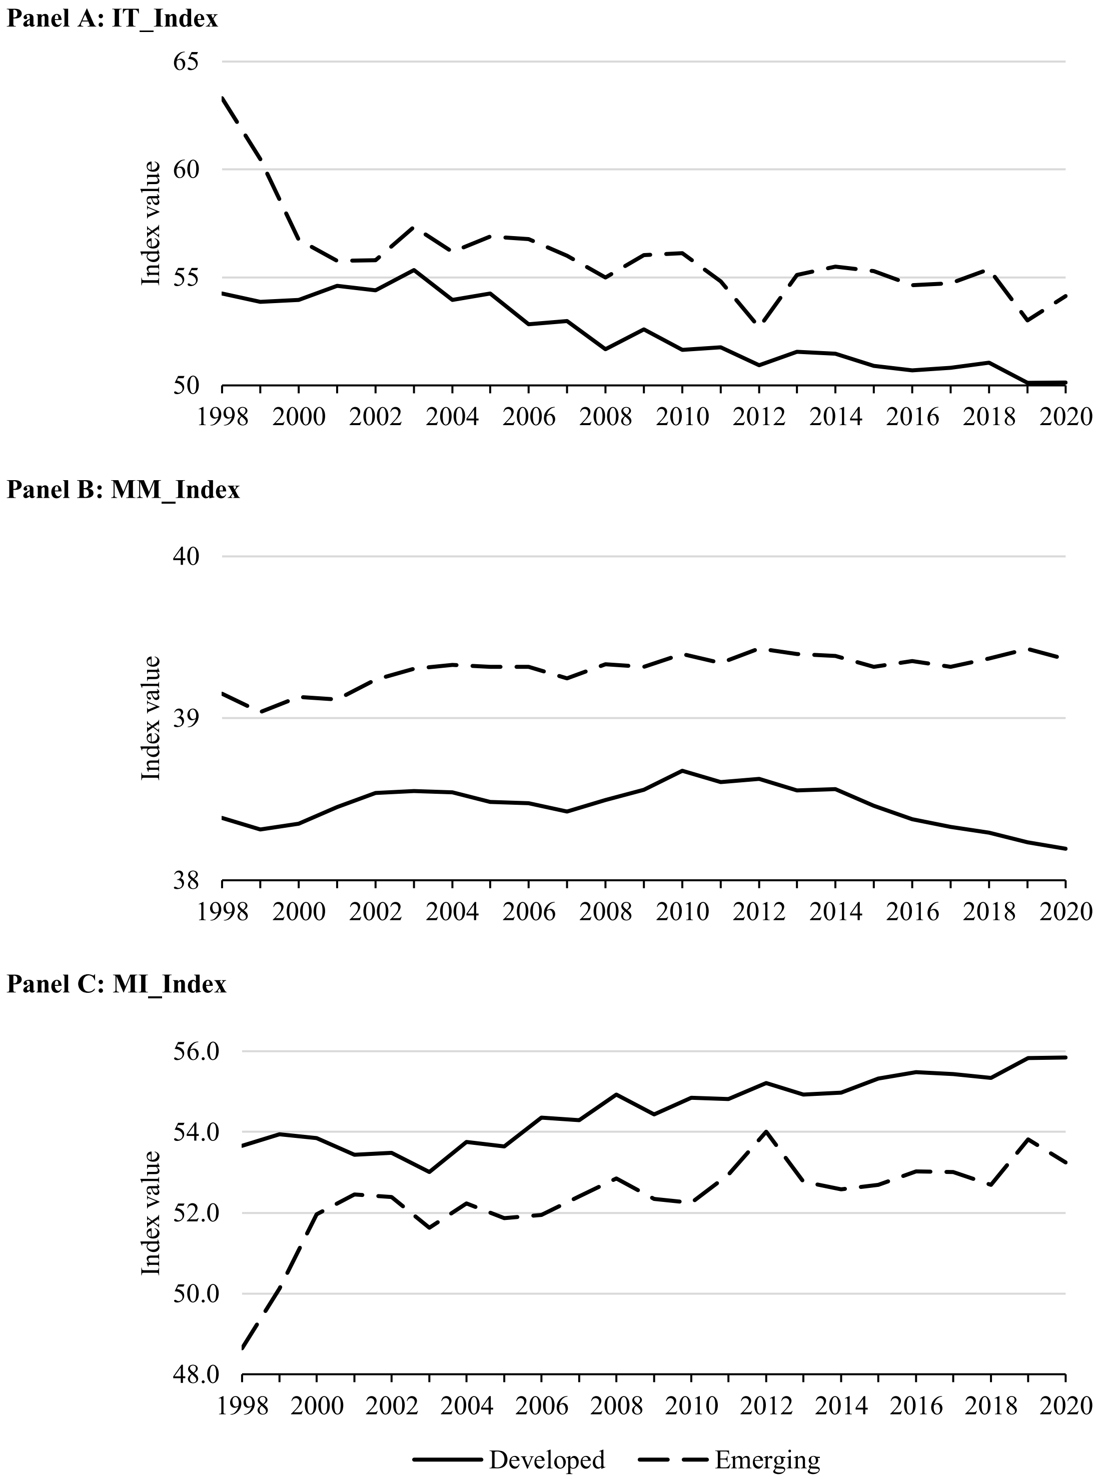

Figure 5.2
Market integrity over time by country development
This figure illustrates the values for our insider trading index (*IT_Index*), market manipulation index (*MM_Index*), and market integrity index (*MI_Index*) in developed and emerging countries over time. The *Developed* countries are Canada, US, Australia, France, Germany, Norway, Sweden, Switzerland, UK, Japan, Netherlands, Italy, Poland, and Russia. The *Emerging* countries are China, Hong Kong, South Korea, Taiwan, India, Malaysia, Singapore, Israel, Thailand, South Africa, and Brazil. We report all index values in percent.

Figure 5.2 groups countries by level of development, as opposed to by geographic regions.[104] Developed countries are more sophisticated regarding market surveillance and have more resources available to enforce legislation. Emerging countries also have higher levels of corruption, lessening enforcement agencies' effectiveness.[105] Panel C shows that market integrity in developed countries shows a high (and steadily rising) level of integrity, mostly owing to a decreasing level of insider trading (Panel A). Market manipulation (Panel B) is more prevalent in emerging countries and they have had little success in decreasing it compared with the developed countries, which have seen a slight decrease over the past ten years (2010 to 2020).

Market manipulation in developed countries has dropped significantly after 2010. Most countries included in the "developed" category are G20 members, who started implementing IOSCO directives to improve market integrity in 2011. The market manipulation index drops substantially in 2014, which is when IOSCO reported that a significant majority (22 jurisdictions) had implemented the directives.[106]

5.5 Drivers of market integrity

5.5.1 Stock characteristics

What types of stocks are more prone to insider trading and market manipulation? To answer this question, we start by analyzing how the market integrity indices vary across quintiles of market capitalization (company size), liquidity (daily traded volume), and volatility (standard deviation of returns). Figure 5.3 illustrates the results.

[104] The grouping comes from the International Monetary Fund (IMF). See the "*World Economic Outlook Databases*" at imf.org/en/Publications.

[105] According to the World Bank's *Control of Corruption* index, the developed countries in our sample are about 30% less corrupt.

[106] For more information on the recommendations and implementation of the IOSCO directives, see the IOSCO final report "G20/FSB Recommendations related to Securities Markets."

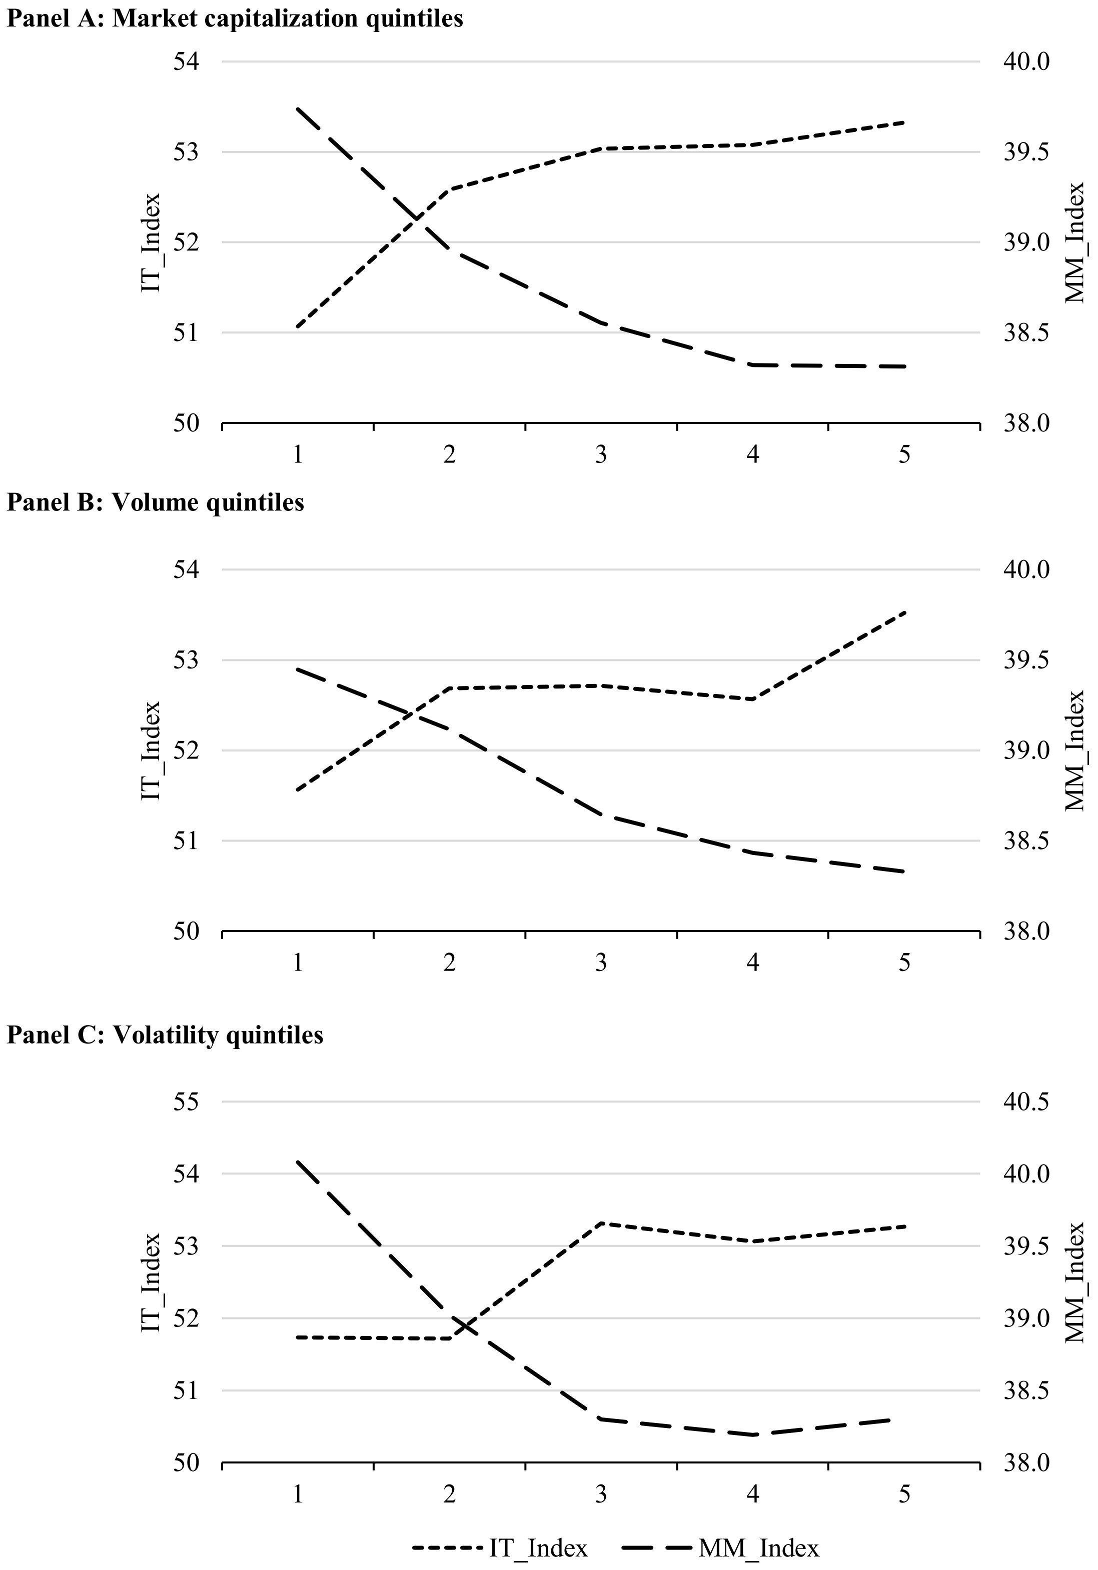

Panel A: Market capitalization quintiles

Panel B: Volume quintiles

Panel C: Volatility quintiles

Figure 5.3
Market integrity across stock characteristics
This figure illustrates the market integrity indices for stocks split into quintiles by three stock characteristics including market capitalization (Panel A), average daily traded volume (Panel B), and average daily volatility (Panel C). The quintiles are calculated on a pooled dataset consisting of all 25 jurisdictions from Section 5.3, where the unit of observation is stock-day. The left-hand-side (right-hand-side) axis values are the average of the insider trading index (market manipulation) index values by quintile.

Figure 5.3 Panel A shows that insider trading is more pervasive in large market capitalization stocks (IT_Index increases monotonically from 51% in the first quintile to above 53% in the fifth). This is consistent with hypothesis $H1$ and Chung and Charoenwong (1998) who argue that larger firms have more insiders and, therefore, more insider trading. In Panel B, we also see that there is more insider trading in firms with a higher average daily volume. Insiders typically prefer trading in liquid firms because their trades have a smaller impact and reveal less private information.[107] This means that they can submit more trades and extract higher profits with reduced detection risk. Consistent with $H3$, insider trading is also more common in volatile stocks (Panel C), where the level of asymmetric information is higher (the IT_Index is above 53% in quintiles three to five and below 52% in quintiles one and two).[108] In these stocks, the insider's private information is more valuable because there is greater uncertainty about the stock's underlying value.

The results support hypotheses $H1$, $H2$, and $H3$. Stocks most susceptible to market manipulation have characteristics that are different, often opposite, to those of insider trading. Small (low market capitalization), illiquid (low trading volume), and stable (low daily volatility) stocks are the most susceptible to manipulation. Small stocks with low liquidity (Panels A and B) have higher levels of market manipulation in the first quintile (MM_Index is above 39.5%) than in the fourth and fifth quintiles (MM_Index is below 38.5%). Manipulators prefer illiquidity stocks because it is easier to impact the price. The small size (or market capitalization) of these stocks means that the analyst coverage is lower and that manipulation is less likely to be detected (e.g., Comerton-Forde and Putniņš, 2014).

In summary, Figure 5.3 suggests that insider trading is more likely to happen in large, liquid, and volatile stocks, while small, illiquid, and stable stocks seem to attract market manipulation.

To test these relations in a multivariate setting, we regress the indices on market capitalization, volume, and volatility and account for country, year, and industry variation with fixed effects. Table 5.3 reports the results. In regression model (4), both *Market cap* and *Volume* are statistically significant; a 1% increase in market capitalization increases the level of the insider trading index by almost 0.4 index points (*Market cap* coefficient is 0.35) and 0.2% (*Volume* coefficient is 0.19), respectively. Similarly, when regressing manipulation in (8), both *Market cap* and *Volume* are significant at the 99% level, with negative coefficients; a 1% increase in the stock's market capitalization and volume decreases the manipulation index by about 0.10 index points (*Market cap* coefficient is -0.09) and around 0.15 index points (*Volume* coefficient is -0.17).

[107] This supports *H3*. The notion that liquidity helps insiders hide their informational advantage is in line with Kyle (1985), Admati and Pfleiderer (1988), and Collin-Dufresne and Fos (2015, 2016).

[108] Van Ness, Van Ness, and Warr (2001) and Wang (1993) use stock price volatility as a proxy for asymmetric information.

Table 5.3
How market integrity varies with stock characteristics
This table reports results of regressions in which the dependent variables are the insider trading index ("*IT_Index*") and market manipulation index ("*MM_Index*") and the independent variables are the natural log of stock characteristics. *Market cap* is the company market capitalization, *Volume* is the daily traded dollar volume, and *Volatility* is the standard deviation of returns (as defined in Table 5.1). The regressions are on a global stock-year sample including the 25 jurisdictions from Section 5.3 from 1996 to 2020. Numbers not in brackets are coefficient estimates. Numbers in brackets are t-statistics computed using standard errors double clustered by country and year. Significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Explanatory variables	Dependent variable							
	IT_Index				MM_Index			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln(Market cap)	0.32***	0.27**	0.32***	0.35***	-0.14*	-0.14**	-0.09***	-0.09***
	(3.15)	(2.47)	(2.94)	(3.01)	(-1.96)	(-1.97)	(-4.37)	(-4.63)
ln(Volume)	0.03	0.07	0.14*	0.19**	-0.07	-0.07	-0.18***	-0.17***
	(0.35)	(0.95)	(1.80)	(2.15)	(-1.08)	(-1.13)	(-5.52)	(-5.48)
ln(Volatility)	-0.30**	-0.28**	0.01	-0.08	0.07	0.08	-0.05	-0.05
	(-2.18)	(-2.02)	(0.06)	(-0.78)	(0.69)	(0.75)	(-0.95)	(-0.99)
Intercept	45.84***	47.14***	45.43***	43.46***	42.87***	42.89***	42.09***	41.82***
	(32.36)	(31.45)	(33.99)	(23.90)	(35.34)	(35.20)	(78.36)	(77.09)
Industry FE		Yes	Yes	Yes		Yes	Yes	Yes
Country FE			Yes	Yes			Yes	Yes
Year FE				Yes				Yes
R-squared	1%	1%	2%	4%	0%	0%	2%	2%
Observations	13,122	13,122	13,122	13,122	20,658,518	20,658,518	20,658,518	20,658,518

5.5.2 Country-level drivers

To understand what drives cross-country differences in market integrity, we regress the misconduct and integrity indices on variables that proxy countries' development, corruption, culture, and location. Table 5.4 reports the results.

Starting with the development variables, *Market cap to GDP* is a widely used indicator of stock market development. Interestingly, we see that more developed markets according to this metric have a tendency to have lower levels market integrity. An extra $1000 of market capitalization in the economy per $1 of GDP, increase the insider trading and market manipulation indices by 1.74 and 2.22 index points respectively, and drop the market integrity index by 1.98 index points. The magnitudes of the coefficients may also be interpreted by their impacts on the variability of the indices: A standard deviation in *Market cap to GDP* is associated with 4%, 36%, and 10% of a standard deviation in the insider trading index, market manipulation index, and market integrity index respectively.

Looking at *Trade value*, we get a sense of the overall activity or "size" of the country's financial market. The results support *H4* in that larger markets have lower levels of market manipulation, which translates into higher level of market integrity—a $1 trillion increase in the yearly traded dollar value drops the market manipulation index by 0.04 index points and increases the market integrity index by

0.03 index points. A standard deviation in *Trade value* accounts for 24% of the standard deviation in the market manipulation index and 5% of the standard deviation in the market integrity index.

GDP per capita (or income level) increases market integrity both in terms of insider trading and market manipulation. Wealthier countries have more resources to enforce legislation, thus discouraging misconduct, which is in line with *H4* and Bhattacharya et al. (2002), who find that it is the enforcement of legislation (and not the existence of legislation) that discourages insider trading. An increase of the GDP per person of $1 thousand lowers the market manipulation index by 0.07 index points and increases the market integrity index by 0.04 index points. A standard deviation in *GDP per capita* accounts for almost 1.5 standard deviations in the market manipulation index and one fourth of a standard deviation in the market integrity index.

We also look at the overall perception of the country's control over corruption. The coefficients on *Control of corruption* support *H5* and show that more control over (or lower levels of) corruption leads to cleaner markets because 1 standard deviation lowers insider trading by 1.42 index points and increases market integrity by almost 1 index point (coefficient equals -0.81). When compared to the variation of the index, a standard deviation in *Control of corruption* is associated with 1.3 standard deviations in the market manipulation index and 0.22 standard deviations in the market integrity index.

When adding our *Corruption control* variable, *GDP per capita* becomes insignificant because the two variables are highly correlated.[109] The correlation between country corruption and wealth is in line with Moiseev et al. (2020), who argue that individuals in wealthier societies are less likely to engage in corruption schemes than poor societies, where scarce opportunities make corruption a more attractive option.

[109] *Corruption control* and *GDP per capita* have a correlation of 0.77.

Table 5.4
How market integrity varies with country characteristics
This table reports results of regressions in which the dependent variables are the market integrity index ("*MI_Index*"), insider trading index ("*IT_Index*"), and market manipulation index ("*MM_Index*") and the independent variables are country characteristics. The country characteristics variables are defined in Table 5.1 and proxy for country development, corruption, culture, and location. First, our country development variables are *Market cap to GDP*, *Trade value to market cap*, *Trade value*, and *GDP per capita*; second, our equality variable are *Corruption control* and *GINI*; third, our culture variables are *Power distance, Uncertainty avoidance*, and *Masculinity*; and fourth, our location variables are *Asia-Pacific*, *Europe* and *North America*, where *North America* is the benchmark. Numbers not in brackets are coefficient estimates. The regressions are on a global country-year sample including the 25 jurisdictions from Section 5.3 from 1996 to 2020. Numbers in brackets are t-statistics computed using standard errors double clustered by country and year. Significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Explanatory variables	Dependent variable											
	MI_Index				IT_Index				MM_Index			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Market cap to GDP	-1.98***	-2.53***	-1.05**	-0.97	1.74*	2.71**	0.61	0.47	2.22***	2.35***	1.43***	1.43***
	(-3.75)	(-4.19)	(-2.06)	(-1.56)	(1.78)	(2.43)	(0.58)	(0.39)	(8.92)	(9.55)	(4.27)	(4.99)
Trade value to market cap	1.66	2.73	0.42	2.18	0.00	-0.90	2.57	0.27	-2.52	-2.71	-2.29	-4.14***
	(0.43)	(0.72)	(0.12)	(0.73)	(0.14)	(-0.13)	(0.37)	(0.04)	(-1.21)	(-1.30)	(-1.43)	(-2.72)
Trade value	0.03**	0.04***	0.04***	0.02	-0.02	-0.03	-0.02	0.00	-0.04**	-0.04**	-0.05***	-0.03**
	(2.50)	(2.69)	(2.61)	(1.47)	(-0.78)	(-1.02)	(-0.90)	(-0.09)	(-2.57)	(-2.53)	(-3.80)	(-2.32)
GDP per capita	0.04***	0.01	-0.01	-0.01	-0.07***	-0.02	0.00	0.00	-0.01	0.00	0.01	0.01
	(2.62)	(0.40)	(-0.67)	(-0.54)	(-3.23)	(-0.83)	(0.21)	(0.09)	(-0.87)	(-0.03)	(0.97)	(1.03)
Control of corruption		0.81**	0.58**	0.69***		-1.42**	-1.06**	-1.20***		-0.17	-0.15	-0.26
		(2.39)	(2.16)	(2.84)		(-2.37)	(-2.47)	(-3.09)		(-0.73)	(-0.57)	(-1.00)
Power distance			-0.42***	-0.32**			0.62**	0.49**			0.17*	0.11*
			(-2.70)	(-2.41)			(2.49)	(2.15)			(1.81)	(1.69)
Uncertainty avoidance			0.13*	0.10			-0.19	-0.14			-0.12***	-0.09**
			(1.68)	(1.18)			(-1.15)	(-0.82)			(-2.63)	(-2.03)
Masculinity			0.04	0.08			-0.09	-0.13			0.03	0.00
			(0.46)	(0.75)			(-0.51)	(-0.69)			(0.46)	(-0.01)

Table 5.4 (continued)
How market integrity varies with country characteristics
This table reports results of regressions in which the dependent variables are the market integrity index ("*MI_Index*"), insider trading index ("*IT_Index*"), and market manipulation index ("*MM_Index*") and the independent variables are country characteristics. The country characteristics variables are defined in Table 5.1 and proxy for country development, corruption, culture, and location. First, our country development variables are *Market cap to GDP*, *Trade value to market cap*, *Trade value*, and *GDP per capita*; second, our equality variable are *Corruption control* and *GINI*; third, our culture variables are *Power distance, Uncertainty avoidance*, and *Masculinity*; and fourth, our location variables are *Asia-Pacific*, *Europe* and *North America*, where *North America* is the benchmark. Numbers not in brackets are coefficient estimates. The regressions are on a global country-year sample including the 25 jurisdictions from Section 5.3 from 1996 to 2020. Numbers in brackets are t-statistics computed using standard errors double clustered by country and year. Significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Explanatory variables	Dependent variable											
	MI_Index				IT_Index				MM_Index			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Asia-Pacific				-0.74				1.11				0.66**
				(-1.39)				(1.08)				(2.34)
Europe				-0.14				0.29				0.27
				(-0.40)				(0.40)				(0.82)
Other				0.48				-0.45				-0.42
				(0.61)				(-0.31)				(-1.06)
Intercept	52.74***	53.06***	55.04***	54.53***	55.63***	55.07***	52.18***	52.75***	38.88***	38.78***	38.12***	38.32***
	(71.61)	(69.74)	(32.61)	(39.00)	(46.45)	(45.24)	(18.18)	(21.43)	(98.09)	(94.88)	(41.81)	(54.31)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	17%	18%	22%	23%	15%	16%	19%	19%	28%	29%	43%	50%
Observations	508	508	508	508	508	508	508	508	587	587	587	587

Turning to the effects of culture on market integrity (using Hofstede's cultural dimensions), societies with an acceptance of inequality among its members (measured by *Power distance*) tend to have higher levels of insider trading, more market manipulation, and lower overall integrity, consistent with *H6*. An increase of 10 (out of 100) dimension points in *Power distance*, increases the insider trading and market manipulation indices by 0.62 and 0.17 respectively, and decreases the market integrity index by 0.42 points. To put the magnitude of the coefficient into perspective we compare the standard deviation of *Power distance* to the standard deviation of the indices: A one standard deviation in *Power distance* accounts for 19%, 34%, and 25% of the standard deviation in the insider trading index, market manipulation index, and market integrity index respectively.

Societies with higher levels of *Uncertainty avoidance* such that they minimize uncertainty with strict behavioral codes, laws, and rules also have higher market integrity consistent with *H7* but only in terms of less market manipulation—a 10 (out of 100) dimension point increase in *Uncertainty avoidance* decreases the market manipulation index by 0.12 index points. A one standard deviation in *Uncertainty avoidance* accounts for 26% of the standard deviation in the market manipulation index and 8% of the standard deviation in the market integrity index.

The results do not support *H8* because the level of masculinity in the country has little to no relation with the country's stock market integrity. Overall, societies that accept that power is distributed unequally (measured by *Power distance*) in countries guided by strict behavioral codes and laws (measured by *Uncertainty avoidance*) are more transparent, have more rule-abiding citizens, and, thus, lower levels of financial market misconduct. Interestingly, when controlling for location, more liquid markets (measured by *Trade value to market cap*) have lower levels of market manipulation, which is consistent with the rankings in Table 5.2 and *H2* that more liquid stocks are harder to manipulate. For a \$1000 increase in a country's yearly traded value per \$1 dollar of market capitalization, the market manipulation index decreases by 4.14 index points. A one standard deviation in the *Trade value to market cap* accounts for one fourth of a standard deviation in the market manipulation index.

Turning to the region variables,[110] *Asia-Pacific* tends to have higher levels (0.66 index points) of market manipulation than North America, controlling for all other country characteristics. This result is in line with *H9*.[111]

5.5.3 Market structure

Three major changes in stock market structure over the past three decades are that financial markets have "fragmented" with the emergence of competition among stock exchanges and trading

[110] We include North America in the regression but do not include the North America dummy variable so that this region becomes the baseline against which other regions are compared. Table 5.1 lists the countries included in each location dummy variable.

[111] Thailand and Malaysia have a GDP per capita around \$8 and \$11 thousand, respectively, while the US is close to \$71 thousand.

venues, dark trading has lowered trade transparency by introducing “dark pools” where orders remain undisclosed until they are matched, and high-frequency algorithmic trading has increased trade speed and impacted market trading characteristics. We use our indices to measure the impact of these changes on insider trading, market manipulation, and overall market integrity. We use the countries affected by the market microstructure event as the “treatment group” in regressions and all other countries as the “control group”. Given the changes occur in a staggered manner in different countries through time, we include country and year fixed effects in all regressions, effectively making them into difference-in-differences models.

In the first set of regressions, we test the effect of market fragmentation on market integrity. The countries included in the treatment groups are the US, Canada, the EU member countries, and Australia. In 2005, RegNMS fragmented the US market with the “order protection rule,” ensuring that that the markets offer investors the best price. In the EU, MiFID I fragmented the market by abolishing the “concentration rule,” which ensured orders could only submitted to national exchanges. The Australian and Canadian markets fragmented when other exchanges (most notably Chi-X) entered their markets in 2008 and 2011, respectively.[112] We group the US and Canada together into one “North America” dummy (*Frag NA*), while the fragmentation dummies for EU and Australia are *Frag EU* and *Frag AU,* respectively.

The results in Table 5.6 partially support *H10* in that fragmentation is associated with a lower level of market manipulation. Fragmentation is associated with a decreased in the level of manipulation in North America by almost 1 index point (coefficient on *Frag_NA* equals -0.86 and -0.80 with the controls). The effect of fragmentation in Australia was about 50% larger, with a drop of almost 1.4 index points (the *Frag_AU* coefficient is -1.20 and -1.39 with the controls). The coefficients may appear small at first glance, but when compared to the variations in the market manipulation index, they are economically meaningful. In North America (*Frag_NA*) and Australia (*Frag_AU*), the impact of fragmentation is around 2 and 3.3 times larger than the standard deviation changes in the market manipulation index respectively.

O’Hara and Ye (2011) argue that fragmentation leads to increased competition, lower transactions costs, and increased market quality liquidity. The increased level of liquidity from fragmentation is likely to contribute to the drop in the amount of manipulation in the North American and Australian markets.[113] In contrast, fragmentation in Europe, following EU’s MiFID I regulation, has no significant effect on integrity. Unlike the US, Canada, and Australia, the EU is a collection of many governments, each with their own financial regulator. Although the EU requires member states to follow EU legislation, the manner and extent to which they implement it is not always consistent.

[112] Before 2008 and 2011, the national stock exchanges of Canada (TSX) and Australia (ASX) were monopolies.

[113] Liquid stocks are harder to manipulate because the manipulator competes more trades when setting the closing price (Comerton-Forde and Putniņš, 2011).

Table 5.5
The impact of stock market fragmentation on market integrity
This table reports results of regressions in which the dependent variables are the market integrity index ("*MI_Index*"), insider trading index ("*IT_Index*"), and market manipulation index ("*MM_Index*") and the independent variables are fragmentation dummy variables. *Frag_NA*, *Frag_EU*, and *Frag_AU* are dummy variables that take the value of one after stock trading in North America, the EU, and Australia fragmented across competing trading venues, respectively (as defined in Table 5.1). The regressions are on a global country-year sample including the 25 jurisdictions from Section 5.3 from 1996 to 2020. Numbers not in brackets are coefficient estimates. Numbers in brackets are t-statistics computed using standard errors double clustered by country and year. Significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Explanatory variables	Dependent variable					
	MI_Index		IT_Index		MM_Index	
	(1)	(2)	(3)	(4)	(5)	(6)
Frag_NA	0.02	0.24	0.77	0.20	-0.86***	-0.80***
	(0.04)	(0.34)	(0.76)	(0.15)	(-5.41)	(-3.86)
Frag_EU	-0.05	-0.15	-0.27	-0.03	0.37	0.35
	(-0.06)	(-0.16)	(-0.16)	(-0.02)	(1.09)	(0.96)
Frag_AU	0.68*	0.53	-0.19	0.32	-1.20***	-1.39***
	(1.70)	(0.80)	(-0.24)	(0.25)	(-7.76)	(-5.41)
Intercept	55.57***	57.29***	51.35***	47.91***	37.51***	37.34***
	(86.75)	(27.36)	(39.70)	(11.07)	(235.11)	(70.47)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes		Yes		Yes
R-squared	27%	28%	23%	23%	74%	75%
Observations	523	508	523	508	612	587

In the second set of market structure regressions, we test the effect of dark pools on market integrity (*H11*). Dark pools (or dark trading venues) have long been a cause for concern among regulators due to their lack of pre-trade transparency, their reliance on using prices from other transparent markets as reference points, and the manipulation incentives that this reference pricing creates. On October 15, 2012, Canada introduced rules that restrict dark trading and Australia followed suit on May 26, 2013. The rules required dark trades to provide a meaningful price improvement,[114] and after their inception, dark trading dropped by approximately one-third in both countries (Foley and Putniņš, 2016). We use this exogenous shock to dark trading to test the effect of dark trading on integrity.

The results are in Table 5.6. *Dark trading restrictions* is a dummy variable for after the Canadian and Australian dark trading restrictions. The implementation of the dark trading restrictions is associated with a significant decrease in the level of manipulation; the Canadian and Australian stock markets saw a drop of more than 1 index point in the manipulation index (coefficient of *Dark trading*

[114] The rules require dark trades to provide at least one tick of price improvement (half a tick if the spread is constrained at one tick).

restrictions is -1.08 and -1.16 with the controls). Compared to the variation in the market manipulation index, this impact is economically meaningful: the effect of the dark trading restrictions on market manipulation is almost three times the standard deviation of changes in the market manipulation index. Dark pools are vulnerable to stock price manipulation because they derive their prices from the publicly quoted prices in the lit market; therefore, a manipulator can manipulate trade prices in the dark pool by trading on a regular exchange.

Table 5.6
The impact of dark trading on market integrity
This table reports results of regressions in which the dependent variables are the market integrity index ("*MI_Index*"), insider trading index ("*IT_Index*"), and market manipulation index ("*MM_Index*") and the independent variable is dark trading restrictions. *Dark trading restrictions* is a dummy variable that takes the value of one in Canada and Australia after those countries implement restrictions on dark trading (as defined in Table 5.1). The regressions are on a global country-year sample including the 25 jurisdictions from Section 5.3 from 1996 to 2020. Numbers not in brackets are coefficient estimates. Numbers in brackets are t-statistics computed using standard errors double clustered by country and year. Significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Explanatory variables	Dependent variable					
	MI_Index		IT_Index		MM_Index	
	(1)	(2)	(3)	(4)	(5)	(6)
Dark trading restrictions	0.46	0.31	0.14	0.51	-1.08***	-1.16***
	(1.12)	(0.64)	(0.16)	(0.49)	(-5.67)	(-5.18)
Intercept	55.57***	57.03***	51.81***	47.76***	37.03***	37.74***
	(204.7)	(30.42)	(94.67)	(12.21)	(588.89)	(57.52)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes		Yes		Yes
R-squared	27%	28%	23%	23%	72%	74%
Observations	523	508	523	508	612	587

Our last set of market structure tests analyzes the effects of high-frequency trading (HFT). One criticism against HFT firms is that their ability to submit and cancel orders faster than many other traders and their ability to automate trading strategies allows them to deploy various manipulative and predatory trading strategies designed to exploit slower or less sophisticated market participants. However, a counterargument is that HFTs tend to make markets more liquid, which can make it more difficult for others to manipulate market prices.[115] To identify the effects of HFT activity, we use the launch of exchange colocation services as an exogenous shock to analyze their effect on market

[115] On October 16, 2014, in the first manipulation case against a HFT firm (Athena Capital), the SEC received a $1 million settlement amount. The SEC charged Athena Capital with manipulating closing prices by flooding the market with buy and sell orders (see sec.gov for more information on the case).

integrity.[116] Colocation lets HFTs place their trading servers at the same location as the exchange servers, thereby giving them low trade latency (for a fee). The introduction of colocation in a country naturally leads to a spike in HFT activity (Boehmer, Fong, and Wu, 2020).

The results are in Table 5.7. High frequency trading reduces the manipulation index by almost 0.5 index points (*HFT* coefficient equals -0.39, or -0.42 with the controls). Compared to the variation in the market manipulation index, this impact is economically meaningful: HFT decreases the manipulation index by around one standard deviation of changes in the market manipulation index. The result supports *H12* and is consistent with Aitken, Cumming, and Zhan (2015), who find that the benefits of HFTs as liquidity providers outweigh any role that they may play in closing price manipulation.

Table 5.7
The impact of high-frequency trading on market integrity
This table reports results of regressions in which the dependent variables are the market integrity index ("*MI_Index*"), insider trading index ("*IT_Index*"), and market manipulation index ("*MM_Index*") and the independent variable is high-frequency trading. *HFT* is a dummy variable that takes the value of one when a country's market has colocation and zero otherwise (as defined in Table 5.1). The regressions are on a global country-year sample including the 25 jurisdictions from Section 5.3 from 1996 to 2020. Numbers not in brackets are coefficient estimates. Numbers in brackets are t-statistics computed using standard errors double clustered by country and year. Significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

Explanatory variables	Dependent variable					
	MI_Index		IT_Index		MM_Index	
	(1)	(2)	(3)	(4)	(5)	(6)
HFT	0.13	0.16	0.09	0.04	-0.39**	-0.42**
	(0.22)	(0.22)	(0.08)	(0.03)	(-1.98)	(-2.34)
Intercept	55.47***	56.99***	51.72***	47.76***	37.38***	37.93***
	(64.92)	(30.16)	(30.96)	(12.13)	(194.24)	(62.14)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls		Yes		Yes		Yes
R-squared	27%	28%	23%	23%	72%	73%
Observations	523	508	523	508	612	587

5.5.4 Regulation and enforcement

There are various regulatory and policy strategies that can be used to try and improve market integrity. The four that we test are whistleblowing, enforcement of insider trading laws, increased

[116] We use colocation dates from Boehmer, Fong, and Wu (2020) found in the last column in Table G1 of Appendix G. We add colocation dates for the remaining jurisdictions, including (colocation year in parenthesis) Israel (2018), South Africa (2014), Poland (2013), South Korea (2016), Taiwan (2017), Hong Kong (2012), and China (2016).

penalties, and intercontinental regulatory cooperation. We regress the insider trading and market manipulation indices on variables capturing the four regulatory strategies. Table 5.8 reports the results.

The first variable in the regressions (*Whistleblowing*) is a dummy variable that takes the value of one in a country after it implements substantial whistleblowing schemes and zero otherwise.[117] Whistleblower schemes are a widespread practice among regulators as a way to detect (and deter) misconduct. The schemes often incentivize the whistleblower with rewards in the form of either cash or immunity.[118] The results suggest that whistleblowing tends to decreases a country's level of insider trading by almost 1.3 index points (*Whistleblowing* coefficient in regression (2) is -1.25) and increases market integrity by almost 1 index point (coefficient on *Whistleblowing* in (1) is 0.78), consistent with *H13*. The coefficients may appear small, but when compared to the variation of the indices they are economically meaningful: the *Whistleblowing* coefficients accounts for 14% of the standard deviation changes in the insider trading index and 17% of the standard deviation changes in the market integrity index.

Bhattacharya and Daouk (2002) argue that it is not the existence of insider trading laws but, rather, their enforcement that matters. They find that the first insider trading prosecution significantly lowers the cost of capital in the country. With our indices, we test if this also holds for market integrity: does a country's first insider trading prosecution decrease its level of insider trading? The dummy variable *IT enforcement* equals one after a country's first insider trading prosecution and zero otherwise. Of the four variables relating to regulatory strategies, IT enforcement has (by far) the biggest impact on insider trading (in support of *H14*). On average, the first prosecution in a country decreases the level of insider trading by more than 4.5 index points (*IT enforcement* coefficient in regression (2) is -4.71) and increases the level of market integrity in the country by more than 2 index points (coefficient on *IT enforcement* in regression (1) is 2.21). When compared to the variation in the index, the *IT enforcement* coefficients accounts for 50% of the standard deviation changes in both the insider trading index and market integrity index.

We also test the effect of increasing penalties on market integrity. In 2014, the EU increased the jail time and penalties for insider trading and market manipulation in member states.[119] Consistent with *H15*, the directive had a significant effect and decreased the level of insider trading in member countries by almost 2 index points (coefficient on *Penalty increase* in regression (1) equals -1.78), leading to a rise in integrity of about 0.5 index points (*Penalty increase* coefficient in regression (2) equals 0.63). The *Penalty increase* coefficients accounts for 20% and 14% of the standard deviation changes in the insider trading and market integrity indices respectively.

[117] For whistleblower dates, see Table G2 in Appendix G.

[118] The SEC has awarded $676 million to 108 individuals since the first award in 2012. The Australian ASIC announced on February 24, 2021 their immunity policy protecting whistleblowers from any legal action related to the crime (see SEC press release 2020-266).

[119] The legislation stipulates insider trading and market manipulation culprits in all member states will face imprisonment for four years and fines of at least €5 million. See MEMO/14/77 on ec.europa.eu.

At the G20 summit in Seoul in 2010, the members agreed to "promote market integrity and efficiency to mitigate the risks posed to the financial system by the latest technological developments"[120] and called on the International Organization of Securities (IOSCO) to make recommendations to the Financial Stability Board (FSB). In October of 2011, IOSCO published the report, and the G20 and the FSB committed to implement the recommendations.[121] The event is particularly interesting because it allows us to observe the effects of intercontinental cooperation. Such cooperation may present a considerable challenge considering the diversity of the jurisdictions and markets involved. The results support *H16*, suggesting that the G20 members saw the level of market manipulation decrease by 0.65 index points (coefficient of *Intercontinental cooperation* in regression (3) equals -0.65) following this directive. The *Intercontinental cooperation* coefficient accounts for more than 1.5 standard deviation changes in the market manipulation index.

[120] See IOSCO July 2011 consultation report (CR02/11) "Regulatory Issues Raised by the Impact of Technological Changes on Market Integrity and Efficiency".

[121] For more information on the recommendations, see IOSCO final report "G20/FSB Recommendations related to Securities Markets."

Table 5.8
Effect of regulatory strategies on market integrity
This table reports results of regressions in which the dependent variables are the market integrity index ("*MI_Index*"), insider trading index ("*IT_Index*"), and market manipulation index ("*MM_Index*") and the independent variables are regulatory strategy dummy variables. The regulatory strategy dummy variables are defined in table 5.1 and are as follows: *Whistleblowing* is a dummy variable that takes the value of one when a country has legislative whistleblower incentives in place; *IT enforcement* is a dummy variable that takes the value of one after the country first enforces insider trading laws; *Penalty increase* is a dummy variable for EU member states that takes the values of one after they increased their fines and jail sentences for insider trading and market manipulation in 2014; and *Intercontinental cooperation* is a dummy variable for G20 member states that takes the value of one when they started implementing IOSCO market integrity directives. The regressions are on a global country-year sample including the 25 jurisdictions from Section 5.3 from 1996 to 2020. Numbers not in brackets are coefficient estimates. Numbers in brackets are t-statistics computed using standard errors double clustered by country and year. Significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively.

	Dependent variable		
Explanatory variables	MI_Index	IT_Index	MM_Index
	(1)	(2)	(3)
Whistleblowing	0.78***	-1.25***	-0.27
	(3.18)	(-2.6)	(-1.51)
IT enforcement	2.21**	-4.71***	0.01
	(2.28)	(-2.58)	(0.04)
Penalty increase	0.63**	-1.78***	0.37
	(2.52)	(-3.08)	(1.50)
Intercontinental cooperation	-0.13	1.06	-0.65**
	(-0.17)	(0.76)	(-2.24)
Intercept	52.44***	57.61***	37.67***
	(36.65)	(21.44)	(76.77)
Country FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
R-squared	28%	24%	73%
Observations	523	523	612

5.6 Conclusion

The market integrity index shows that the five countries with the highest market integrity are the US, Japan, Netherlands, Canada, and Australia, and the five lowest ranking are Hong Kong, Singapore, Malaysia, Thailand, and China. We find that large, liquid stocks tend to attract insider trading with one explanation being that there are more insiders in such stocks and large stocks offer more liquidity for the insider to hide their trades. Manipulators, on the other hand, tend to prefer small, illiquid stocks, which are easier to manipulate and detection may be less likely.

We also find that whistleblower programs, high penalties, intercontinental cooperation, and, most importantly, enforcement are excellent tools in a regulator's toolbox to deter misconduct. We also find that market integrity is affected by equity market structure—fragmentation of trading across competing trading venues, dark pools, and high-frequency trading. Market fragmentation and high-frequency trading activity tend to increase the overall market liquidity, making the market's underlying stocks less susceptible to manipulation. Dark pools, on the other hand, are a conduit for manipulation strategies, and the results suggest that countries that have imposed restrictions on dark pools have had a small increase in market integrity.

The findings regarding which types of stocks and market characteristics are the most vulnerable to insider trading and market manipulation can help make more efficient use of scarce regulatory resources.

Chapter 6

Conclusions

This chapter summarizes the conclusions from Chapters 2 through 5, including the following: (i) the prevalence of illegal activity in bitcoin; (ii) the characteristics of users involved in illegal activity and the topology of their network; (iii) effective measures of market integrity; and (iv) the trend of market integrity around the world, stocks that are susceptible to low market integrity levels, and effective ways to improve market integrity.

6.1 How pervasive is illegal activity in bitcoin?

Approximately one-fourth (around 28 million) of bitcoin users are illegal, and their activity accounts for almost half of all bitcoin transactions (almost 280 million), one-quarter of transacted volume (just below $430 billion), and half of all bitcoin holdings as of 2017 (about $1,460 million). These results come from two separate models that rely on independent assumptions.

6.2 How has illegal activity in bitcoin changed over time?

In the beginning of 2009, around the creation of bitcoin, there are only few users (below one hundred) in the bitcoin network. However, of those users, a significant proportion (around 60%) are illegal. Bitcoin provides participants with an anonymous means of transaction, and with it, a lucrative opportunity for criminals to bring illegal commerce online. Then in 2011, one of the first darknet markets to use bitcoin, the "Silk Road", becomes operational and the absolute number of illegal users starts to rise. From 2012 to early 2015, the online illegal community grows tremendously and darknet market activity increases from around 0.5 million to more than 10 million. This period marks the "growth phase" of the illegal community, and while law enforcement agencies close many prominent darknet markets, new ones take their place. Darknet market activity in bitcoin continues to rise towards 2017. However, during this period bitcoin also gains popularity among investors and speculators and the proportion of illegal bitcoin activity decreases as they enter the bitcoin network.

6.3 What characterizes illegal bitcoin users and their network?

Bitcoin users associated with illegal activity tend to use bitcoin as a transactional tool to buy and sell illegal goods and services online, so they transact more frequently, especially when more darknet markets are operational. The illegal community also relies more on trust than conventional online commerce where legislation safeguards consumer rights. Participants in the illegal community therefore transact repeatedly with the same counterparty after establishing rapport from successful past transactions. The cost of capital and the constant threat of law enforcement intervention encourages users in the illegal community to hold few bitcoin to limit their losses in potential darknet market seizures. Seizures also increase activity among illegal community participants as they take precautionary measure to hide their bitcoin or reallocate their bitcoin holdings to competing darknet markets.

6.4 How do countries compare in market integrity?

When measuring market integrity around the world, the five highest ranking countries (from high to low) are the US, Japan, Netherlands, Canada, and Australia, and the five lowest ranking (from low to high) are China, Thailand, Malaysia, Singapore, and Hong Kong. The "developed" countries tend to rank higher in market integrity than "emerging" countries because of their sophisticated regulatory bodies and low corruption activity. When grouping the countries by geographical location, North America ranks first, Europe second, and Asia-Pacific last. Regulators often quote insider trading and market manipulation as the main sources poor market integrity. When measuring both types of misconduct we find that cumulative abnormal returns, large trading volumes, and positive order imbalances indicate insider trading while manipulation exhibits intraday order imbalances, stock price volatility, and price reversals.

6.5 How has market integrity changed over time?

In the US, there are three noticeable changes in market integrity. US market integrity improves substantially from 1998 to 2007. This improvement is likely due to the US regulator, the SEC, becoming more effective as the *International Securities Enforcement Cooperation Act of 1990* increased the SEC's ability to address issues in international securities. Market integrity improves again from 2010 to 2013, which may be because of the increased detection rates caused by the Dodd–Frank whistleblower program. The program promised whistleblowers awards between 10% and 30% of the sanctions collected and may have significantly improved the incentive for individuals to provide evidence against misconduct. Finally, integrity drops again after the 2014 U.S. v. Newman ruling, increases insider trading. The ruling made prosecution less likely for individuals who are several links away from the original source of private information, thus increasing the incentive for insider trading activity.

When examining the trends in the world, the results indicate that market manipulation levels drop substantially among developed countries after the G20 countries agreed to follow IOSCO directives to combat market misconduct. The first decrease occurs from 2011 to 2013 during the first years of the implementation phase and at an increasing rate in 2014 after 22 jurisdictions reported completed implementation of the directives. From 1998 to 2006, there also appears to be a significant improvement in market integrity in Asia, which may be because of a large number of Asian countries advancing in terms of development. In Asia, Hong Kong, Singapore, South Korea, and Taiwan transitioned from "developing" countries to "developed" in 1997, while China, India, Malaysia, and Thailand became newly industrialized countries (NICs) in the late part of the 2000s.

6.6 What determines market integrity?

Large, liquid stocks attract insider trading because they have more insiders and offer higher order flow that effectively hides the insider's trades. Manipulators, on the other hand, prefer small, illiquid stocks because they are easier to manipulate and detection is less likely given the lower analyst coverage of small companies.

The level of (i) market development, (ii) corruption, (iii) culture, and (iv) the geographical location of countries also significantly affect their market integrity levels. In terms of market development, there is a high correlation between market overvaluation (the market's propensity to crashes) and low market integrity levels. High-liquidity markets, however, have low levels of market manipulation (and high market integrity) because manipulators compete with more trades when setting the closing price. Similarly, larger and richer markets generally have more sophisticated (although costly) regulation, so their market integrity levels are lower. Conversely, corruption acts as a major contributor to countries' insider trading activity; markets with high corruption levels are less transparent, in turn encouraging financial misconduct through low perceived detection rates among offenders.

Culture, as measured by Hofstede's culture dimensions, also significantly affects country market integrity. Countries where the poorer parts of the population share a general acceptance of their unfortunate societal standing have lower market integrity. Superiors and subordinates in these markets commonly trade favors for loyalty; this transactional relationship in the country's population leads to lower market integrity levels because subordinates are less prone to report misconduct committed by their superiors. Uncertainty-avoiding societies, where members follow strict behavioral codes, laws, and rules, have higher levels of market integrity. In these societies, the population is less inclined to commit misconduct, especially given their propensity to avoid "uncertainty."

6.7 How does equity market structure and regulation affect market integrity?

Market design changes such as fragmenting concentrated markets or allowing high frequency trading, may discourage market manipulation by increasing market liquidity levels as the closing prices of liquid stocks are harder to manipulate. Market fragmentation increases liquidity through competition, while high-frequency traders act as liquidity providers on days of extreme price movements.

Implementing dark trading restrictions may also improve market integrity—dark pools provide manipulators with alternative manipulation strategies, so decreasing the presence of dark pools decreases the strategies available to manipulators.

Regulatory strategies may also deter financial misconduct in markets. Whistleblower schemes can increase detection rates by providing incentives such as monetary awards or immunity to individuals who provide evidence of financial misconduct. Countries can also deter misconduct by increasing the perceived likelihood of detection among offenders by enforcing (not just passing) financial misconduct legislation. High penalties for misconduct, such as long prison sentences or large fines, can also act as an effective misconduct deterrent because criminal activity becomes less attractive. Finally, countries can increase their market integrity by cooperating with other jurisdictions when combating misconduct in their respective markets.

6.8 Implications for misconduct legislation and enforcement

This thesis provides government agencies, such as legislators, regulators, and law enforcement, with the techniques to measure misconduct in their markets. Using the DCE and SLM models developed in this thesis, regulators can measure the prevalence of illegal activity in bitcoin and other cryptocurrencies. The model's findings can help guide cryptocurrency legislation and inform regulators on where to devote resources when combating misconduct in cryptocurrencies. Law enforcement agencies can also use the models to measure the effect of their strategies when fighting darknet activity and to benchmark their overall progress. The DCE model results provided in this thesis also show the key characteristics of the users involved in illegal activity and their network, which law enforcement may use to identify key individuals (or hubs) in the illegal network.

The high proportion of illegal bitcoin activity (46.17% of all transactions) also has considerable ethical implications for bitcoin (and other cryptocurrencies) as an investment. The high dependence of the value of bitcoin on how actively it is used (coupled with a high proportion of illegal usage) means that bitcoin investments appreciate when the illegal network is active. Therefore, bitcoin investors profit when illegal bitcoin activity increases and make losses when law enforcement successfully shut down darknet marketplaces.

This thesis also provides government agencies with a measure of market integrity in traditional financial markets. Similar to the techniques developed for the bitcoin network, law enforcement agencies can use the index to keep a finger on the "pulse" of illegal activity and measure the effect of their regulatory strategies. The results also give insights into the types of stocks that are susceptible to misconduct and the determinants of market integrity, which can help regulators devote their resources efficiently. Finally, the findings of this thesis show what market designs improve market integrity and what regulatory strategies can effectively deter and combat misconduct.

6.9 Avenues of future research

The techniques developed in this thesis take the first step in measuring illegal activity in cryptocurrencies and the thesis develops new measures of market integrity for traditional equity markets.

Future research may use the SLM and DCE models developed in Chapter 2 to measure illegal activity in newly developed cryptocurrencies, the characteristics of users involved in illegal activity, and the determinants of their detection. This research is especially important with the rapid increase of shadow coins, which feature increasingly sophisticated techniques to disguise their user's illegal activity.

Researchers can also shed light on the welfare consequences of illegal bitcoin activity by measuring illegal activity in conventional black markets and other cryptocurrencies. Bitcoin may facilitate illegal trade on darknet markets, but the effect of the emergence of cryptocurrencies on conventional (cash-based) black markets is not yet understood; the rising popularity of bitcoin (and darknet markets) may, for example, mean decreased activity in illegal cash-based markets as cryptocurrencies become more popular as a means of illegal transactions. Therefore, the key question is if total (cash and cryptocurrency based) black market activity has increased because of cryptocurrencies or if the rise in illegal darknet activity simply reflects a shift—black market participants adopting bitcoin rather than conventional cash. Future research might attempt to quantify the total black market activity—in both darknet markets and conventional black markets. If the rise in darknet markets merely reflects a migration of activities that would have otherwise occurred in conventional markets, then illegal online activity may be good from a welfare perspective. Online darknet markets may, for example, be safer and lead to reduced violence that would have otherwise occurred "on the street." From a public health perspective, darknet markets also rely on user feedback and the vendor's online reputation, which may increase the overall quality and safety of the goods sold online.

Future research may also expand the results of this thesis on illegal cryptocurrency activity to include more recent years. It may be especially interesting to observe the effects of the boom in decentralized finance (DEFI) since 2020 and with it the increased regulatory attention on illegal cryptocurrency activity. The growth of DEFI may have prompted regulators to devote significant resources to manage illegal darknet markets and the effects of their regulatory efforts would be interesting to measure.

The market integrity measures developed in Chapter 4 also provide researchers with a much-needed tool to measure the effects of market legislation and regulatory enforcement on market integrity. Chapter 5 takes the first step in demonstrating their application in measuring the effect of market design and regulation on market integrity; researchers can use the techniques to measure the effect of such events in their local jurisdictions. The market integrity measures also focus on two types of market misconduct; future research could add more types of misconduct to the indices, thereby further increasing their value as forensic tools.

The market integrity measures developed in this thesis may also provide researchers with a tool to answer questions about the effects of a range of topics on market integrity. Researchers can for example use the indices to measure the effects of market participation, corporate decision-making, and asset pricing on market integrity.

Appendix A

Darknet sites and bitcoin seizures

Table A1: Darknet sites accepting bitcoin, current and past

This table reports the 30 known darknet marketplaces with the longest operational history. For sites that remain operational (as at May 2017), the *End date* column states "Operational" and thus there is no *Closure reason*. *Days operational* is the number of days the site was operational before closure. Data are sourced from www.gwern.net.

Market	Launch date	End date	Closure reason	Days operational
Dream	November 15, 2013	Operational		>1,207
Outlaw	December 29, 2013	May 16, 2017	Hacked	1234
Silk Road 1	January 31, 2011	October 2, 2013	Raided	975
Black Market Reloaded	June 30, 2011	December 2, 2013	Hacked	886
AlphaBay	December 22, 2014	July 4, 2017	Raided	925
Tochka	January 30, 2015	Operational		>766
Crypto Market / Diabolus	February 14, 2015	Operational		>751
Real Deal	April 9, 2015	Operational		>697
Darknet Heroes	May 27, 2015	Operational		>649
Agora	December 3, 2013	September 6, 2015	Voluntary	642
Nucleus	October 24, 2014	April 13, 2016	Scam	537
Middle Earth	June 22, 2014	November 4, 2015	Scam	500
BlackBank	February 5, 2014	May 18, 2015	Scam	467
Evolution	January 14, 2014	March 14, 2015	Scam	424
Silk Road Reloaded	January 13, 2015	February 27, 2016	Unknown	410
Anarchia	May 7, 2015	May 9, 2016	Unknown	368
Silk Road 2	November 6, 2013	November 5, 2014	Raided	364
The Marketplace	November 28, 2013	November 9, 2014	Voluntary	346
Blue Sky Market	December 3, 2013	November 5, 2014	Raided	337
Abraxas	December 13, 2014	November 5, 2015	Scam	327
Pandora	October 21, 2013	August 19, 2014	Scam	302
BuyItNow	April 30, 2013	February 17, 2014	Voluntary	293
TorBazaar	January 26, 2014	November 5, 2014	Raided	283
Sheep	February 28, 2013	November 29, 2013	Scam	274
Cloud-Nine	February 11, 2014	November 5, 2014	Raided	267
Pirate Market	November 29, 2013	August 15, 2014	Scam	259
East India Company	April 28, 2015	January 1, 2016	Scam	248
Mr Nice Guy 2	February 21, 2015	October 14, 2015	Scam	235
Andromeda	April 5, 2014	November 18, 2014	Scam	227
Topix 2	March 25, 2014	November 5, 2014	Voluntary	225

Table A2: Bitcoin seizures

This table reports major bitcoin seizures, the seizing authority, the owner of the seized bitcoin, the date of the seizure, and the amount (in bitcoin) seized.

Seizing authority	Seized entity	Owner of seized bitcoins	Date of seizure	Bitcoin seized
Australian Government	Individual	Richard Pollard	December 1, 2012	24,518
US government	Individual	Matthew Luke Gillum	July 23, 2013	1,294
ICE and HSI	Individual	Cornelius Jan Slomp	August 27, 2013	385,000
FBI	Individual	Ross William Ulbricht	October 1, 2013	144,000
FBI	Site	Silk Road escrow accounts (many users)	October 2, 2013	29,655

Appendix B

Derivations for the DCE model

We define $I(.)$ and $D(.)$ to be monotonic link functions that map $x_{1i}\beta_1$ and $x_{2i}\beta_2$ to latent probabilities of a bitcoin user being involved predominantly in illegal activity, and detection of an illegal user, respectively.[122] That is,

$$I(x_{1i}\beta_1) = \Pr(L_{1i} = 1) \tag{B.1}$$

$$D(x_{2i}\beta_2) = \Pr(L_{2i} = 1 \mid L_{1i} = 1) \tag{B.2}$$

Table B1 reports the probabilities of various joint outcomes (represented by cells in the table). The joint outcomes are mutually exclusive and exhaustive, so the probabilities in Table B1 sum to one.

Table B1: Two-stage DCE model probability matrix

	Illegal user	Legal user
Detected	$I(x_{1i}\beta_1)D(x_{2i}\beta_2)$	0
Not detected	$I(x_{1i}\beta_1)[1-D(x_{2i}\beta_2)]$	$1-I(x_{1i}\beta_1)$

The log likelihood of the users that end up in the detected (seized) illegal users category (A) is the log of the sum (over users in A) of the probabilities of that joint outcome:

$$log\ L_A = \sum_{i\epsilon A} \log[I(x_{1i}\beta_1)D(x_{2i}\beta_2)] \tag{B.3}$$

Similarly, the log likelihood of the users that end up in the other category (A^C) is the log of the sum (over users in A^C) of the probabilities of that joint outcome (the probability that the user is a legal one plus the probability that an illegal user is not detected):

$$log\ L_{A^C} = \sum_{i\epsilon A^C} \log[I(x_{1i}\beta_1)[1-D(x_{2i}\beta_2)] + 1 - I(x_{1i}\beta_1)] \tag{B.4}$$

$$log\ L_{A^C} = \sum_{i\epsilon A^C} \log[1 - I(x_{1i}\beta_1)D(x_{2i}\beta_2)] \tag{B.5}$$

Sets A and A^C constitute the universe of all bitcoin users. Therefore the full-sample log likelihood is:

$$log\ L = \sum_{i\epsilon A} \log[I(x_{1i}\beta_1)D(x_{2i}\beta_2)] + \sum_{i\epsilon A^C} \log[1 - I(x_{1i}\beta_1)D(x_{2i}\beta_2)] \tag{B.6}$$

Maximum likelihood estimation involves selecting parameter vectors β_1 and β_2 such that the function $log\ L$ is maximized.

[122] In our implementation, the link functions are cumulative logistic distribution functions, that is, $I(x_{1i}\beta_1) = \frac{1}{1+e^{-x_{1i}\beta_1}}$, $D(x_{2i}\beta_2) = \frac{1}{1+e^{-x_{2i}\beta_2}}$.

Appendix C

Estimated illegal activity in bitcoin from DCE and SLM models

Panel A: SLM estimated number of illegal and legal bitcoin users

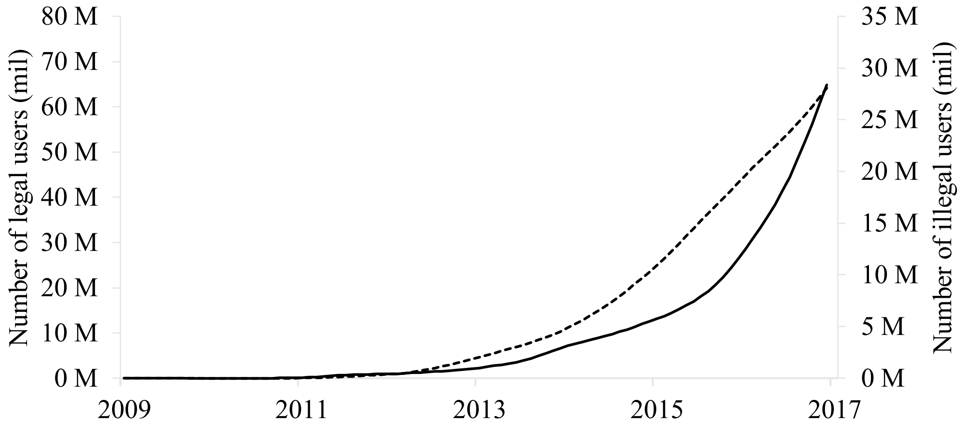

Panel B: SLM estimated percentage of illegal bitcoin users

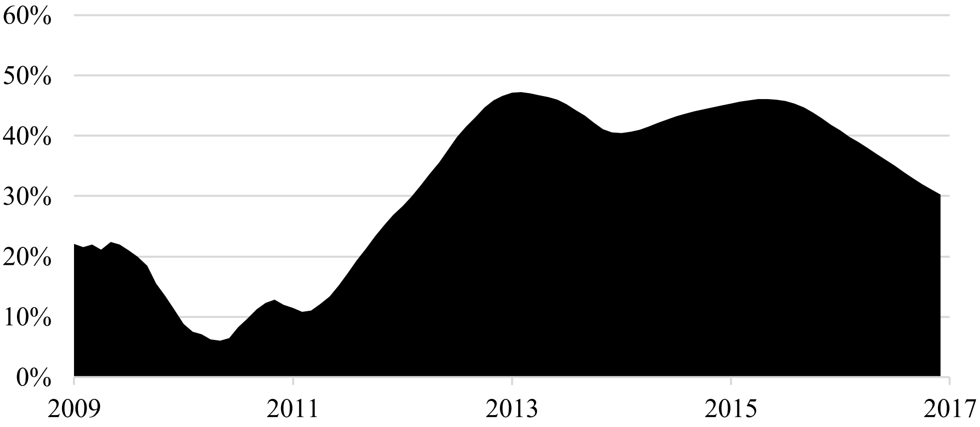

Figure C1
SLM estimated number and percentage of bitcoin users involved in illegal activity
This figure illustrates the time-series of the estimated number of illegal and legal bitcoin users (Panel A) and the percentage of illegal users (Panel B). In Panel A, the number of legal users is plotted with the solid line using the left-hand-side axis and the number of illegal users is plotted with the dashed line using the right-hand-side axis. In Panel B, the black area is the estimated percentage of illegal users and the white area is the estimated percentage of legal users. The estimates come from a network cluster analysis algorithm (SLM). Values are smoothed with a moving average.

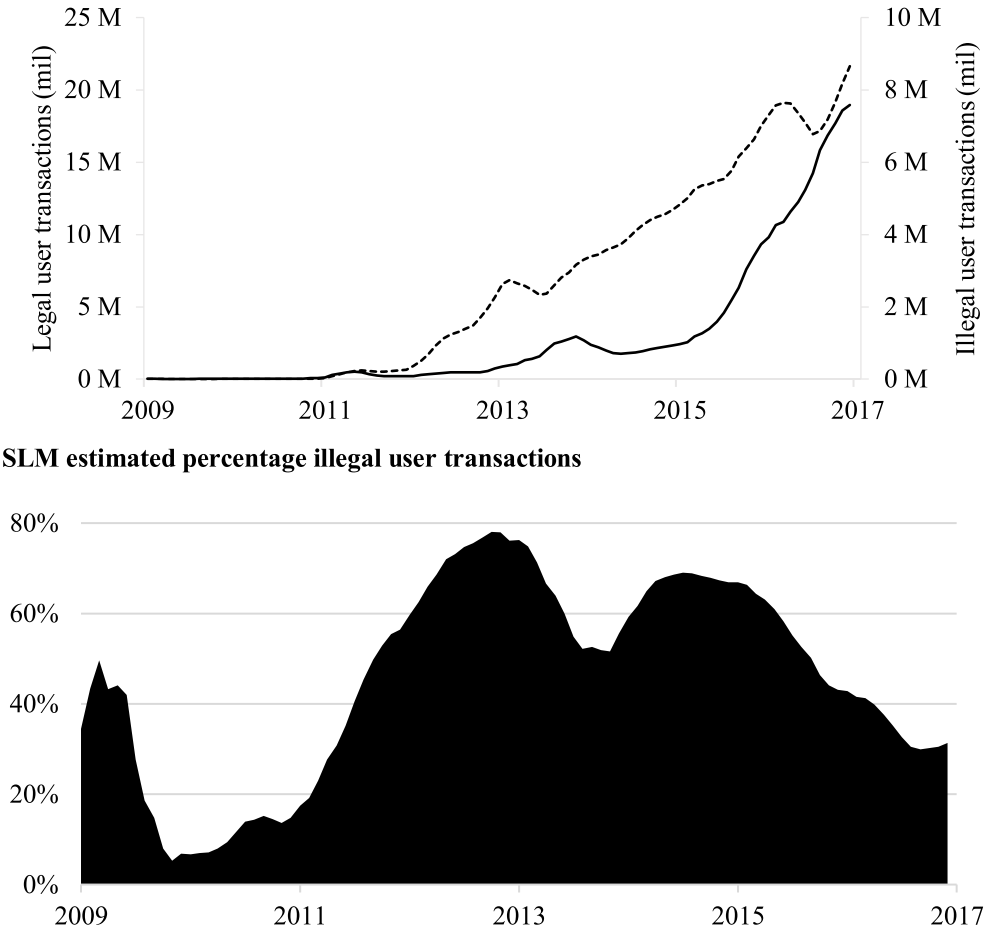
Panel A: SLM estimated number of illegal and legal transactions per month

Panel B: SLM estimated percentage illegal user transactions

Figure C2

SLM estimated number and percentage of illegal bitcoin user transactions

This figure illustrates the time-series of the estimated number of illegal and legal bitcoin user transactions per month (Panel A) and the percentage of illegal user transactions (Panel B). In Panel A, the number of legal user transactions is plotted with the solid line using the left-hand-side axis and the number of illegal user transactions is plotted with the dashed line using the right-hand-side axis. In panel B, the black area is the estimated percentage of illegal user transactions and the white area is the estimated percentage of legal user transactions. The estimates come from a network cluster analysis algorithm (SLM). Values are smoothed with a moving average.

Panel A: SLM estimated dollar volume of illegal and legal bitcoin transactions per month

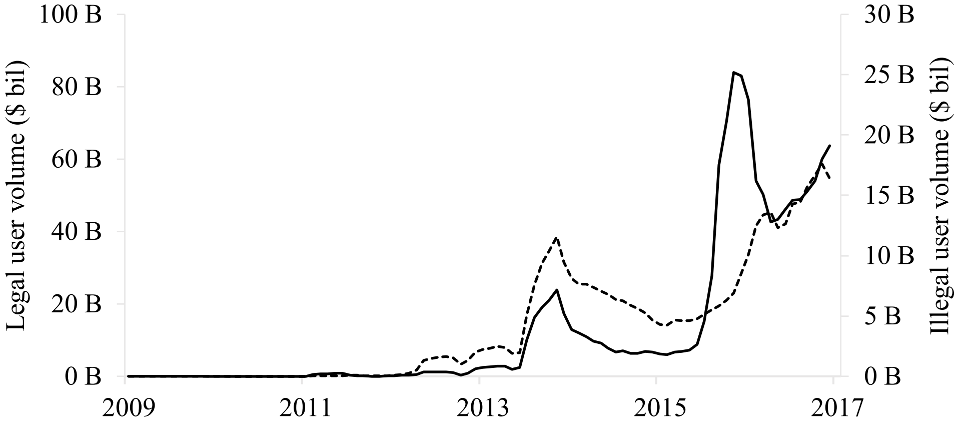

Panel B: SLM estimated percentage illegal user dollar volume

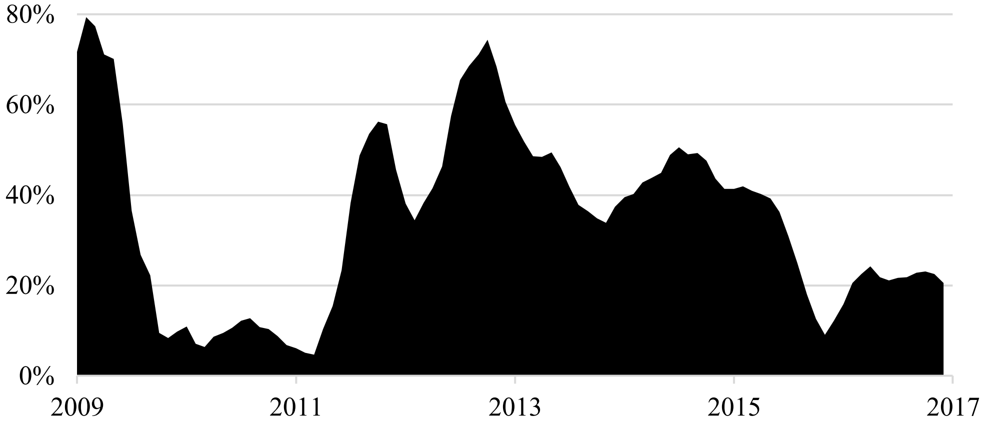

Figure C3

SLM estimated volume and percentage dollar volume of illegal bitcoin user transactions

This figure illustrates the time-series of the estimated dollar volume of illegal and legal bitcoin user transactions per month (Panel A) and the percentage of illegal dollar volume (Panel B). In Panel A, the dollar volume of legal user transactions is plotted with the solid line using the left-hand-side axis and the dollar volume of illegal user transactions is plotted with the dashed line using the right-hand-side axis. In panel B, the black area is the estimated percentage of illegal dollar volume and the white area is the estimated percentage of legal dollar volume. The estimates come from a network cluster analysis algorithm (SLM). Values are smoothed with a moving average.

Panel A: SLM estimated dollar value of illegal and legal user bitcoin holdings

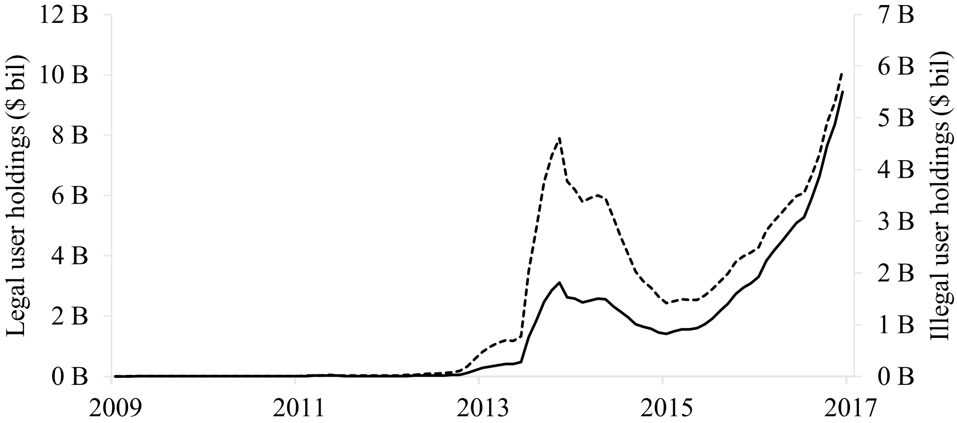

Panel B: SLM estimated percentage illegal user bitcoin holdings

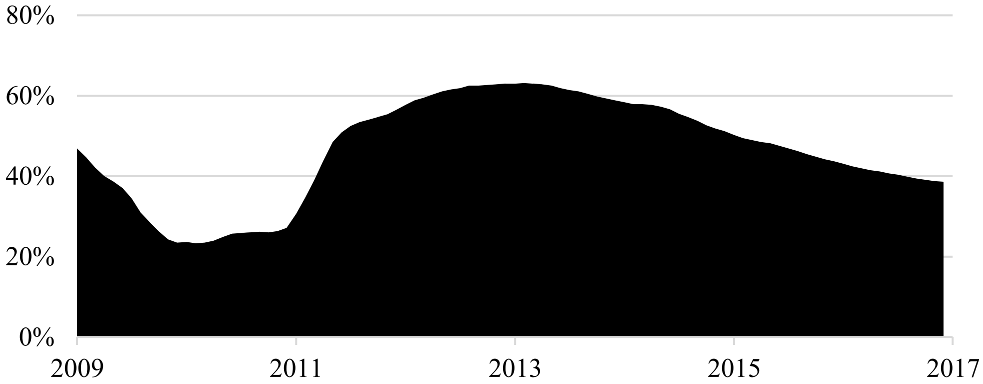

Figure C4
SLM estimated dollar value and percentage of illegal user bitcoin holdings
This figure illustrates the time-series of the estimated dollar value of illegal and legal user bitcoin holdings (Panel A) and the percentage of illegal bitcoin holdings (Panel B). In Panel A, the dollar value of legal user bitcoin holdings is plotted with the solid line using the left-hand-side axis and the dollar value of illegal user bitcoin holdings is plotted with the dashed line using the right-hand-side axis. In panel B, the black area is the estimated percentage of illegal user holdings and the white area is the estimated percentage of legal user holdings. The estimates come from a network cluster analysis algorithm (SLM). Values are smoothed with a moving average.

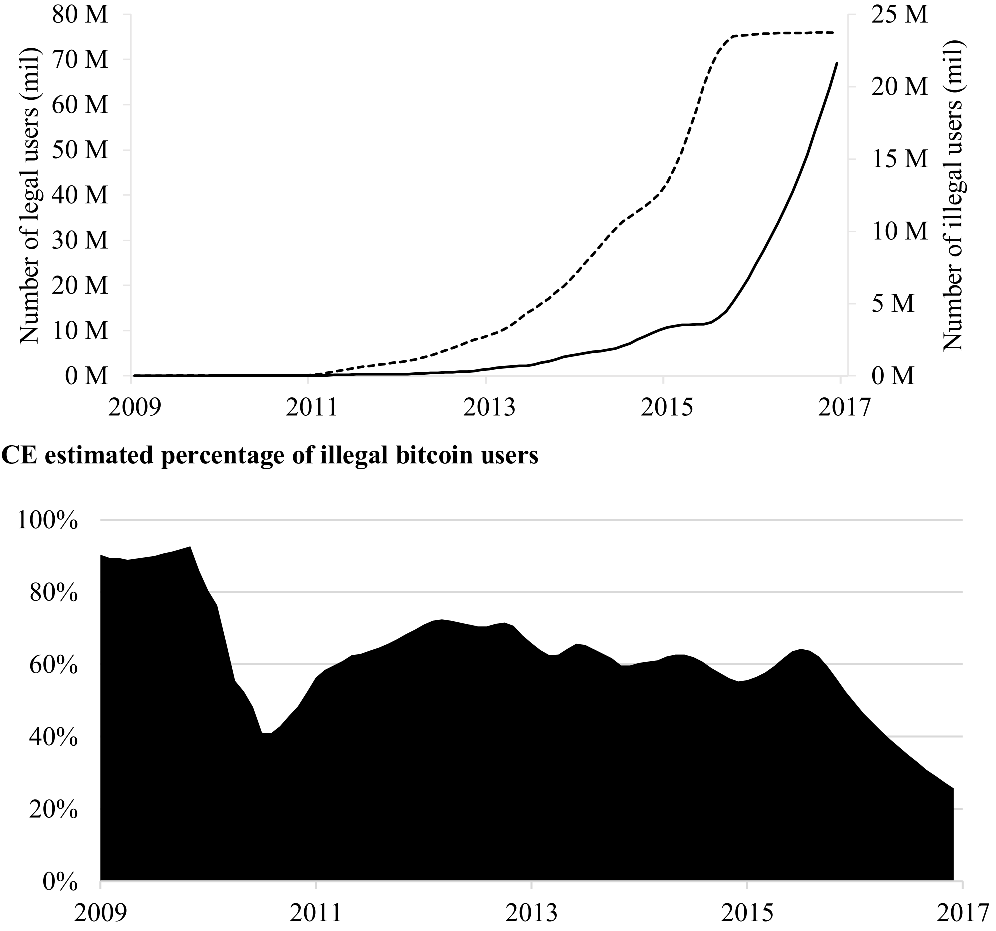

Panel A: DCE estimated number of illegal and legal bitcoin users

Panel B: DCE estimated percentage of illegal bitcoin users

Figure C5

DCE estimated number and percentage of bitcoin users involved in illegal activity

This figure illustrates the time-series of the estimated number of illegal and legal bitcoin users (Panel A) and the percentage of illegal users (Panel B). In Panel A, the number of legal users is plotted with the solid line using the left-hand-side axis and the number of illegal users is plotted with the dashed line using the right-hand-side axis. In panel B, the black area is the estimated percentage of illegal users and the white area is the estimated percentage of legal users. The estimates come from detection-controlled estimation (DCE) model. Values are smoothed with a moving average.

Panel A: DCE estimated number of illegal and legal transactions per month

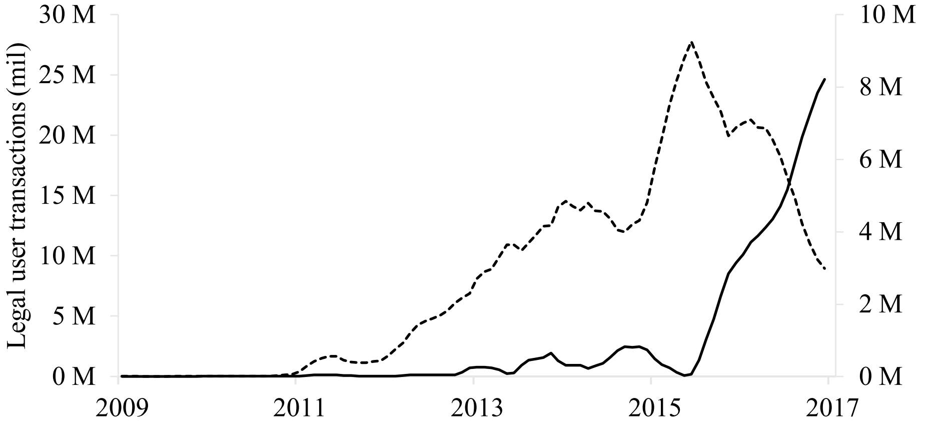

Panel B: DCE estimated percentage illegal user transactions

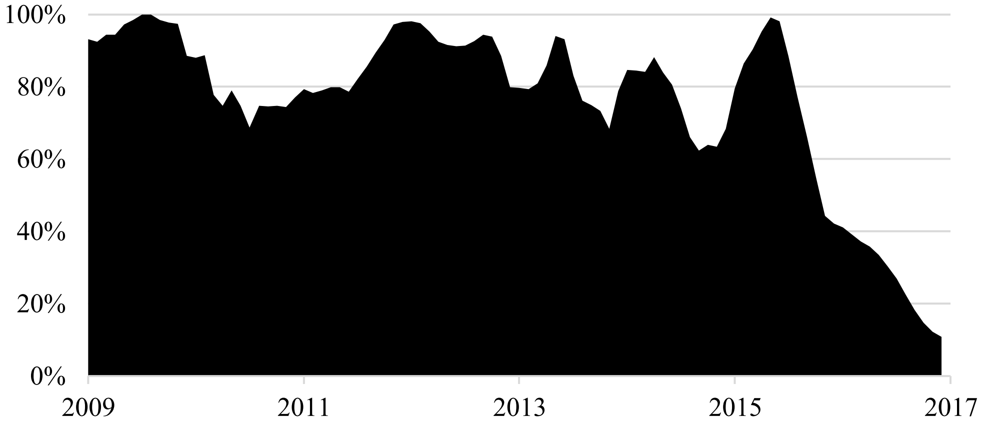

Figure C6
DCE estimated number and percentage of illegal bitcoin user transactions
This figure illustrates the time-series of the estimated number of illegal and legal bitcoin user transactions per month (Panel A) and the percentage of illegal user transactions (Panel B). In Panel A, the number of legal user transactions is plotted with the solid line using the left-hand-side axis and the number of illegal user transactions is plotted with the dashed line using the right-hand-side axis. In panel B, the black area is the estimated percentage of illegal user transactions and the white area is the estimated percentage of legal user transactions. The estimates come from a detection-controlled estimation (DCE) model. Values are smoothed with a moving average.

Panel A: DCE estimated dollar volume of illegal and legal bitcoin transactions per month

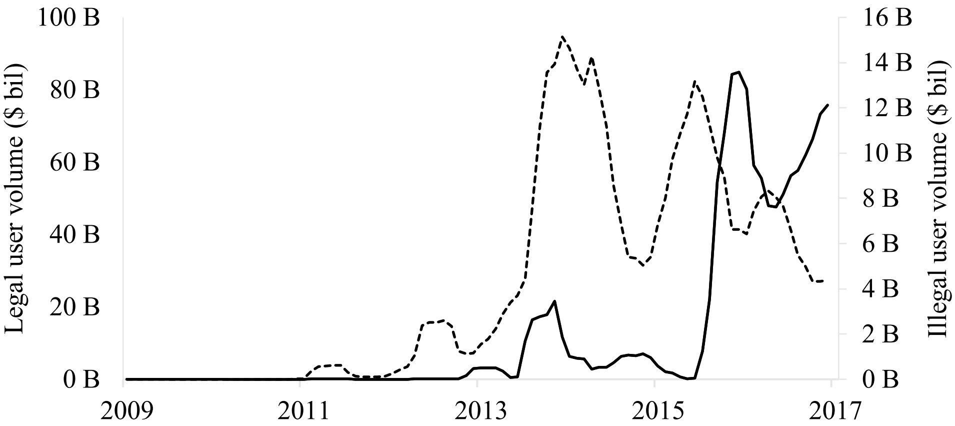

Panel B: DCE estimated percentage illegal user dollar volume

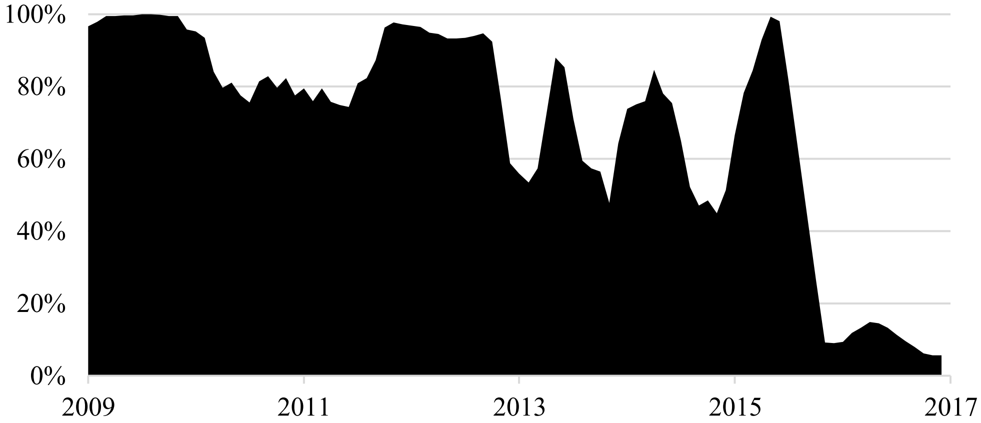

Figure C7

DCE estimated volume and percentage dollar volume of illegal bitcoin user transactions

This figure illustrates the time-series of the estimated dollar volume of illegal and legal bitcoin user transactions per month (Panel A) and the percentage of illegal dollar volume (Panel B). In Panel A, the dollar volume of legal user transactions is plotted with the solid line using the left-hand-side axis and the dollar volume of illegal user transactions is plotted with the dashed line using the right-hand-side axis. In panel B, the black area is the estimated percentage of illegal dollar volume and the white area is the estimated percentage of legal dollar volume. The estimates come from a detection-controlled estimation (DCE) model. Values are smoothed with a moving average.

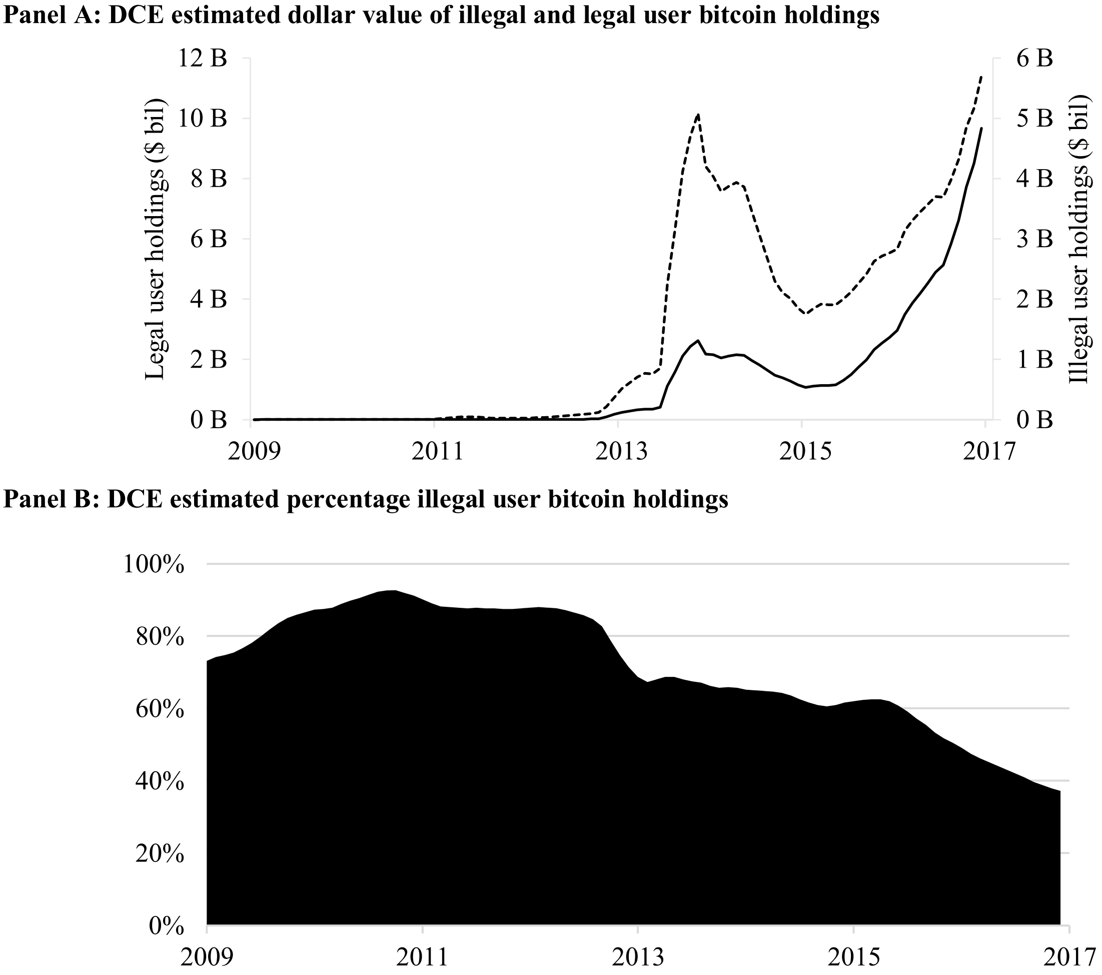

Panel A: DCE estimated dollar value of illegal and legal user bitcoin holdings

Panel B: DCE estimated percentage illegal user bitcoin holdings

Figure C8

DCE estimated dollar value and percentage of illegal user bitcoin holdings

This figure illustrates the time-series of the estimated dollar value of illegal and legal user bitcoin holdings (Panel A) and the percentage of illegal bitcoin holdings (Panel B). In Panel A, the dollar value of legal user bitcoin holdings is plotted with the solid line using the left-hand-side axis and the dollar value of illegal user bitcoin holdings is plotted with the dashed line using the right-hand-side axis. In panel B, the black area is the estimated percentage of illegal user holdings and the white area is the estimated percentage of legal user holdings. The estimates come from a detection-controlled estimation (DCE) model. Values are smoothed with a moving average.

Appendix D

Impact of AlphaBay's adoption of Monero on illegal activity in bitcoin

Table D1: Impact of AlphaBay's adoption of Monero on illegal activity in bitcoin

This table reports difference-in-differences estimates of how illegal activity in bitcoin was impacted by the dark market AlphaBay's adoption of Monero for payments from August 22, 2016. For each day eight weeks either side of the event, we sum the number of transactions (*Transaction count*, measured in thousands) made by illegal and legal users, creating two daily time-series: legal transaction activity and illegal transaction activity in bitcoin. We regress daily *Transaction count* on a dummy variable that is one after AlphaBay started accepting Monero (*Post* a*nnouncement*), a dummy variable that equals one for the illegal transactions (*Illegal*), and an interaction between *Post announcement* and *Illegal* (*Interaction*). The interaction term provides the difference-in-differences estimate. We use three different estimates of the illegal transaction activity: *SLM* is based on the network cluster analysis algorithm (SLM) to identify illegal users, *DCE* is based on user classifications from our detection-controlled estimation (DCE) model, and *Midpoint* is the average of the estimates from the SLM and DCE models. The *t*-statistics are reported in parentheses. ***, **, and * indicate statistical significance at 1%, 5%, and 10% levels respectively.

	Transaction count (Thousands)		
Variable	SLM	DCE	Midpoint
Intercept	388.65***	405.63***	397.14***
	(73.88)	(70.28)	(73.53)
Post announcement	69.52***	79.26***	74.39***
	(9.26)	(9.62)	(9.65)
Illegal	-163.68***	-197.63***	-180.66***
	(-22.00)	(-24.21)	(-23.65)
Interaction	-79.31***	-98.80***	-89.06***
	(-7.47)	(-8.48)	(-8.17)
R^2 (%)	87.54%	89.55%	89.07%

Estimated dollar value and percentage of illegal user bitcoin holdings

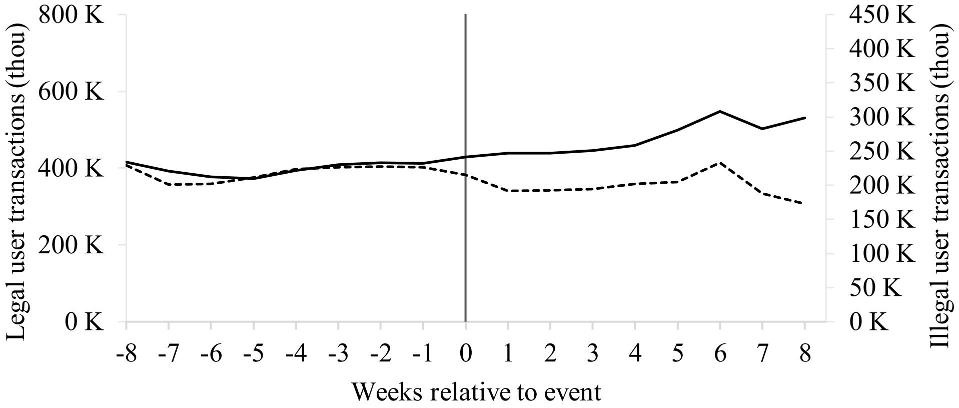

Figure D1
Legal and illegal user transaction activity in bitcoin around AlphaBay's adoption of Monero for payments
This figure illustrates the time-series of legal and illegal user transaction activity in bitcoin eight weeks either side of AlphaBay's adoption of Monero on August 22, 2016. The number of legal user transactions per day is plotted with the solid line using the left-hand-side axis and the number of illegal user transactions per day is plotted with the dashed line using the right-hand-side axis. The estimates are the averages of the SLM and DCE model estimates.

Appendix E

Detection controlled estimation robustness tests

Table E1: Detection-controlled estimation model (DCE) robustness tests

This table reports the coefficient estimates and marginal effects of two DCE models in which various exclusion restrictions imposed in the baseline model are relaxed. Both models use the two-equation structure given by Equations (2.1-2.4) of Chapter 2. Model 1 differs from the baseline specification in that it includes the variable *Darknet sites* in the detection equation. Model 2 differs from the baseline specification in that it includes the variable *Pre-Silk-Road user* in the violation equation. Model 3 differs from the baseline specification in that it includes the variable *Tumbling* in the detection equation. I() is the probability that a given user predominantly uses bitcoin for illegal activity. D() is the conditional probability of detection. Variables are defined in Table 2.1 of Chapter 2. Numbers not in brackets are the coefficient estimates. Numbers in parentheses are the marginal effects (partial derivatives of the corresponding probability with respect to each of the variables, reported as a fraction of the estimated corresponding probability). Pseudo R^2 is McFadden's likelihood ratio index (one minus the ratio of the log-likelihood with all predictors and the log-likelihood with intercepts only). Significance at the 1%, 5%, and 10% levels is indicated by ***, **, and *, respectively using bootstrapped standard errors.

	Model 1		Model 2		Model 3	
Variable	I()	D()	I()	D()	I()	D()
Intercept	-0.255***	0.206***	-1.165***	0.287***	-1.141***	0.248***
	(-0.160)	(0.093)	(-0.771)	(0.135)	(-0.750)	(0.118)
Darknet sites	0.122***	0.358***	0.938***		1.008***	
	(0.076)	(0.162)	(0.620)		(0.662)	
Tumbling	0.006		0.094***		0.105***	-0.071**
	(0.004)		(0.063)		(0.069)	(-0.034)
Bitcoin market cap	-1.494***		-1.694***		-1.612***	
	(-0.936)		(-1.121)		(-1.060)	
Shadow coins	-0.675***		-0.699***		-0.651***	
	(-0.423)		(-0.462)		(-0.428)	
Alt coins	-3.524***		0.551***		0.593***	
	(-2.206)		(0.364)		(0.390)	
Darknet shock volume	0.152***		0.379***		0.447***	
	(0.095)		(0.250)		(0.294)	
Pre-Silk-Road user		0.195***	-0.254***	0.494***		0.429***
		(0.088)	(-0.168)	(0.233)		(0.205)
Transaction frequency	0.591***	0.338***	0.454***	0.462***	0.431***	0.482***
	(0.370)	(0.153)	(0.300)	(0.218)	(0.284)	(0.230)
Transaction size	-0.119***	-0.136***	-0.004	-0.171***	0.014	-0.181***
	(-0.074)	(-0.061)	(-0.003)	(-0.081)	(0.009)	(-0.087)
Concentration	0.457***	0.322***	0.303***	0.524***	0.291***	0.541***
	(0.286)	(0.145)	(0.201)	(0.247)	(0.191)	(0.258)
Existence time	1.694***	0.059***	0.122***	1.728***	0.091***	1.735***
	(1.061)	(0.026)	(0.081)	(0.815)	(0.060)	(0.828)
Pseudo R^2	24.82%		22.02%		21.92%	

Appendix F

Differences in characteristics of illegal users after the Silk Road seizure

Table F1: Differences in characteristics of illegal users after the Silk Road seizure

This table reports difference-in-differences estimates of how the characteristics of illegal bitcoin users change after the FBI's seizure of the Silk Road darknet marketplace in October 2013. We first estimate the average of each of the user-level characteristics for illegal users and legal users separately in each month of our sample, after which we estimate the difference-in-differences regressions. Each user is assigned to the month in which they started transacting in bitcoin. Each of the columns in the table corresponds to a different definition of illegal users: *SLM* uses the classifications of the network cluster analysis (SLM) algorithm, *DCE* uses the classifications of the detection-controlled estimation (DCE) model, and *Observed illegal* uses the directly "observed" illegal users that exist before and after the Silk Road seizure (corresponding to 2B and 2C in Table 2.3 of Chapter 2). The difference-in-differences models regress each user characteristic separately on a dummy variable that is one after the Silk-Road seizure (*Post-Silk-Road*), a dummy variable that is one for the illegal group (*Illegal*), and an interaction term between *Post-Silk-Road* and *Illegal* (*Interaction*). The interaction term provides the difference-in-differences estimate and is reported in the table. User characteristics are defined in Table 2.1 of Chapter 2. The *t*-statistics are reported in parentheses below the coefficient estimates. ***, **, and * indicate statistical significance at 1%, 5%, and 10% levels respectively.

		Interaction	
Dependent Variable	SLM	DCE	Observed illegal
Transaction count	-0.02	0.96***	-1.00***
	(-0.15)	(3.71)	(-4.42)
Transaction size	-782.58***	-654.50***	-867.85***
	(-4.62)	(-4.03)	(-3.94)
Transaction frequency	-3.39**	-0.05	-4.02***
	(-2.50)	(-0.02)	(-2.79)
Holding value	-791.01***	19.33	-1,325.32***
	(-5.70)	(0.11)	(-4.08)
Concentration	0.01	-0.08	0.12**
	(0.23)	(-1.38)	(2.56)
Existence time	2.59*	20.09***	-3.60*
	(1.71)	(10.90)	(-1.78)
Darknet sites	-0.04	1.52*	-0.29
	(-0.04)	(1.75)	(-0.32)
Tumbling	0.02**	0.06*	0.01
	(2.21)	(1.92)	(0.69)
Darknet shock volume	0.00	0.17***	-0.01
	(-0.09)	(3.08)	(-0.27)
Bitcoin market cap	0.00	-0.08*	-0.01
	(-0.08)	(-1.72)	(-0.18)
Shadow coins	0.00	-0.28	-0.10
	(0.00)	(-0.56)	(-0.20)
Alt coins	-0.06	-0.10	-0.26
	(-0.20)	(-0.29)	(-0.76)

Appendix G

HFT trading, colocation, and whistleblower dates

Table G1: HFT and Colocation start dates

The table lists the colocation and HFT start dates extracted from Aitken, Cumming, and Zhan (2015) and Boehmer, Fong, and Wu (2020).

Country	Exchange	Aitken et al. (2015)		Boehmer et al. (2020)
		HFT	Colocation	Colocation
Australia	ASX	200604	2008Q4	200811
Belgium	Euronext Brussels	N/A	N/A	200804
Brazil	BM&FBOVESPA	N/A	N/A	20090629
Canada	TSE	200505	200804	200811
China	Shanghai SE	N/A	N/A	N/A
China	Shenzhen SE	N/A	N/A	N/A
Denmark	Nasdaq Copenhagen	N/A	N/A	20080625
Finland	Nasdaq Helsinki	N/A	N/A	20080625
France	Euronext Paris	N/A	N/A	200804
Germany	XETRA Germany	200301	200608	2006Q4
Hong Kong	Hongkong SE	N/A	2012Q4	N/A
India	Bombay SE	200905	201002	20101115
India	NSE India	200905	201001	200908
Italy	Italian Bourse	N/A	N/A	200909
Japan	Tokyo SE	200505	201001	200905
Japan	Osaka SE	N/A	N/A	200811
Korea	Korea SE	N/A	N/A	N/A
Korea	KOSDAQ	N/A	N/A	N/A
Malaysia	Bursa Malaysia	N/A	N/A	N/A
Netherlands	Euronext Amsterdam	N/A	N/A	200804
New Zealand	New Zealand SE	200411	N/A	N/A
Norway	Oslo SE	200504	201004	N/A
Portugal	Euronext Lisbon	N/A	N/A	200804
Saudi Arabia	Saudi Arabia SE	N/A	N/A	N/A

Table G1: HFT and Colocation start dates (continued)

The table lists the colocation and HFT start dates extracted from Aitken, Cumming, and Zhan (2015) and Boehmer, Fong, and Wu (2020).

Country	Exchange	Aitken et al. (2015)		Boehmer et al. (2020)
		HFT	Colocation	Colocation
Singapore	Singapore SE	N/A	201107	201104
Sweden	Nasdaq Stockholm	200504	201103	20080625
Switzerland	Swiss SE	200401	201204	20080624
Taiwan	Taiwan SE	N/A	2010Q4	2010Q4
UK	Chi-X London	200701	200811	N/A
UK	LSE	200602	200909	200809
UAE	Dubai SE	N/A	N/A	N/A
USA	Nasdaq	200301	200703	200504
USA	NYSE	200305	200804	200701

Table G2: Whistleblower laws

The table lists the first whistleblower law for each country.

Country	Date	Legislation
Australia	2004	Part 9.4AAA the Corporations *Act 2001* (Corporations Act)
Brazil	N/A	N/A
Canada	2005	Public Servants Disclosure Protection Act, SCC 2005, c 46
China	N/A	N/A
France	2017	Law No. 2016-1691 (known as "Sapin II Law")
Germany	N/A	N/A
Hong Kong	N/A	N/A
India	2002	The public interest disclosure (protection of informers) bill 2002
Israel	July, 2008	Protection of Workers (Disclosure of Offenses and Harm to Integrity or to Proper Administration) Law (Amendment No. 2), 5768-2008
Italy	2013	Protection of the Whistleblower Act
Japan	June 18, 2004	Whistleblower Protection Law (Law No. 122 of June 18, 2004).
Malaysia	2010	Whistleblower Protection Act 2010 (Act 711) [P.U.(B) 537/2010]
Netherlands	2001	Art 125 quinquies 1.f. of the Ambtenarenwet
Norway	2005	Working Environment Act (WEA, 2005)
Poland	N/A	N/A
Russia	N/A	N/A
Singapore	2005	Code of corporate governance (2005)
South Africa	2008	Companies Act 71 of 2008
South Korea	February 29, 2008	Act No. 8878, Feb. 29, 2008
Sweden	2016	Proposition 2015/16:128
Switzerland	N/A	N/A
Taiwan	2011	The Anti-Corruption Informant Rewards and Protection Regulation
Thailand	N/A	N/A
UK	1998	Public Interest Disclosure Act 1998
USA	2008	The Sarbanes–Oxley Act of 2002

References

Aboody, D., and B. Lev, 2000, Information asymmetry, R&D, and insider gains, *Journal of Finance* 55, 2747–2766.

Admati, A., and P. Pfleiderer, 1988, A theory of intraday patterns: Volume and price variability, *Review of Financial Studies* 1, 3–40.

Agrawal, A., and T. Nasser, 2012, Insider trading in takeover targets, *Journal of Corporate Finance* 18, 598–625.

Ahern, K.R., 2017, Information networks: Evidence from illegal insider trading tips, *Journal of Financial Economics* 125, 26–47.

Ahern, K.R., 2020, Do proxies for informed trading measure informed trading? Evidence from illegal insider trades, *Review of Asset Pricing Studies* 10, 397–440.

Aitken, M., D. Cumming, and F. Zhan, 2015, High frequency trading and end-of-day price dislocation, *Journal of Banking & Finance* 59, 330–349.

Akey, P., V. Gregoire, and C. Martineau, 2020, Price revelation from insider trading: Evidence from hacked earnings news, *Unpublished manuscript.*

Aldridge, J., and D. Décary-Hétu, 2014, 'Not an Ebay for drugs': The cryptomarket 'Silk Road' as a paradigm shifting criminal innovation, *Unpublished manuscript*.

Aldridge, J., and D. Décary-Hétu, 2016, Hidden wholesale: The drug diffusing capacity of online drug cryptomarkets, *International Journal of Drug Policy* 35, 7–15.

Allingham, M.G., and A. Sandmo, 1972, Income tax evasion: A theoretical analysis, *Journal of Public Economics* 1, 323–338.

Aloosh, A., and J. Li, 2019, Direct evidence of bitcoin wash trading, *Unpublished manuscript*.

Amara, I., and H. Khlif, 2018, Financial crime, corruption and tax evasion: a cross-country investigation, *Journal of Money Laundering Control* 21, 545–553.

Anderson, J.P., 2015, Solving the paradox of insider trading compliance, *Temple Law Review* 88, 273–311.

Anderson, J.P., 2021, Regulatory ritualism and other lessons from the global experience of insider trading law, *Unpublished manuscript.*

Androulaki, E., G. Karame, M. Roeschlin, T. Scherer, and S. Capkun, 2013, Evaluating user privacy in bitcoin, In *Proceedings of Financial Cryptography 2013*.

Augustin, P., M. Brenner, and M. Subrahmanyam, 2015, Informed options trading prior to M&A announcements: Insider trading, *Management Science* 65, 5697–5720.

Austin, J., 2016, What exactly is market integrity: An analysis of one of the core objectives of securities regulation, *William and Mary Business Law Review* 8, 215–240.

Bancroft, A., 2017, Responsible use to responsible harm: Illicit drug use and peer harm reduction in a darknet cryptomarket, *Health, Risk and Society* 19, 336–50.

Barratt, M.J., J.A. Ferris, and A.R. Winstock, 2016a, Safer scoring? Cryptomarkets, social supply and drug market violence, *International Journal of Drug Policy* 35, 24–31.

Barratt, M.J., S. Lenton, A. Maddox, and M. Allen, 2016b, 'What if you live on top of a bakery and you like cakes?'—Drug use and harm trajectories before, during and after the emergence of Silk Road, *International Journal of Drug Policy* 35, 50–57.

Baruch, S., M. Panayides, and K. Venkataraman, 2017, Informed trading and price discovery before corporate events, *Journal of Financial Economics* 125, 561–588.

Basu, G., 2014, The strategic attributes of transnational smuggling: Logistics flexibility and operational stealth in the facilitation of illicit trade, *Journal of Transportation Security* 7, 99–113.

Basu, S., D. Easley, and M. O'Hara, 2019, From mining to markets, *Journal of Financial Economics* 134, 91–109.

Ben-Sasson, E., A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, 2014, Zerocash: Decentralized anonymous payments from bitcoin, In *2014 IEEE Symposium on Security and Privacy.*

Beny, L.N., 2006, Insider trading laws and stock markets around the world: An empirical contribution to the theoretical law and economics debate, *Journal of Corporation Law* 32, 237–300.

Bernile, G., J. Hu, and Y. Tang, 2016, Can information be locked up? Informed trading ahead of macro-news announcements, *Journal of Financial Economics* 121, 496–520.

Bhattacharya, U, H. Daouk, B. Jorgenson, and C.H. Kehr, 2000, When an event is not an event: The curious case of an emerging market, *Journal of Financial Economics* 55, 69–101.

Bhattacharya, U., and H. Daouk, 2002, The world price of insider trading, *Journal of Finance* 57, 75–108.

Boehmer, E., K.Y. Fong, and J. Wu, 2020, Algorithmic trading and market quality: International evidence, *Journal of Financial and Quantitative Analysis*.

Bray, C., 2010, The Galleon case: Kumar says he was paid for tips, *Wall Street Journal*.

Bray, C., and J. Strasburg, 2009, Six charged in vast insider-trading ring—billionaire financier, IBM, McKinsey executives in alleged plot to profit on Google, Hilton; echoes of Ivan Boesky, *Wall Street Journal*.

Bris, A., 2005, Do insider trading laws work, *European Financial Management* 11, 267–312.

Brownbridge, M., and C. Kirkpatrick, 2000, Financial regulation in developing countries, *The Journal of Development Studies* 37, 1–24.

Cao S., L.W. Cong, and B. Yang, 2019, Auditing and blockchains: Pricing, misstatements, and regulation, *Unpublished manuscript*.

Carlton, D.W., and D.R. Fischel, 1983, The regulation of insider trading, *Stanford Law Review* 35, 857–895.

Chakravarty, S., and J.J. McConnell, 1999, Does insider trading really move stock prices?, *Journal of Financial and Quantitative Analysis* 34, 191–209.

Chod, J., N. Trichakis, G. Tsoukalas, H. Aspegren, and M. Weber, 2020, On the financing benefits of supply chain transparency and blockchain adoption, *Management Science* 66, 4359–4919.

Christin, N., 2013, Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace, In *Proceedings of the 22nd International Conference on World Wide Web*.

Christophe, S.E., M.G. Ferri, and J. Hsieh, 2010, Informed trading before analyst downgrades: Evidence from short sellers, *Journal of Financial Economics* 95, 85–106.

Chung, K.H., and C. Charoenwong, 1998, Insider trading and the bid-ask spread, *Financial Review* 33, 1–20.

Collin-Dufresne, P., and V. Fos, 2015, Do prices reveal the presence of informed trading?, *Journal of Finance* 70, 1555–1582.

Collin-Dufresne, P., and V. Fos, 2016, Insider trading, stochastic liquidity, and equilibrium prices, *Econometrica* 84, 1441–1475.

Comerton-Forde, C., and J. Rydge, 2006, Market integrity and surveillance effort, *Journal of Financial Services Research* 29, 149–172.

Comerton-Forde, C., and T.J. Putniņš, 2011, Measuring closing price manipulation, *Journal of Financial Intermediation* 20, 135–158.

Comerton-Forde, C., and T.J. Putniņš, 2014, Stock price manipulation: Prevalence and determinants, *Review of Finance* 18, 23–66.

Cormen, T.H., C.E. Leiserson, R.L. Rivest, and C. Stein, 2001, *Introduction to Algorithms*, Cambridge: MIT Press.

Cox, J., 2016, Staying in the shadows: The use of bitcoin and encryption in cryptomarkets, In *The internet and drug markets, Edited by EMCDDA*. 41–47, *Lisbon: EMCDDA*.

Cumming, D., and S. Johan, 2008, Global market surveillance, *American Law and Economics Review* 10, 454–506.

Daouk, H., C.M. Lee, and D. Ng, 2005, Capital market governance: How do security laws affect market performance?, *Journal of Corporate Finance* 12, 560–593.

De Vries, A., 2018, Bitcoin's Growing Energy Problem, *Joule* 2, 801–805.

Del Guercio, D., E.R. Odders-White, and M.J. Ready, 2017, The deterrent effect of the securities and exchange commission's enforcement intensity on illegal insider trading: Evidence from run-up before news events, *Journal of Law and Economics* 60, 269–307.

Dhawan, A., and T.J. Putniņš, 2020, A new wolf in town? Pump-and-dump manipulation in cryptocurrency markets, *Unpublished manuscript*.

Dissanaike, G., and K.H. Lim, 2015, Detecting and quantifying insider trading and stock manipulation in Asian markets, *Asian Economic Papers* 14, 1–20.

Easley, D., M. López de Prado, and M. O'Hara, 2016, Discerning information from trade data, *Journal of Financial Economics* 120, 269–285.

Emmons, S., S. Kobourov, M. Gallant, and K. Börner, 2016, Analysis of network clustering algorithms and cluster quality metrics at scale, *PLOS ONE* 11, 1–18.

Fajnzylber, P., D. Lederman, and N. Loayza, 1998, Empirical implementation, In *Determinants of crime rates in Latin America and the World: An empirical assessment*, World Bank.

Fajnzylber, P., D. Lederman, and N. Loayza, 2002a, Inequality and violent crime, *Journal of Law and Economics* 45, 1–39.

Fajnzylber, P., D. Lederman, and N. Loayza, 2002b, What causes violent crime?, *European Economic Review* 46, 1323–1357.

Fanusie, Y., and T. Robinson, 2018, Bitcoin laundering: An analysis of illicit flows into digital currency services. Center on Sanctions and Illicit Finance memorandum.

Feinstein, J.S., 1989, The safety regulation of US nuclear power plants: Violations, inspections, and abnormal occurrences, *Journal of Political Economy* 97, 115–54.

Feinstein, J.S., 1990, Detection controlled estimation, *Journal of Law and Economics* 33, 233–76.

Feinstein, J.S., 1991, An econometric analysis of income tax evasion and its detection, *RAND Journal of Economics* 22, 14–35.

Fernandes, N., and M.A. Ferreira, 2008, Insider trading laws and stock price informativeness, *Review of Financial Studies* 22, 1845–1887.

Fishe, R., and M. Robe, 2004, The impact of illegal insider trading in dealer and specialist markets: Evidence from a natural experiment, *Journal of Financial Economics* 71, 461–488.

Fishman, M.J., and K.M. Hagerty, 1992, Insider trading and the efficiency of stock prices, *RAND Journal of Economics* 26, 106–122.

Foley, S., and T.J. Putniņš, 2016, Should we be afraid of the dark? Dark trading and market quality, *Journal of Financial Economics* 122.3, 456–481.

Foley, S., J.R. Karlsen, and T.J. Putniņš, 2019, Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies?, *Review of Financial Studies* 32, 1798–1853.

Franklin, J., A. Perrig, V. Paxson, and S. Savage, 2007, An inquiry into the nature and causes of the wealth of internet miscreants, In *Proceedings of the 14th ACM Conference on Computer and Communications Security.*

Frantz, D., 1987, Levine & Co: Wall Street's insider trading scandal, *Holt*, 1–352.

Goldstein, I., and A. Guembel, 2008, Manipulation and the allocation role of prices, *Review of Economic Studies* 75, 133–164.

Gregoire, P., and H. Huang, 2009, Informed trading, noise trading and the cost of equity, *International Review of Economics & Finance* 17, 13–32.

Hanson, R., and R. Oprea, 2009, Manipulators increase information market accuracy, *Economica* 76, 304–314.

Hardy, R.A., and J.R. Norgaard, 2016, Reputation in the internet black market: An empirical and theoretical analysis of the deep web, *Journal of Institutional Economics* 12, 515–539.

Hillion, P., and M. Suominen, 2004, The manipulation of closing prices, *Journal of Financial Markets* 7, 351–375.

Hofstede, G., 1997, Cultures and organizations, *New York: McGraw-Hill*.

Hofstede, G., 2001, Culture's consequences: Comparing values, behaviors, institutions, and organizations across nations, *Sage Publications*.

Hofstede, G., 2011, Dimensionalizing cultures: The Hofstede model in context, *Online readings in psychology and culture* 2, 1–26.

Hsieh, C.C., and M.D. Pugh, 1993, Poverty, income inequality, and violent crime: a meta-analysis of recent aggregate data studies, *Criminal Justice Review* 18, 182–202.

Huberman, G., J.D. Leshno, and C. Moallemi, 2017, Monopoly without a monopolist: An economic analysis of the bitcoin payment system, *Bank of Finland Research Discussion Paper* 27, 1–52.

Husted, B.W., 1999, Wealth, culture, and corruption, *Journal of International Business Studies* 30, 339–359.

Irvine, P., M. Lipson, and A. Puckett, 2007, Tipping, *Review of Financial Studies* 20, 741–768.

Iyengar, G., F. Saleh, J. Sethuraman, and W. Wang, 2021, Economics of permissioned blockchain adoption, *Unpublished manuscript*.

Kacperczyk, M., and E.S. Pagnotta, 2019, Chasing private information, *Review of Financial Studies* 32, 4997–5047.

Kahan, M., 1992, Securities laws and the social costs of "inaccurate" stock prices, *Duke Law Journal* 41, 977–1044.

Kappos, G., H. Yousaf, M. Maller, and S. Meiklejohn, 2018, An empirical analysis of anonymity in Zcash, In 27th *USENIX Security Symposium, 463–477. Baltimore*.

Kelly, M., 2000, Inequality and crime, *Review of Economics and Statistics* 82, 530–539.

Keown, A.J., and J.M. Pinkerton, 1981, Merger announcements and insider trading activity: An empirical investigation, *Journal of Finance* 36, 855–869.

Khapko, M., and M. A. Zoican, 2017, "Smart" settlement, *Unpublished manuscript*.

Kinsella, D., 2006, The black market in small arms: Examining a social network, *Contemporary Security Policy* 27, 100–117.

Klöck, F., A. Schied, and Y. Sun, 2017, Price manipulation in a market impact model with dark pool, *Applied Mathematical Finance* 24, 417–450.

Koshy, P., D. Koshy, and P. McDaniel, 2014, An analysis of anonymity in bitcoin using p2p network traffic, In *18th International Conference on Financial Cryptography and Data Security*.

Kruithof, K., J. Aldridge, D. Décary-Hétu, M. Sim, E. Dujso, and S. Hoorens, 2016, Internet-facilitated drugs trade, *RAND Corporation*, 21–32.

Kruithof, K., J. Aldridge, D. Décary-Hétu, M. Sim, E. Dujso, and S. Hoorens, 2016, Internet-facilitated drugs trade: An analysis of the size, scope and the role of the Netherlands, *RAND Europe*, 1–167.

Kyle, A.S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.

Ladegaard, I., 2018, Instantly hooked? Freebies and samples of opioids, cannabis, MDMA, and other drugs in an illicit e-commerce market, *Journal of Drug Issues* 48.2, 226–45.

Lavorgna, A., 2016. How the use of the internet is affecting drug trafficking practices, In *Internet and Drug Markets, EMCDDA Insights.*

Leland, H., 1992, Insider trading: Should it be prohibited?, *Journal of Political Economy* 100, 859–887.

Lewman, A., 2016, Tor and links with cryptomarkets, In *Internet and Drug Markets, EMCDDA Insights.*

Malinova, K., and A. Park, 2016, Market design for trading with blockchain technology, *Unpublished manuscript*.

Manne, H., 1966, Insider trading and the stock market, *Free Press*, 99–104.

Martin, J., 2014a, Drugs on the dark net: How cryptomarkets are transforming the global trade in illicit drugs. Berlin: Springer.

Martin, J., 2014b, Lost on the Silk Road: Online drug distribution and the "cryptomarket", *Criminology and Criminal Justice* 14, 351–67.

Martin, J., 2018, Cryptomarkets, systemic violence and the "gentrification hypothesis", *Addiction* 113.5, 797–98.

Matthews, A., R. Sutherland, A. Peacock, J. Van Buskirk, E. Whittaker, L. Burns, and R. Bruno, 2017, I like the old stuff better than the new stuff? Subjective experiences of new psychoactive substances, *International Journal of Drug Policy* 40, 44–49.

Meiklejohn, S., M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G.M. Voelker, and S. Savage, 2013, A fistful of bitcoins: Characterizing payments among men with no names, In *13th ACM Internet Measurement Conference*.

Meulbroek, L.K., 1992, An empirical analysis of illegal insider trading, *Journal of Finance* 47, 1661–1699.

Mittal, H., 2008, Are you playing in a toxic dark pool?: A guide to preventing information leakage, *Journal of Trading* 3, 20–33.

Moiseev, N., A. Mikhaylov, I. Varyash, and A. Saqib, 2020, Investigating the relation of GDP per capita and corruption index, *Entrepreneurship and Sustainability Issues* 8, 780–794.

Morgenson, G., 2006, Whispers of mergers set off bouts of suspicious trading, *New York Times*.

Morselli, C., D. Décary-Hétu, M. Paquet-Clouston, and J. Aldridge, 2017, Conflict management in illicit drug cryptomarkets, *International Criminal Justice Review* 27, 237–54.

Möser, M., K. Soska, E. Heilman, K. Lee, H. Heffan, S. Srivastava, K. Hogan, J. Hennessey, A. Miller, A. Narayanan, and N. Christin, 2018, An empirical analysis of traceability in the monero blockchain, In *Proceedings on Privacy Enhancing Technologies 2018.*

Nakamoto, S., 2008, Bitcoin: A peer-to-peer electronic cash system, *Unpublished manuscript.*

Noether, S., 2015, Ring signature confidential transactions for monero, In *IACR Cryptology ePrint Archive*, 1098.

O'Hara, M., and M. Ye, 2011, Is market fragmentation harming market quality?, *Journal of Financial Economics* 100, 459–474.

Patel, V., and T. Putniņš, 2021, How much insider trading happens in stock markets?, *Unpublished manuscript*.

Pirrong, S.C., 1995, The self-regulation of commodity exchanges: The case of market manipulation, *Journal of Law and Economics* 38, 141–206.

Porta, R.L., F. Lopez-de-Silanes, A. Shleifer, and R.W. Vishny, 1998, Law and finance, *Journal of Political Economy* 106, 1113–11155.

Porta, R.L., F. Lopez-de-Silanes, A. Shleifer, and R.W. Vishny, 1997, Legal determinants of external finance, *Journal of Finance* 52, 1131–1150.

Putniņš, T.J., and A. Sauka, 2015, Measuring the shadow economy using company managers, *Journal of Comparative Economics* 43, 471–490.

Putniņš, T.J., and A. Sauka, 2016, The components and determinants of the shadow economy: Evidence from the Baltic countries, In *Entrepreneurship and the shadow economy*, Edward Elgar Publishing.

Rogoff, K., 2016, The Curse of Cash. (Princeton, NJ: Princeton University Press).

Ron, D., and A. Shamir, 2013, Quantitative analysis of the full bitcoin transaction graph, In *17th Financial Cryptography and Data Security International Conference.*

Saleh, F., 2021, Blockchain without waste: Proof-of-Stake, *Review of Financial Studies* 34, 1156–1190.

Sanyal, R., 2005, Determinants of bribery in international business: The cultural and economic factors, *Journal of Business Ethics* 59, 139–145.

Sharma, A., and S. Pulliam, 2009, Galleon case prompts firms to plug leaks; Intel assures Clearwire on confidential data; Google cuts ties with investor-relations firm, *Wall Street Journal*.

Soares, R.R., 2004, Development, crime and punishment: Accounting for the international differences in crime rates, *Journal of Development Economics* 73, 155–184.

Soska, K., and N. Christin, 2015, Measuring the longitudinal evolution of the online anonymous marketplace ecosystem, In *Proceedings of the 24th USENIX Conference on Security Symposium.*

Stein, R., 2005, The relationship between default prediction and lending profits: Integrating ROC analysis and loan pricing, *Journal of Banking and Finance* 29, 1213–1236.

Takyi-Asiedu, S., 1993, Some socio-cultural factors retarding entrepreneurial activity in sub-Saharan Africa, *Journal of Business Venturing* 8, 91–98.

Tang, T.C., and L. Chi, 2005, Predicting multilateral trade credit risks: Comparisons of logit and fuzzy logic models using ROC curve analysis, *Expert Systems with Applications* 28, 547–556.

Tasca, P., S. Liu, and A. Hayes, 2018, The evolution of the bitcoin economy: Extracting and analyzing the network of payment relationships, *Journal of Risk Finance* 19.2, 94–126.

Tsakumis, G.T., A.P. Curatola, and T.M. Porcano, 2007, The relation between national cultural dimensions and tax evasion, *Journal of International Accounting, Auditing and Taxation* 16, 131–147.

Tzanetakis, M., G. Kamphausen, B. Werse, and R. Von Laufenberg, 2016, The transparency paradox. Building trust, resolving disputes and optimising logistics on conventional and online drugs markets, *International Journal of Drug Policy* 35, 58–68.

Van Buskirk, J., A. Roxburgh, M. Farrell, and L. Burns, 2014, The closure of the Silk Road: What has this meant for online drug trading?, *Addiction* 109, 517–18.

Van Buskirk, J., S. Naicker, A. Roxburgh, R. Bruno, and L. Burns, 2016, Who sells what? Country specific differences in substance availability on the agora cryptomarket, *International Journal of Drug Policy* 35, 16–23.

Van Hout, M.C., and T. Bingham, 2013, 'Surfing the Silk Road': A study of users' experiences, *International Journal of Drug Policy* 24, 524–29.

Van Ness, B.F., R.A. Van Ness, and R.S. Warr, 2001, How well do adverse selection components measure adverse selection?, *Financial Management* 30, 77–98.

Van Slobbe, J., 2016, The drug trade on the deep web: A law enforcement perspective, In *Internet and Drug Markets, EMCDDA Insights.*

Vitell, S.J., S.L. Nwachukwu, and J.H. Barnes, 1993, The effects of culture on ethical decision-making: An application of Hofstede's typology, *Journal of Business Ethics* 12, 753–760.

Waltman, L., and N. Jan Van Eck, 2013, A smart local moving algorithm for large-scale modularity-based community detection, *European Physical Journal B* 86, 1–14.

Wang, J., 1993, A model of intertemporal asset prices under asymmetric information, *Review of Economic Studies* 60, 249–282.

Wang, T.Y., A. Winton, and X. Yu, 2010, Corporate fraud and business conditions: Evidence from IPOs, *Journal of Finance* 65, 2255–92.

Wijaya, D.A., J. Liu, R. Steinfeld, and D. Liu, 2018, Monero ring attack: Recreating zero mixin transaction effect, In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*, 1196–1201.

Yamen, A., A. Al Qudah, A. Badawi, and A. Bani-Mustafa, 2019, The impact of national culture on financial crime, *Journal of Money Laundering Control* 22, 373–387.

Ye, M., 2012, Price manipulation, price discovery and transaction costs in the crossing network, *Unpublished manuscript.*

Yermack, D., 2017, Corporate governance and blockchains, *Review of Finance* 21, 7–31.

Yin, H.S., and R. Vatrapu, 2017, A first estimation of the proportion of cybercriminal entities in the Bitcoin ecosystem using supervised machine learning, In *Proceedings of the 2017 IEEE International Conference on Big Data*, 3690–99.

Yitzhaki, S., 1974, A note on income tax evasion: A theoretical analysis, *Journal of Public Economics* 3, 201–202.