How Strong Are Anti-Money Laundering Regulations in Practice?

Evidence from Cryptocurrency Transactions

Karen Nershi *

December 24, 2020

Abstract

Money laundering creates massive international transfers of wealth and poses a significant international security risk due to its connection to organized crime and the illegal drug trade. The primary way that countries can combat money laundering – anti-money laundering enforcement – requires international coordination, since one country's lax enforcement enables criminals to continue laundering funds; however, one hindrance to achieving widespread international cooperation is the lack of reliable data with which to compare countries' enforcement. To address this problem, I collect a new dataset from a novel data source – cryptocurrency transactions from over 60 cryptocurrency exchanges – to measure cross-national enforcement of a new customer due diligence law for cryptocurrency. I use bunching estimation and instrumental variable estimation to measure how users adjust the amount of their crypto-to-fiat (government-issued) transactions to avoid a 1,000 dollar/euro threshold that triggers customer due diligence. I find significant evidence of suspicious activity (consistent with efforts to avoid due diligence screening) across exchanges, suggesting that both individuals and private sector businesses respond strategically to customer due diligence laws. I also find significant variation in cross-national levels of suspicious activity, suggesting that national laws do not serve as a good measure of how well countries address money laundering risk. Lastly, I find significant evidence of suspicious activity in exchanges in developed countries, suggesting that lax intermediary enforcement in developed countries is a major weakness in the international anti-money laundering regime.

1 Introduction

In February 2014, news broke that Mt. Gox – the biggest cryptocurrency business in the world at the time – had been hacked. The hack resulted in the theft of over $460 million in customer cryptocurrency, which criminals later laundered (i.e., integrated the stolen funds

*Ph.D. Candidate in Political Science at the University of Pennsylvania

into the legal economy) through another exchange.[1] Six years later, the stolen funds have still not been recovered.[2] Although cryptocurrency makes up only a fraction of all the laundered funds each year, it is an important new area of money laundering risk. Moreover, it presents a useful tool for the study of money laundering, which, despite its significance, has been held back by a lack of reliable data. In this paper, I seek to address this gap in the literature by collecting a new dataset of cryptocurrency transactions directly from exchanges and measuring cross-national enforcement of a new anti-money laundering law for cryptocurrency.

Money laundering, which can be defined simply as the integration of illegal funds into the legitimate economy, first emerged as a major international problem in the 1980s stemming from the illegal drug trade. In the decades since, increased capital mobility has made it easier for criminals to hide wealth abroad.[3] Money laundering primarily harms society through its connection to "predicate crimes" that generate funds for laundering, which also affect international finance and security. Corruption, a common source of predicate crimes, leads to depleted national coffers in many developing countries, while tax evasion limits the ability of governments to pursue redistributive policies. Crimes associated with the illegal drug trade and organized crime, meanwhile, contribute to instability, violence, and loss of government control in some countries; this has led some scholars to argue that organized crime poses a greater threat to international security in the twenty-first century than traditional wars.[4]

To address these externalities, states have banded together to form an anti-money laundering regime. The regime centers on states, which pass and enforce national laws designed to catch criminals and deter future crimes. States also supervise a wide range of intermediaries – private-sector actors including bankers, lawyers, and other finance professionals – who are legally bound to screen customers for money laundering risk (customer due diligence laws). Successful enforcement cannot be achieved at the state level alone, though, since one weak link in the global regime allows criminals to continue laundering money. Instead, states must coordinate internationally, and the G7 countries responded to this challenge by founding an international organization dedicated to promoting cooperation in anti-money laundering enforcement – the Financial Action Task Force (FATF). However, over three decades after the FATF's founding, the dynamics of international cooperation for this issue remain poorly understood.

One key reason behind this failure is that scholars have focused on international and state actors while overlooking the role of sub-state actors. Accordingly, the literature has missed an important weakness in the anti-money laundering regime – intermediaries failure to enforce customer due diligence laws.[5] I seek to address this gap by using intermediary-

[1] Cryptocurrency exchanges allows users to convert cryptocurrency to fiat (government-issued) currency (and vice versa).

[2] Baydakova 2020; McMillan 2018.

[3] Reuter 2005, p. 78.

[4] Williams 1994; Levitsky 2003; Engvall 2006; Swanstrom 2007.

[5] One important exception is Findley, Nielson, and Sharman (2014), who use field experiments to measure compliance with customer due diligence laws by corporate service providers. They follow this study with an examination of customer due diligence compliance by banks (Findley, Nielson, and Sharman 2020).

level data – a new dataset of transactions collected from over 60 cryptocurrency exchanges – to study cross-national enforcement of a new customer due diligence. I use bunching estimation and instrumental variable estimation and find evidence of significant variation in levels of suspicious activity (consistent with efforts to launder money) across countries. My findings support the conclusion that national laws *do not* serve as a good measure of a country's efforts to manage money laundering risk, which holds important implications for international cooperation in anti-money laundering enforcement. I also find high levels of suspicious activity in (some) developed countries, which seems to contradict a dominant view in the literature that developed countries effectively enforce anti-money laundering laws. In short, my research suggests that scholars should place greater emphasis on the role of intermediaries and seek new sources of data to empirically test questions about anti-money laundering enforcement.

1.1 Contributions

This research offers three main contributions to the money laundering literature. First, it makes use of a new data source to study anti-money laundering enforcement – cryptocurrency transactions. One of the primary challenges scholars face is a lack of reliable data that can shed light on either money laundering or anti-money laundering enforcement.[6] I address this challenge by creating a new dataset of cryptocurrency transactions collected directly from over 60 exchanges. This dataset presents the opportunity to measure how well intermediaries enforce customer due diligence laws using transaction-level data.

The second main contribution of this paper lies in testing fundamental questions about how actors respond to anti-money laundering regulation. I find evidence that suggests individuals strategically respond to regulation in an effort to avoid due diligence screening, while intermediaries selectively comply with regulations (actively enforcing easily-verified measures while neglecting others). This finding is significant since most of the anti-money laundering literature has failed to analyze the behavior of sub-state actors. Scholars have typically overlooked the role of individuals altogether and undertheorized that of intermediaries – either assuming that intermediaries uniformly enforce national laws or can be divided into good and bad actors. My findings suggest that anti-money laundering laws should take into account strategic behavior by individuals and intermediaries, and scholars should carefully consider the role of sub-state actors in the anti-money laundering regime.

The third main contribution of this research lies in highlighting an important weakness of the global anti-money laundering regime – the failure of intermediaries to enforce customer due diligence laws. While the international community has mounted an impressive coordinated effort to address money laundering through the FATF, I argue that a major flaw in the FATF's approach is its use of national laws (rather than enforcement) to mark defecting countries for blacklisting. This choice is important, since international organizations can promote cooperation by identifying defecting countries for punishment by member states;

[6]See Reuter (2013) for a discussion of the problems with money laundering estimates. Takáts (2011) notes problems with cross-national comparisons of suspicious activity reports. Deleanu (2017) discusses the possibility that some countries may fake their money laundering statistics.

however, the success of this mechanism depends on an organization's ability to identify defection.[7] If an organization cannot identify defection, then its efforts to promote cooperation by sharing information about state enforcement will be severely limited. My research also suggests that intermediaries in industrialized countries often fail to effectively manage money laundering risk, calling into question the strength of anti-money laundering enforcement in industrialized countries.

The rest of this paper is divided into four main sections. First, I provide background on cryptocurrency and why it provides a good case for studying anti-money laundering enforcement. I then detail my predictions and present my research design. Lastly, I discuss the results and provide a few concluding thoughts.

2 Cryptocurrency

Cryptocurrency is a form of digital money that uses encryption to generate new units and secure transactions.[8] Use of cryptocurrency has increased dramatically since Bitcoin's introduction in 2009, with 1,670 unique cryptocurrency coins today worth an estimated $197 billion in market capitalization.[9] Although the vast majority of transactions are legal, cryptocurrency has been used extensively in illegal activity. One major source of illicit cryptocurrency is dark web marketplaces, which allow users to anonymously buy and sell illegal goods like weapons, stolen credit cards, and drugs using cryptocurrency and encrypted IP addresses.[10] These sites have continued to attract new users since the first dark web market debuted in 2011, with sales for these markets generating at least $660 million in 2017.[11][12]

Cybercrime is another major source of illicit cryptocurrency, and many of these crimes are directly enabled by cryptocurrency. During ransomware attacks, for example, hackers encrypt a victim's computer system and demand a cryptocurrency ransom to decrypt the victim's files. These atttacks are often incredibly costly, leading some 17% of private businesses and 45% of public offices to pay a cryptocurrency ransom according to a recent survey.[13] [14] Other types of cybercrime have also been on the rise, including fraud, scams, and thefts connected to cryptocurrency exchanges. In addition to misappropriation and theft by exchange operators, outside hackers have stolen tens of millions of dollars from exchanges, sometimes

[7] Garrett 1992.

[8] Greenberg 2017.

[9] *All Cryptocurrencies* 2019.

[10] Greenberg 2014; Cox 2015.

[11] United Nations 2019; Soska and Christin 2015; Greenberg 2015.

[12] This estimate is provided by Chainalysis Team (2018), but estimates vary. In 2018, FinCen estimated that at least $4 billion had passed through the Dark Web since 2011 (Ott 2019). A 2019 study estimates that $76 billion worth of illegal activity involves Bitcoin annually (Foley, Karlsen, and Putniņš 2019).

[13] CyberEdge 2019.

[14] In 2019, the City of Baltimore became the victim of such an attack that brought down the city's payment system for water bills, property taxes, and vehicle citations and prevented city employees from accessing email and a parking fine database (Chokshi 2019). Though Baltimore refused to pay the 13 Bitcoin ($76,000) ransom, it paid $6 million to restore computer systems and an estimated $18 million in damages (Broadwater 2019).

forcing exchanges to close down and and leave their users bereft of funds. All told, cybercrime generated at least $4.26 billion in 2019, producing significant public and private costs.[15]

Once cyberthieves or dark web sellers obtain illicit cryptocurrency, they must convert the funds into fiat (government-issued) currency to use it in the broader economy. However, cryptocurrency presents a paradox for money launderers. On the one hand, cryptocurrency transactions are pseudonomynous and identified only through digital keys that are not linked to public identities. On the other hand, most cryptocurrencies record information about each transaction – including the time, date, amount, and keys of senders and recipients – in decentralized ledgers that create a permanent public record of all transactions; accordingly, law enforcement and researchers who manage to match a digital key to a public identify can often learn much about a person's previous transactions.[16] Thus, those who wish to launder illicit funds must convert cryptocurrency to fiat currency without drawing attention to their legal identities or previous criminality, which many have done using cryptocurrency exchanges.[17]

2.1 Why Cryptocurrency?

Cryptocurrency presents a good case for studying anti-money laundering enforcement given (1) the regulatory environment and (2) the available data. First, because cryptocurrency is a new kind of financial instrument, it has not been subject to significant regulation until now. Although some in the law enforcement community voiced concerns about cryptocurrency's potential money laundering risk, only Japan had implemented a comprehensive regulatory framework by 2019. This changed in June 2019, when the FATF issued a new set of legal and regulatory recommendations for cryptocurrency exchanges, which allow users to trade cryptocurrency for fiat currency and vice versa, as well as make crypto-to-crypto trades. The FATF's 35 member states agreed to implement these measures within a year.

This new regulation presents a good case for studying cross-national anti-money laundering enforcement since the new law was introduced by an external party (the FATF) and countries agreed to implement the measures within a given time period (by July 2020). This is significant since it is often difficult to study the effect of an international law or regulation given selection bias – countries that are more inclined to enforce a certain kind of regulation are also more likely to adopt it in the first place. Similarly, the fact that countries typically adopt laws over time and these laws often differ from country to country further complicates cross-national comparisons of enforcement. However, the fact that these countries have agreed to implement a similar type of law in a given time period makes cryptocurrency an ideal case for studying cross-national enforcement.

[15] CipherTrace Cryptocurrency Intelligence 2018, p. 4.

[16] Although users can make unlimited keys (up to one per transaction), most are not so careful (Greenberg 2017).

[17] Laundering through exchanges has been documented for ransomware attackers (Apuzzo 2014), cyberthieves, and dark web marketplace sellers (Meiklejohn et al. 2013).

Second, cryptocurrency presents a good case for studying anti-money laundering enforcement given the available data. Although most intermediaries like banks and law firms typically withhold customer data, cryptocurrency exchanges generally share public information about each transaction's price and amount via each site's application programming interface (API). This data presents an unprecedented chance to study anti-money laundering enforcement using transaction-level data.

3 Predictions

The FATF's new cryptocurrency regulation details two main obligations for exchanges: performing customer due diligence for transactions of 1,000 euros/dollars or more and performing risk-based measures. Customer due diligence requires that exchanges obtain information about a customer's identity "using reliable, independent source documents, data or information," understand the nature of a customer's business, and maintain records of this information. Risk-based measures, meanwhile, require exchanges to "identify, assess, and take effective action to mitigate their money laundering/terrorist financing risks" and conduct customer due diligence for additional transactions that they deem high risk.[18] I hypothesize that exchanges are likely to enforce customer due diligence at the required threshold but may fail to perform additional risk-based measures.

I argue that exchanges are likely to enforce customer due diligence at the threshold since this is a concrete task that regulators can easily verify by checking an exchange's records, and exchanges seek to avoid potential fines for failure to carry out due diligence. However, exchanges may not perform additional risk-based measures since these requirements are vague and hard for an external party to verify. I argue that exchanges are unlikely to undertake risk-based measures unless national regulators require it though active regulation; thus, we should expect to see "suspicious activity" in regulated exchanges as some individuals seek to avoid due diligence and exchanges fail to identify them through risk-based measures.[19]

3.1 Suspicious Activity in Regulated Exchanges

First (and fundamental to this research design), I predict that criminals will use regulated exchanges to launder illegal cryptocurrency even after the adoption of the new guidelines. Money launderers typically face a tradeoff between the security of an investment and the risk of detection by authorities since many of the world's safest and most lucrative investments are located in industrialized countries with extensive anti-money laundering regulation. Criminals typically respond by investing in assets that provide the highest premium of secrecy for an acceptable level of security.[20] While there is greater flexibility for criminals laundering other types of funds, cryptocurrency launderers have few options other than to use exchanges.

[18] Mnuchin 2019.

[19] I refer to any measures consistent with efforts to avoid due diligence screening as suspicious activity. Although this is not direct evidence of money laundering, it is consistent with the behavior of those seeking to launder illegal cryptocurrency

[20] Masciandaro, Takats, and Unger 2007, p. 155.

Although some have argued that money launderers will shift operations from exchanges to peer-to-peer trading sites on the dark web,[21] a mass migration is unlikely because peer-to-peer sites are less convenient (users must arrange each transaction) and riskier (there is no third-party guarantee). Thus, in the face of new regulation, cryptocurrency launderers will face a choice between risky, unregulated exchanges or peer-to-peer sites and secure, regulated exchanges. I predict that many criminals will respond strategically by shifting laundering activities to regulated exchanges with lax compliance.

3.2 Cross-National Variation

Second, I predict that we will see significant cross-national variation in levels of suspicious activity. Because anti-money laundering enforcement is costly and offers no direct benefits, intermediaries will likely only enforce hard-to-verify measures if national regulators require it through active regulation. In countries where regulators adequately supervise and sanction financial institutions, we should expect intermediaries to invest a higher premium into compliance to avoid possible punishment from regulators. As a result, these systems should better detect and manage money laundering risks. Conversely, in countries with lax standards, we should expect intermediaries to devote fewer resources to compliance and have less effective systems as a result.

Thus, I predict that exchanges will likely partially comply with due diligence requirements but fail to perform additional risk-based measures that are required to maintain effective compliance systems. Notably, this view differs from previous research, which often performs analysis at the level of national laws and assumes that intermediaries comply with these laws[22] or that intermediaries actively seek to avoid money laundering risk to minimize financial risk and safeguard institutional reputations.[23] Research that *has* considered adverse incentives for intermediaries has sometimes characterized a dichotomy of "criminal" and "honest" intermediaries that either openly collude with criminals or diligently abide by national laws.[24] Rather than a dichotomy, I argue that many intermediaries will only partially comply with due diligence obligations, creating an important weakness in the anti-money laundering regime.

3.3 Industrialized Countries

Third and last, I predict that industrialized countries will show significant levels of suspicious activity. Some have argued that industrialized countries generally do a good job of enforcing anti-money laundering laws because of their extensive capabilities (including existing financial infrastructure), reputation for low levels of corruption and high rule of law, or because money laundering poses financial risks that industrialized countries will seek to avoid.[25] However, these arguments are unconvincing for several reasons. First, there is no

[21] Havilland 2019; Aguilar 2019.
[22] Unger and Ferwerda 2008; Schwarz 2011; Gnutzmann, McCarthy, and Unger 2010; Takáts 2011.
[23] Morse 2019.
[24] Reuter 2005; Masciandaro, Takats, and Unger 2007, p. 30.
[25] Morse 2019.

conclusive evidence that money laundering causes economic harm for states,[26] calling into question self-interest as a possible motivation for industrialized countries' anti-money laundering enforcement. Both data leaks (like the Panama Papers) and a series of experiments by Findley, Nielson, and Sharman (2014; 2020) show evidence of widespread lapses in intermediaries' enforcement of customer due diligence laws, including in developed countries. Lastly, a series of high profile money laundering scandals in the last decade – including the HSBC 2012 case and the Danske Bank case – show systematic failures by banks in industrialized countries to enforce customer due diligence laws. Given this evidence, I predict that exchanges in industrialized countries will show significant levels of suspicious activity.

4 Data

To test these predictions, I have collected a new dataset of cryptocurrency trades using public trade data from 60 of the world's biggest exchanges that offer crypto-to-fiat trades. To automate this process, I set up remote servers using Amazon Web Services that collected trades every 15, 30, 60, or 150 seconds, with increments chosen based on the volume and number of trades available from each site's application programming interface (API). I collected trades from the two most widely used cryptocurrencies – Bitcoin and Ethereum – to FATF member states' fiat currencies between 07/01/2020 and 08/31/2020. Each trade in the dataset includes information about the time of the trade (typically given to the second), the quantity of cryptocurrency, and the price of the currency pair at the time the trade was executed.[27] Because many exchanges had low daily trade volumes that prohibited analysis using these methods, the final sample includes 23 exchanges located in 8 countries. This sample represents all countries with major exchanges offering crypto-to-fiat trades and presents a diverse cross-section of countries, including wealthy, industrialized countries (US, UK, and Japan), a middle-upper income country (Estonia), a developing country (China), and several tax havens (St. Kitts and Nevis, the British Virgin Islands, and Singapore).

Because some experts have raised concerns about the validity of data released by cryptocurrency exchanges,[28] I have taken several additional steps to ensure the validity of this data. First, unlike most cryptocurrency research that uses aggregate exchange data often acquired from a third-party site, I have collected trade-level data directly from exchanges in real-time. This approach minimizes the chance that an exchange might artificially inflate its daily numbers when sharing it with a third-party service. Second, I have examined the data for statistical anomalies – such as distributions that appear to mimic a perfect exponential distribution – and removed these trading pairs from the analysis.[29] Lastly, even if an exchange included some fake trade data, this alone would not bias the results of my analysis unless the fake data systematically increased the number of transactions below the

[26] Reuter 2013.

[27] Prices vary by site and are not pegged to a currency.

[28] Kauflin 2019; *The Anatomy Of A Fake Cryptocurrency Trade: How Exchanges Create Phony Transactions* 2019.

[29] I removed a handful of trading pairs involving the Russian ruble, which closely resembled an exponential distribution.

due diligence threshold, which seems unlikely. Reassuringly, the results of my analysis differ across exchanges, indicating unique underlying trade data from each exchange.

5 Research Design

I use bunching analysis to compare the predicted distribution of trades close to the threshold to the actual distribution. I find that many exchanges consistently show bunching of trades just below the 1,000 dollar/euro due diligence threshold. However, this varies across exchanges, with Chinese and Japanese exchanges showing the highest levels of bunching below the threshold (20 times the predicted amount), while on the low end, US exchanges show no abnormal bunching. I also use instrumental variable regression to measure whether the threshold moving below a bin is associated with a decrease in the proportion of trades in the bin. I find that there is a significant and negative effect when the threshold passes below a bin.

While a few studies have analyzed the laundering process for cryptocurrency transactions directly linked to crime,[30] I choose instead to analyze a much broader sample of trades across all exchanges with high trading volumes by uncovering behavior consistent with efforts to avoid screening (what I term "suspicious activity"). Importantly, this behavior need not be criminal and may instead reflect a person's preference for avoiding screening for other reasons (e.g., privacy). However, we should not expect most people to avoid due diligence screening, because although it may create a momentary annoyance, screening is relatively easy (sharing a copy of a government-issued form of identification and answering a few basic questions) and clears individuals for all future transactions of any amount. Thus, we can think of customer due diligence as akin to TSA PreCheck for airline passengers – it presents an initial hassle, but it enables smoother interactions in the future. Further, while avoiding screening for some trades is straightforward, it can become increasingly complex for others, requiring knowledge of the currency in which the exchange enforces due diligence and conversions across several exchange rates. Thus, it is unlikely that most people would be willing to go through the trouble of avoiding due diligence unless they had something to hide.

5.1 Bunching Estimation

Bunching estimation is an econometric strategy introduced by Saez (2010) and further developed by Chetty et al. (2011). It exploits a discontinuity in incentives at a cutoff point where individuals can sort below the cutoff and uses the mass of a distribution to measure how individuals respond to a change in incentives. It can be used to study phenomenon that involve evasion or avoidance, and most commonly, researchers have used this approach to study how actors respond to changes in tax brackets.[31] Specifically, actors who wish to avoid being taxed at a higher rate may report incomes that fall slightly below the cutoff that marks the start of a higher tax bracket, leading to excess mass ("bunching") in the distribution just before the threshold. One important finding of this literature is that optimization frictions (the cost of switching below the threshold) often prevent actors from adjusting strategically

[30] Meiklejohn et al. 2013; Apuzzo 2014.

[31] Kleven 2016.

Figure 1: Density Distribution

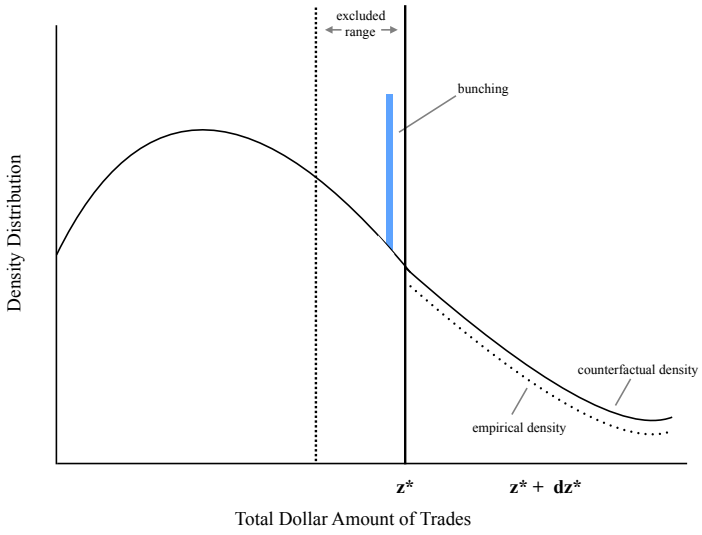

to an incentive. For example, studies have found bunching in the reported income of self-employed workers but not other workers', as self-employed people can more easily adjust the number of hours of worked (or their reported incomes) than those working for companies.[32]

Figure 1 provides a visual representation of this strategy. First, I assume that the number of trades can be represented by a smooth density distribution $h(z)$ across a continuous variable z, which denotes the total fiat amount of a trade.[33] However, because of the introduction of the due diligence threshold at z^*, we see bunching for trades that would have fallen in the range of $[z^*, z^* + d(z)]$ as users adjust the quantity of their transactions so the total fiat amount of a trade falls below z^* (thus avoiding due diligence). This creates excess mass below the threshold, and, consequently, shifts the distribution beyond $z*$ downward. Because there is some randomness in how individuals choose to adjust their transactions, bunching may more closely resemble a hump rather than a spike in some exchanges.[34]

To estimate the level of bunching before the threshold, one must first estimate a counterfactual density that shows what the density would look like in the absence of a due diligence threshold ($h(z)$). I follow Chetty et al. (2011)'s procedure by fitting a flexible polynomial to the distribution of data *excluding* a region below the threshold where bunching may occur. This model takes the following form:

[32]Chetty et al. 2011; Bastani and Selin 2014; Kleven and Waseem 2013.

[33]Importantly, bunching estimation does not require that a researcher know the global distribution of z, but rather can approximate the local distribution within the bunching window (Kleven 2016).

[34]Bastani and Selin 2014.

Figure 2: Simulated Density Distributions with Excess Mass

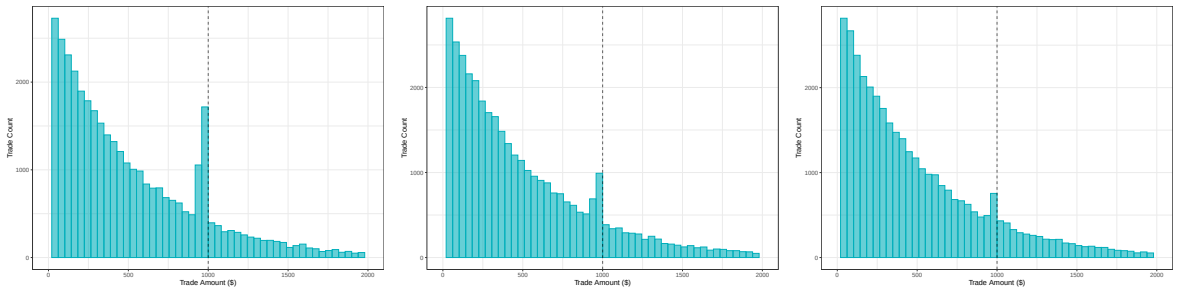

(a) (b) (c)

Caption: These density plots show simulated bunching at 5% (a), 2% (b), and 1% (c) excess mass below the threshold (dashed line). Simulations are made using an exponential distribution of 10,000 trades with a mean of 500.

$$C_j = \sum_{i=0}^{p} \beta_i^0 * (Z_j)^i + \sum_{i=-R}^{0} \gamma_i^0 * 1[Z_j = i] + \epsilon_j^0, \tag{1}$$

where C denotes the number of transactions in bin j, Z denotes the total fiat amount of a trade in 10 euro or dollar increments, p is the order of the polynomial, and $[-R,0]$ is the excluded range (100 dollars or euros) below the threshold. The initial estimate of the counterfactual density is the predicted values from (1) excluding the contribution of the dummies in the range below the threshold $\hat{C_j^0} = \sum_{i=0}^{p} \hat{\beta_i^0} * (Z_j)^i$. Accordingly, the excess number of the transactions in the range below the threshold can be calculated relative to the counterfactual density as $\hat{B_N^0} = \sum_{j=-R}^{0} C_j - \hat{C_j^0} = \sum_{i=-R}^{0} \hat{\gamma_i^0}$. However, this overestimates the amount of excess mass below the threshold because it does not account for the downward shift of the distribution after $z*$. To address this, I compare the excess mass around the kink relative to excess mass in the rest of the distribution.

The Fine Print

While the canonical approach assumes that only individuals in a given range above the threshold will adjust their transactions in response to the threshold, I make no such assumption for cryptocurrency trades. This relates to optimization frictions – the costs that individuals face for sorting below the threshold.[35] While individuals reporting taxable income (the typical setting in which bunching analysis has been applied) face greater adjustment costs for reporting an income below the threshold if their income falls far from the threshold, those making cryptocurrency transactions face no such constraints. Thus, I assume that adjusting the quantity of a cryptocurrency transaction so that the dollar amount falls below the threshold poses no significant adjustment costs for the user aside from the fact that he or she will need to complete more smaller transactions to trade the same amount of cryp-

[35] Chetty et al. 2011; Kleven and Waseem 2013; Chetty 2012.

tocurrency (a relatively small burden).

Finally, there are two major threats to inference using bunching estimation,[36] neither of which pose a problem for this research. First, the presence of another policy that affects transactions at the threshold could confound the estimate; for cryptocurrency, however, there are no other policies that affect transactions at the 1,000 dollar/euro threshold other than due diligence. Second, one may obtain biased estimates if a threshold also serves as a natural reference point. This second concern is unlikely to affect my estimates because I analyze excess bunching in the 100 dollars *below* the due diligence threshold, a range that typically does not serve as a natural reference point for users. Further, while some policies become so well known that the associated threshold can influence individuals' behavior beyond the actual policy itself (e.g., U.S. citizens 65 and older are eligible for Medicare, and this age has become widely associated with retirement), this is not true for cryptocurrency's due diligence threshold.

Estimation Procedure

I first aggregate the number of trades within 10 dollar or euro bins for crypto-to-fiat trading pairs. In order to minimize the potential bias from estimating a counterfactual distribution for a large number of trades, I restrict the analysis to a bandwidth of 250 dollars/euros above and below the due diligence threshold. Because it is not clear *where* below the threshold bunching occurs, I measure excess bunching in the 100 dollars/euros below the threshold.[37] I then fit a third-degree polynomial to the data (excluding the range below the threshold) to estimate the predicted distribution of trades. Next, I calculate the difference between the predicted and actual distribution of trades within the range below the threshold and subtract this from the difference between the predicted and actual values for the rest of the distribution, yielding an estimate of excess bunching in the range below the threshold. I estimate standard errors using parametric bootstrapping by drawing 1,000 samples with replacement from the vector of errors ϵ_j in Equation 1. I calculate a new estimate of excess mass for each sample and derive the standard error as the standard deviation of these coefficients.

Figure 3 illustrates this approach. These graphs show the distribution of trades by 10 dollar increments for two exchanges in the 500 dollars surrounding the threshold. I fit a third-degree polynomial to each sample's data (excluding the range below the threshold) to provide a counterfactual estimate of the distribution. As we can see, the prediction fits the actual data well for Coinsbit (a); Gemini (b), meanwhile, shows a significant spike in the distribution just below the due diligence threshold. This intuition is borne out in the estimates of excess bunching: Coinsbit has an estimate of 0.11 that is not statistically significant, while Gemini has a coefficient of 6.38 with a standard error of 1.96, leading us to reject the null hypothesis of no excess mass in the excluded region. This indicates that the actual trades in the range below the threshold are more than six times greater than predicted and statistically significant ($p < 0.01$).

[36] Kleven 2016.

[37] Since it is inconvenient to conduct many transactions at low amounts (i.e., 10 transactions at $100 each), a savvy user is likely to conduct a transaction relatively close to the threshold (e.g., $950) but still far enough away to avoid transactions that appear very suspicious (e.g., $999 signals likely avoidance).

Figure 3: Excess Bunching by Exchange

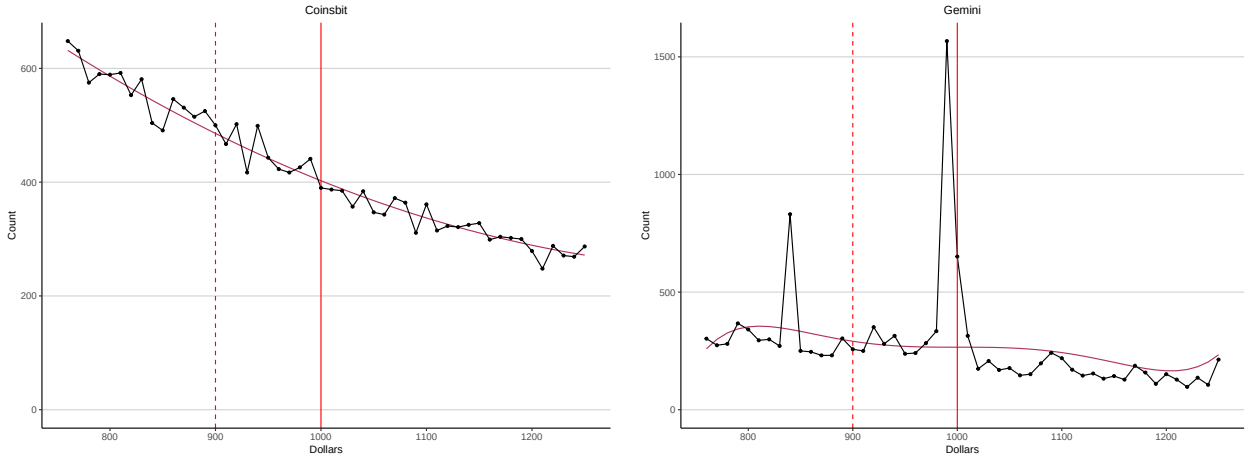

(a) (b)

Caption: Graphs show examples of excess bunching in two exchanges: Coinsbit (a) and Gemini (b).

5.1.1 Results by Trading Pair

Table 1: Excess Bunching by Trading Pair

	Bitcoin			**Ethereum**		
Threshold	*Dollar*	*Euro*	*Yen*	*Dollar*	*Euro*	*Yen*
	(1)	(2)	(3)	(4)	(5)	(6)
700	−1.018	−0.310	5.804	−0.432	−0.567	−0.952
	(0.785)	(0.860)	(6.611)	(0.784)	(0.642)	(4.102)
1,000	5.779***	5.348***	18.135***	2.903**	3.203***	17.871***
	(0.439)	(0.429)	(2.201)	(0.887)	(0.438)	(2.490)
1,300	−0.140	−0.463	12.332	−0.113	−0.713	13.107
	(0.414)	(0.402)	(7.860)	(0.544)	(0.446)	(140.415)
Count Exchanges	11	7	3	11	8	3

Notes: Excess bunching by trading pair between 07/01/20 and 08/03/20; standard errors are in parentheses and stars indicate the statistical significance level: *p<0.05; **p<0.01; ***p<0.001.

First, I estimate the excess mass below the threshold for all exchanges with a given trading pair. Table 1 presents these results; each coefficient can be interpreted as the number of times greater the actual number of trades was relative to the predicted number of trades in the 100 units below the threshold. Each coefficient represents the average excess

bunching across trades for all exchanges offering that particular trade, thus offering a way to compare levels of excess bunching for all trades between cryptocurrency and fiat currency in the dataset. Interestingly, all trading pairs show statistically significant excess mass below the due diligence threshold; however, the level of excess mass varies significantly across these pairs. Notably, Bitcoin and Ethereum trades to Japanese yen show the highest levels of excess mass below the threshold, at 18 and 17 times greater than predicted respectively. Bitcoin to dollar trades show the next highest level of excess mass with nearly 6 times greater trades than predicted, followed by Bitcoin to euro trades (5 times greater). Lastly, Ethereum to dollar and euro trades show nearly 3 times greater excess mass than predicted.

I also include estimates of two placebo thresholds for each trading pair. For each trading pair, I follow the same steps outlined above to estimate excess mass in the 100 units below 700 and 1,300 dollars/euros, which are not subject to due diligence screening and were chosen based on their proximity to the actual threshold. I do not find statistically significant excess mass below these placebo thresholds, confirming our intuitions and highlighting how unusual the mass below the actual threshold is.

5.1.2 Results by Country

Table 2: Bitcoin Trades Excess Bunching by Country

Threshold	*China*	*Japan*	*St. Kitts & Nevis*	*United Kingdom*	*BVI†*	*Estonia*	*United States*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
700	−2.412	3.854	−5.903	−0.051	−3.216	−0.639	−2.102
	(1.452)	(118.545)	(2.907)	(0.950)	(0.927)	(0.677)	(2.051)
1,000	20.600*	16.605***	11.033***	5.394***	1.947**	1.857***	−0.188
	(7.675)	(2.129)	(2.961)	(0.408)	(0.692)	(0.525)	(3.249)
1,300	1.441	−5.210*	4.041	−0.443	0.506	−0.487	−24.257
	(1.154)	(2.162)	(9.874)	(0.446)	(0.370)	(0.352)	(520.638)
Exchanges	1	4	1	3	3	7	3
Pairs	2	4	1	5	3	8	3

Notes: Excess bunching by country between 07/01/20 and 08/03/20; standard errors are in parentheses and stars indicate the statistical significance level: *p<0.05; **p<0.01; ***p<0.001; †British Virgin Islands

Next, I analyze excess mass in trades by country. Similar to my analysis by trading pair, I grouped all the Bitcoin to dollar, euro, or yen trades by the country in which the exchange is located. Accordingly, these estimates can be interpreted as the average excess mass across

all Bitcoin-to-fiat trades for a given country, which shows a "snapshot in time" for the 34 days of trades included in the dataset. I find statistically significant excess mass for all countries except one, the United States. At the extreme, China's two Bitcoin trading pairs show 20 times the predicted number of trades in the range below the threshold, followed by Japan, where exchanges show an average of 16.5 times more trades in the range below the threshold than predicted. In the mid range, St. Kitts and Nevis and the United Kingdom show excess mass 11 and 5 times greater than predicted, while the British Virgin Islands and Estonia show nearly twice the number of trades as predicted. Finally, exchanges in the United States show no statistically significant excess mass in the region below the threshold. For each country, I also include placebo estimates of excess bunching at 700 and 1,300 dollars/euros and find that most of these estimates are not statistically significant.

Table 3: Ethereum Trades Excess Bunching by Country

Threshold	*Japan*	*BVI†*	*Estonia*	*United Kingdom*	*Singapore*	*United States*
	(1)	(2)	(3)	(4)	(5)	(6)
700	2.278	−0.325	0.004	−0.271	0.239	−1.218
	(3.278)	(0.818)	(1.035)	(0.776)	(0.218)	(1.720)
1,000	15.475***	6.446***	5.192***	3.838***	1.336***	0.217
	(2.496)	(1.135)	(1.242)	(0.329)	(0.341)	(1.749)
1,300	2.439**	−1.731*	−2.980	−0.148	−0.077	0.309
	(0.680)	(0.655)	(0.652)	(0.465)	(0.551)	(0.734)
Count Exchanges	3	3	5	3	1	6
Trading Pairs	4	3	7	5	1	6

Notes: Excess bunching by country between 07/01/20 and 08/03/20; standard errors are in parentheses and stars indicate the statistical significance level: *p<0.05; **p<0.01; ***p<0.001; †British Virgin Islands

I also analyze excess mass by country for trades from Ethereum to fiat currency, with the results displayed in Table 3. Similar to Bitcoin trades, I find statistically significant excess bunching below the threshold for all countries but one – the United States. I also find a roughly similar ranking of countries in terms of excess mass, with Japan again on the extreme end showing excess mass 15.5 times greater than predicted in the range below the threshold, followed by the British Virgin Islands, Estonia, and the United Kingdom in the middle range. On the low end, Singapore shows trades below the threshold relatively close to the predicted number (about 134%), while U.S. exchanges do not show statistically significant excess mass in the region below the threshold. These results coupled with those for Bitcoin-to-fiat trades support the conclusion that there is often statistically significant excess mass in the region below the threshold, although this varies significantly across coun-

Figure 4: Ethereum Trades in Extstock Exchange

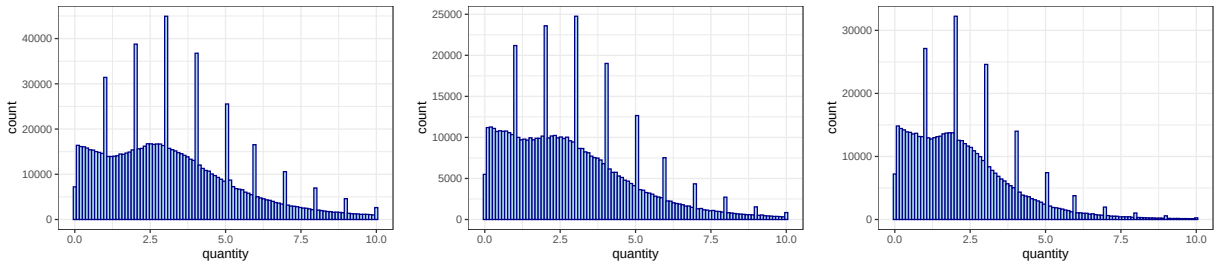

Dollars Euros Bitcoin

(A) Distribution in Ethereum

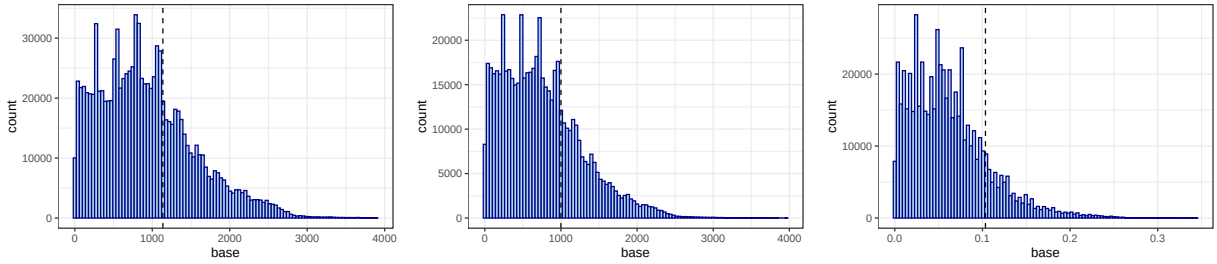

Dollars Euros Bitcoin

(B) Distribution in Euros

Caption: The height of each bin marks the number of transactions at that currency (or cryptocurrency) amount for all trades between 07/01/20 and 08/03/20. The dashed line denotes the 1,000 euro threshold (for trades to euros) or the average 1,000 euro threshold based on the exchange rate between euros and dollars or euros and Bitcoin during the 34 day period.

tries. At the extreme, China and Japan show high levels of excess, the British Virgin Islands, Estonia, and the United Kingdom typically fall in the middle range, and the United States does not show statistically significant excess mass.

5.1.3 Bunching at Round Numbers

One possible alternative explanation for this finding is that bunching at round numbers – rather than effort to avoid due diligence screening – accounts for excess mass below the due diligence threshold. Round number bunching is a common phenomenon and has been documented in a wide range of fields,[38] including high school test scores, baseball batting

[38]Kleven 2016.

averages,[39] house prices,[40] taxable income,[41] and even odometer mileage in used cars.[42] The same holds true for cryptocurrency trades, as we see noticeable bunching at round numbers of Ethereum and round increments of Bitcoin across exchanges in the sample. Importantly, bunching at round numbers does not preclude the use of bunching estimation,[43] and, in fact, bunching at round quantities of cryptocurrency may sometimes be part of efforts to avoid due diligence screening.

Take, for example, patterns of bunching in Extstock, a United Kingdom based exchange that enforces due diligence at 1,000 euros regardless of the currency involved in a trade. Figure 4 (A) shows significant bunching at round numbers of cryptocurrency for trades between Ethereum and dollars, euros, and Bitcoin. Although this bunching might appear random at first, examining the trade distribution in terms of the currency Ethereum is traded into (B) provides interesting results. Noticeably, the distribution of trades from Ethereum to dollars shows significant bunching around 1,115 dollars – just below the average 1,000 euro threshold (marked by the dashed line) – while trades to euros show bunching just below 1,000 euros; thus, Ethereum trades to both currencies show bunching below the due diligence threshold. Trades from Ethereum to Bitcoin, meanwhile, show no significant bunching below 1,000 euros, which is unsurprising given that due diligence *is not* enforced for trades that only involve cryptocurrency. These findings suggest that although bunching at round numbers is a naturally occurring social phenomenon, there appears to be something unusual happening in the crypto-to-fiat trades that drive how users select their transaction amounts.

This example from Extstock highlights several common features that make purely "coincidental" bunching at round numbers of cryptocurrency below the threshold unlikely. First, it is unlikely that users simply prefer conducting transactions at round numbers of cryptocurrency that happen to fall below the threshold as we find evidence of excess bunching below the threshold for both Bitcoin and Ethereum trades. Second, I do not find consistent evidence of excess mass below the placebo thresholds for trades from either cryptocurrency, highlighting that the mass below the due diligence threshold is, indeed, unusual. Third, it is unlikely that people simply prefer conducting transactions that result in certain fiat currency amounts that happen to lie below the threshold as the total dollar and euro amounts with the highest bunching varies across exchanges and over time (e.g., the total dollar value of a 0.1 Bitcoin trade varies due to the Bitcoin-to-dollar exchange rate). Fourth and perhaps most convincingly, I find that exchanges that enforce the threshold in euros tend to have bunching below 1,000 euros, while exchanges that enforce it in dollars tend to have bunching below 1,000 dollars. These points can lead us to conclude that an explanation based on coincidental bunching at round numbers of cryptocurrency does not adequately explain the data's consistent pattern of excess mass below the due diligence threshold.

[39] Pope and Simonsohn 2011.
[40] Pope, Pope, and Sydnor 2015.
[41] Kleven and Waseem 2013.
[42] Lacetera, Pope, and Sydnor 2012.
[43] Kleven and Waseem (2013) suggest estimating a model using the full sample and again with fixed effects for round numbers that show significant bunching.

Importantly, I do not argue that all bunching at round numbers of cryptocurrency stems from individuals' efforts to avoid due diligence screening. Rather, I suggest that bunching at round numbers is a natural feature of cryptocurrency transactions, and a small number of individuals may exploit this feature in their efforts to avoid screening. I also do not rule out the possibility that some of the bunching at round numbers of cryptocurrency below the threshold may be driven by amplification – users may be more likely to conduct transactions at 0.1 Bitcoin if they know that many other transactions are conducted at this amount. However, in order for amplification to occur, there must be an initial trend driving user behavior, which I argue may come from users seeking to avoid due diligence.

5.2 Instrumental Variable Estimation

Figure 5: Variation in Price and the Proportion of Trades in Coinbase

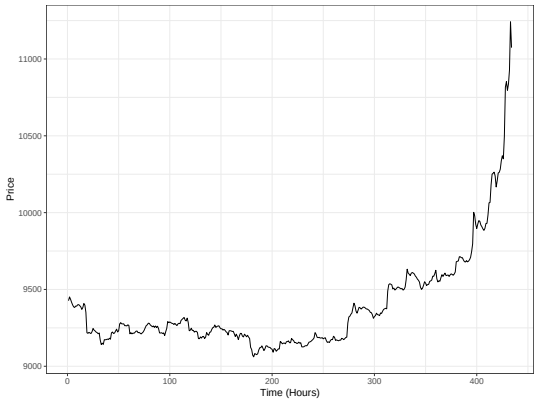

(a)

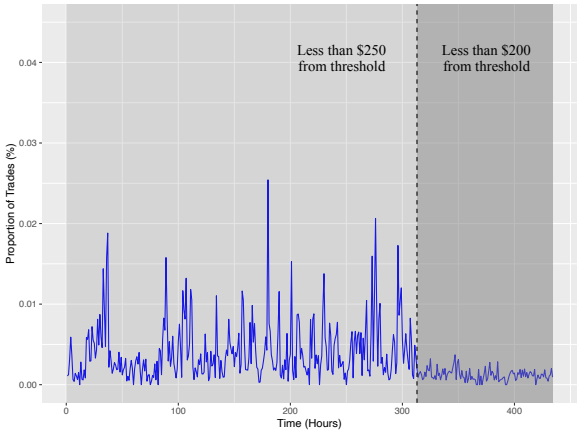

0.085 Bitcoin

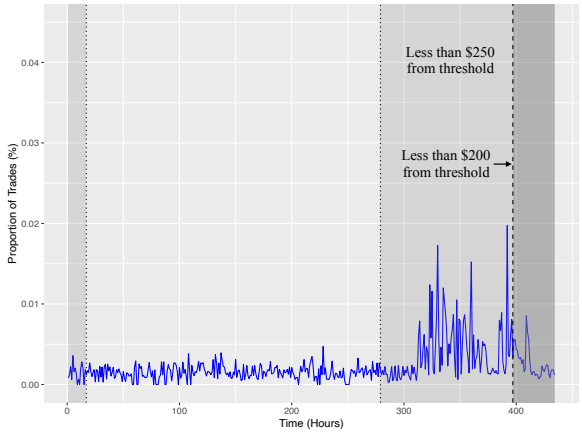

0.080 Bitcoin

Caption: These graphs show that the proportion of hourly trades at 0.085 Bitcoin decreases as the crypto-to-fiat exchange rate drives the total dollar amount of these transactions closer to the threshold. This decrease coincides with an uptick in the proportion of trades at 0.080 Bitcoin (the next major round quantity of Bitcoin). This behavior may result from users switching to conducting transactions at a lower round number of cryptocurrency in response to changes in the threshold.

Second, I perform instrumental variable regression to analyze how users adjust the price of their transactions in response to changes in the threshold. Specifically, I look at trades in exchanges that enforce the due diligence threshold in another currency than the one involved

in the trade.[44] This allows me to isolate how users respond to changes in the threshold that result from exogenous fluctuations in the fiat-to-fiat exchange rate while controlling for the crypto-to-fiat exchange rate. I predict that once the due diligence threshold passes below a given bin in the distribution, the proportion of trades in that bin will decrease as users switch to conducting transactions at smaller amounts. Indeed, I find a negative and often statistically significant coefficient for bins above the threshold.

The Setup

First, I assume that there is a distribution of cryptocurrency quantities held by individuals that they may convert to fiat currency using an exchange. We can write the total fiat currency amount of one such trade as $D_i = p_t \cdot q$, where q is the quantity of cryptocurrency and p is the crypto-to-fiat exchange rate at time t. The crypto-to-fiat exchange rate is an important factor that influences the proportion of trades at a given quantity, since a price increase between two time periods ($p_{t+1} > p_t$) will *increase* the the total fiat value of the trade (D_i) for a given quantity of cryptocurrency. However, since the crypto-to-fiat exchange rate also decreases the amount of cryptocurrency a new user can buy, we should expect a roughly net zero effect of the crypto-to-fiat exchange rate on the proportion of trades at a given amount over time.[45]

Another factor that might influence the proportion of trades is the behavior of individuals seeking to avoid due diligence. Specifically, a change in the crypto-to-fiat exchange rate may drive the total fiat amount of a trade above 1,000 dollars/euros for some cryptocurrency quantities. Accordingly, users who wish to keep the total fiat value of the transaction below the due diligence threshold (K) may adjust the quantity of a cryptocurrency transaction by some value b, so that $D_i = (q - b) * p < K$. However, movement of the threshold and the proportion of trades at a given amount of cryptocurrency could have common causes or be driven by reverse causality. For example, an increase in demand for cryptocurrency could drive the proportion of trades in a bin along with the exchange rate (and the due diligence threshold) higher. Thus, I seek to isolate *exogenous* variation in the threshold in order to measure individuals' response by using fluctuations in fiat-to-fiat exchange rates for trades in exchanges that enforce the due diligence threshold in a different currency (e.g., Bitcoin-to-dollar trades in an exchange that enforces due diligence at 1,000 euros). If users respond to changes in the threshold, we should expect the proportion of trades for a given amount of cryptocurrency to decrease once the threshold passes below that bin.

This strategy rests on the assumption that the exchange rate between two fiat currencies affects the proportion of trades in a given bin *only through* manipulation of the due diligence threshold; this assumption would be violated, for example, if individuals sought to perform a crypto-to-fiat trade and then convert that currency into the currency used to enforce due diligence. However, there is no evidence and little reason to suspect that users engage in these sorts of trades. A second main assumption of this strategy is that fluc-

[44] Recall that exchanges in Europe enforce the due diligence threshold at €1,000 while exchanges everywhere else enforce due diligence at $1,000.

[45] This depends on individuals buying and selling roughly the same amount of cryptocurrency over time.

tuations in the exchange rate between two currencies are exogenous. This is a reasonable assumption because if people could predict variation in fiat-to-fiat exchange rates, then they would be able to manipulate exchange markets (something for which we do not see evidence).

Estimation Procedure
To isolate the effect of passing below the threshold, I restrict the analysis to a bandwidth of 250 dollars/euros above and below the threshold and aggregate the number of trades across quantities of one-hundredth of a Bitcoin and one-tenth of a unit of Ethereum. I then calculate the proportion of trades in each bin j for each hour t relative to the total number of trades in that hour (Y). This specification of Y as a proportion allows me to control for variation in hourly trade volumes within exchanges as well as variation in trade volumes across exchanges. I specify the dependent variable (X) as a dummy indicating whether the average fiat value of a bin is equal to or above the due diligence threshold in a given hour, resulting in the model

$$Y_{jt} = \beta_0 + \beta_1 X_{jt} + \beta_2 Price_t + \epsilon_{jt}, \tag{2}$$

where $Price$ is a control for the crypto-to-fiat exchange rate in hour t.

To isolate the exogenous variation of the threshold independent of the proportion of trades in a given bin, I estimate a first-stage model using the fiat-to-fiat exchange rate (Z) as an instrument. Because I wish to estimate how the threshold passing below the bin affects the proportion of trades at a given amount rather than an average effect across all bins, I include fixed effects for each quantity of cryptocurrency (Q_j). Thus, the first-stage model is

$$X_{it} = \beta_0 + \beta_1 Z_t + \beta_2 Q_j + \epsilon_{it}. \tag{3}$$

I then substitute the predicted values from Equation 3 for X in Equation 2.

5.2.1 Results

Table 4 presents the main results of the instrumental variable regression with exchanges aggregated by country. These results show that there is a statistically significant *decrease* in the proportion of trades in a bin above the threshold, although this varies in magnitude across countries. At the extreme, China's exchange shows an average 1.8% decrease in the proportion of hourly trades in a bin once the due diligence threshold passes below it. Other countries show a smaller but still significant decreases in the proportion of hourly trades, including a 0.4% decrease for Japan, 0.3% for the United Kingdom, 0.2% for the United States, and 0.04% for Estonia. Notably, the threshold coefficients are larger in absolute terms than coefficients for the crypto-to-fiat exchange rate ($Price$), suggesting that the threshold has a bigger influence than price on the proportion of trades at a given amount when controlling for both factors. The country rankings of these results also closely mirror those obtained using bunching: China and Japan are at the far end followed by the United Kingdom, with the United States and Estonia showing smaller absolute coefficients for the decrease of trades in bins above the threshold. All models have F-statistics well above 10, indicating that the instruments are strong; I also include fixed effects by exchange for countries with more than

Table 4: Proportion of Trades in Bin by Country

	China	Japan	United Kingdom	United States	Estonia
	(1)	(2)	(3)	(4)	(5)
Threshold	-0.018^{***}	-0.004^{***}	-0.003^{***}	-0.002^{***}	-0.0004^{***}
	(0.001)	(0.0001)	(0.0001)	(0.0001)	(0.000)
Price	0.000^{**}	-0.000^{***}	0.000^{***}	-0.000^{***}	-0.000^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Constant	-0.006	0.023^{***}	-0.003	0.025^{***}	0.004^{***}
	(0.017)	(0.002)	(0.0005)	(0.001)	(0.0004)
Observations	2,874	42,088	62,717	6,069	41,164
Site FEs	N	Y	Y	N	Y
Number of Exchanges	1	4	3	1	5
R^2	0.03	0.07	0.10	0.09	0.17
Residual Error	0.02	0.01	0.007	0.005	0.002
F Statistic	190.5	138.6	91.9	66.6	214.9

Notes: The unit of observation is the hour-bin. For countries with more than one exchange, I include exchange fixed effects. Standard errors are in parentheses and stars indicate the statistical significance level: *p<0.05; **p<0.01; ***p<0.001.

Figure 6: Coefficient Plot by Threshold and Country

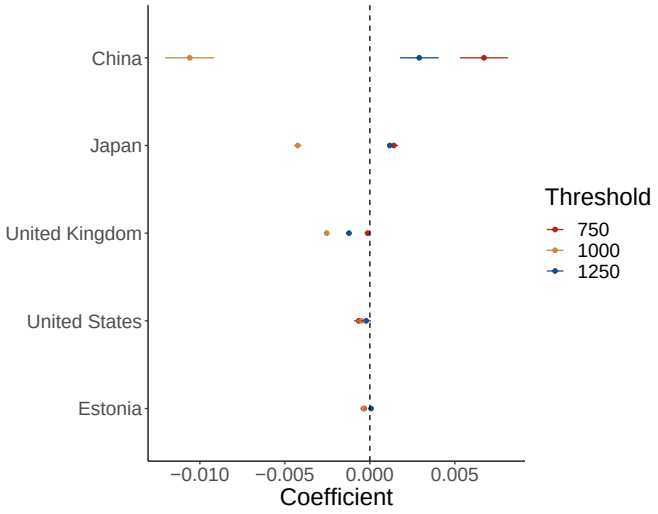

one exchange included in the estimate.

Similar to the bunching analysis, I include estimates for two placebo thresholds – 750 and 1,250 dollars or euros. Figure 6 shows that the due 1,000 euro/dollar threshold is consistently much larger in magnitude and negatively signed when compared to the placebo thresholds. The spread between the actual threshold and the placebo thresholds also follows the same ranking as the magnitude of the coefficients. These results suggest that behavior close to the due diligence threshold is significantly different than at other amounts, with significant divergence between the actual and placebo thresholds in Chinese and Japanese exchanges. For the United States and Estonia, meanwhile, coefficients are grouped close to zero, suggesting that there is little change in the number of transactions for a bin above the threshold.

6 Discussion

My findings support two micro-level conclusions about how actors respond to customer due diligence regulation and two macro-level conclusions related to the dynamics of international cooperation in anti-money laundering enforcement.

Individual Response

First and foremost, these results suggest that a small minority of users in some exchanges avoid due diligence screening by adjusting the quantity of their transactions to fall below the due diligence threshold. Both bunching estimation and instrumental variable estimation show unusual activity below the due diligence threshold, suggesting that individuals strategically respond to regulation in an effort to avoid government screening. This is a

notable finding since most of the anti-money laundering literature does not treat individuals as strategic actors, although some money laundering cases show evidence of this. These results suggest that anti-money laundering laws should carefully consider strategic behavior by individuals and a regulation's potential unintended consequences.

Exchange Response

These results suggest that many exchanges selectively enforce regulations — complying with easily verifiable measures (due diligence at the threshold) while neglecting others (risk-based measures). I base this assessment on a first-order conclusion – the presence of excess mass below the threshold indicates that exchanges are enforcing customer due diligence at the threshold (individuals have an incentive to conduct transactions below the threshold). I also draw the second-order conclusion that the persistence of these spikes over time suggests that exchanges are not performing additional risk-based measures that would lead them to conduct additional screening and detect suspicious trends. Indeed, this interpretation fits with findings by Findley, Nielson, and Sharman that corporate service providers (2014) and banks (2020) show little sensitivity to a customer's potential money laundering risk when opening accounts, which suggests that there are significant weaknesses in intermediaries' application of risk-based measures. Importantly, although one could interpret the *absence* of unusual activity near the threshold in some countries as evidence that exchanges are not enforcing customer due diligence at the threshold (presenting no incentive to sort below it), I argue this is unlikely given that customer due diligence records can be verified by national regulators.

These findings suggest that exchanges respond strategically to regulation with an eye toward minimizing costs. This conceptualization differs from common ones in the anti-money laundering literature, which typically either view intermediaries as uniformly in compliance with national laws or diverging into categories of good and bad intermediaries. This research suggests that scholars should consider the role of intermediaries as cost-minimizing, strategic actors.

Variation Across Countries

These results also suggest that there is significant variation in how well countries combat money laundering risk despite having similar laws in place. These differences vary in orders of magnitude; for example, China's excess mass below the due diligence threshold is 20 times the predicted distribution while the U.S. shows no significant excess mass. Similarly, the instrumental variable estimation shows that Japanese exchanges had an average decrease of 0.4% of the hourly trades for a quantity of cryptocurrency above the threshold, while U.S. exchanges showed a decrease of half that (0.2%) and Estonian exchanges showed one-tenth the decrease (0.04%). These cross-national differences are important, particularly since many in the international community and some scholars have used national laws as a measure of countries' efforts to combat money laundering. And although the FATF has noted significant deficiencies in the national enforcement of anti-money laundering laws, it has used national laws to measure compliance for its most effective form of punishment – the blacklist.

Developed Countries

Finally, my results suggest that – contrary to widely held views in some academic and policy circles[46] – enforcing due diligence laws poses a significant challenge for *developed* countries. My analysis shows significant levels of suspicious activity across countries, including developed countries like Japan and the United Kingdom. Indeed, this finding fits with a wealth of evidence from recent events that show significant failures in the enforcement of customer due diligence laws by wealthy, industrialized countries. These findings suggest that intermediaries failure to enforce customer due diligence laws (including in developed countries) is a major weakness in the global anti-money laundering regime, above and beyond (I argue) the failure of some developing countries to implement comprehensive anti-money laundering legal frameworks.

7 Conclusion

This research aims to contribute to the anti-money laundering literature by testing questions about how actors respond to new regulation. I collect a new dataset of cryptocurrency transactions that provides ground-level data about how intermediaries enforce a new customer due diligence law. I find that individuals respond strategically to the new regulation, and that many intermediaries seem to selectively enforce customer due diligence laws.

These findings hold important implications for cooperation in anti-money laundering enforcement. Specifically, they suggest that the FATF system (and much of the academic literature) has done a poor job of identifying defection because it has focused on national laws, thus overlooking a major problem – the failure of intermediaries to enforce customer due diligence laws. This, in turn, calls into question the fairness of FATF blacklisting, which has caused severe economic harm to listed countries. Indeed, some have raised allegations of bias, arguing that the FATF has blacklisted countries based on political factors rather than an accurate measure of states' efforts to address money laundering.[47] However, there are signs that the FATF has begun to recognize this problem. Most tellingly, the FATF has introduced a new measure of effectiveness – defined as "the extent to which the national [anti-money laundering] system is achieving the objectives of the FATF standards" – (based on qualitative assessments) to its most recent round of evaluation reports.[48] And while this represents a step in the right direction, the FATF has yet to use an enforcement-based assessment for blacklisting and continues using national laws instead.

Finally, my findings highlight how anti-money laundering enforcement differs structurally from many other areas of international cooperation. Although state actors play a primary role in many other areas of international cooperation like trade, human rights, and security, anti-money laundering enforcement involves not only state actors but also sub-state actors – namely individuals (who are screened by customer due diligence laws) and intermediaries

[46] The IMF, for example, suggests that countries with "extensive financial infrastructure" provide the most effective anti-money laundering enforcement.

[47] Sharman 2010.

[48] FATF 2019, p. 15.

(the private sector actors who carry them out). Accordingly, my findings calls for greater theorizing and empirical testing of anti-money laundering laws at the level of sub-state actors.

References

Aguilar, Diana (2019). "Regulators Debate Cryptocurrency Legislation Ahead of G20 Summit - CoinDesk". In: *CoinDesk*. URL: https://www.coindesk.com/regulators-begin-to-debate-cryptocurrency-legislation-ahead-of-g20-summit.

All Cryptocurrencies (2019). [Online; accessed 10. Dec. 2019]. URL: https://coinmarketcap.com/all/views/all.

Apuzzo, Matt (2014). "Secret Global Strike Kills 2 Malicious Web Viruses". In: *N.Y. Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/2014/06/03/world/europe/battling-destructive-computer-viruses-agents-seize-networks-used-by-hackers.html.

Bastani, Spencer and Håkan Selin (2014). "Bunching and non-bunching at kink points of the Swedish tax schedule". In: *Journal of Public Economics* 109, pp. 36–49.

Baydakova, Anna (2020). "$2 Billion Lost in Mt. Gox Bitcoin Hack Can Be Recovered, Lawyer Claims - CoinDesk". In: *CoinDesk*. URL: https://www.coindesk.com/2-billion-lost-in-mt-gox-bitcoin-hack-can-be-recovered-lawyer-claims.

Broadwater, Luke (2019). "Baltimore transfers $6 million to pay for ransomware attack; city considers insurance against hacks". In: *baltimoresun.com*. URL: https://www.baltimoresun.com/politics/bs-md-ci-ransomware-expenses-20190828-njgznd7dsfaxbbaglnvnb
story.html.

Chainalysis Team (2018). *The Changing Nature of Cryptocrime*. [Online; accessed 21. Nov. 2019]. URL: https://blog.chainalysis.com/reports/report-the-changing-nature-of-cryptocrime.

Chetty, Raj (2012). "Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply". In: *Econometrica* 80.3, pp. 969–1018.

Chetty, Raj et al. (2011). "Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records". In: *The quarterly journal of economics* 126.2, pp. 749–804.

Chokshi, Niraj (2019). "Hackers Are Holding Baltimore Hostage: How They Struck and What's Next". In: *N.Y. Times*. ISSN: 0362-4331. URL: https://www.nytimes.com/2019/05/22/us/baltimore-ransomware.html.

CipherTrace Cryptocurrency Intelligence (2018). *Cryptocurrency Anti-Money Laundering Report 2018 Q3*. [Online; accessed 21. Nov. 2019]. URL: https://ciphertrace.com/wp-content/uploads/2019/01/crypto_aml_report_2018q3.pdf.

Cox, Joseph (2015). "The Biggest Dark Web Markets Rake in Up to $500,000 a Day, Study Says". In: *Vice*. [Online; accessed 21. Nov. 2019]. URL: https://www.vice.com/en_us/article/kbzme9/the-biggest-dark-web-markets-rake-in-up-to-500000-a-day-study-says.

CyberEdge (May 2019). *2019 Cyberthreat Defense Report*. [Online; accessed 21. Nov. 2019]. URL: https://cyber-edge.com/wp-content/uploads/2019/03/CyberEdge-2019-CDR-Report.pdf.

Deleanu, Ioana Sorina (2017). "Do countries consistently engage in misinforming the international community about their efforts to combat money laundering? Evidence using Benford's law". In: *PloS one* 12.1, e0169632.

Engvall, Johan (2006). "The state under Siege: The drug trade and organised crime in Tajikistan". In: *Europe-Asia Studies* 58.6, pp. 827–854.

FATF (2019). *Methodology for Assessing Technical Compliance with the FATF Recommendations and the Effectiveness of AML/CFT Systems*. URL: http://www.fatf-gafi.org/media/fatf/documents/methodology/fatf%20methodology%2022%20feb%202013.pdf.

Findley, Michael G, Daniel L Nielson, and Jason Campbell Sharman (2014). *Global shell games: Experiments in transnational relations, crime, and terrorism*. 128. Cambridge University Press.

Foley, Sean, Jonathan R Karlsen, and Tālis J Putniņš (2019). "Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies?" In: *The Review of Financial Studies* 32.5, pp. 1798–1853.

Garrett, Geoffrey (1992). "International cooperation and institutional choice: the European Community's internal market". In: *International Organization* 46.2, pp. 533–560.

Gnutzmann, Hinnerk, Killian J McCarthy, and Brigitte Unger (2010). "Dancing with the devil: Country size and the incentive to tolerate money laundering". In: *International Review of Law and Economics* 30.3, pp. 244–252.

Greenberg, Andy (2014). "What Is the Dark Web?" In: *WIRED*. URL: https://www.wired.com/2014/11/hacker-lexicon-whats-dark-web.

— (2015). "Crackdowns Haven't Stopped the Dark Web's $100M Yearly Drug Sales". In: *WIRED*. URL: https://www.wired.com/2015/08/crackdowns-havent-stopped-dark-webs-100m-yearly-drug-sales.

— (2017). "Monero, the Drug Dealer's Cryptocurrency of Choice, Is on Fire". In: *WIRED*. URL: https://www.wired.com/2017/01/monero-drug-dealers-cryptocurrency-choice-fire.

Havilland, Paul de (2019). "FATF Issues Draconian Crypto Recommendations: You Now Have 12 Months To Comply | Crypto Briefing". In: *Crypto Briefing*. [Online; accessed 24. Nov. 2019]. URL: https://cryptobriefing.com/fatf-draconian-crypto.

Kauflin, Jeff (2019). "The Anatomy Of A Fake Cryptocurrency Trade: How Exchanges Create Phony Transactions". In: *Forbes*. URL: https://www.forbes.com/sites/jeffkauflin/2019/07/02/the-anatomy-of-a-fake-cryptocurrency-trade-how-exchanges-create-phony-transactions/#309ccbc83132.

Kleven, Henrik J and Mazhar Waseem (2013). "Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan". In: *The Quarterly Journal of Economics* 128.2, pp. 669–723.

Kleven, Henrik Jacobsen (2016). "Bunching". In: *Annual Review of Economics* 8, pp. 435–464.

Lacetera, Nicola, Devin G Pope, and Justin R Sydnor (2012). "Heuristic thinking and limited attention in the car market". In: *American Economic Review* 102.5, pp. 2206–36.

Levitsky, Melvyn (2003). "Transnational criminal networks and international security". In: *Syracuse J. Int'l L. & Com.* 30, p. 227.

Masciandaro, Donato, Elod Takats, and Brigitte Unger (2007). *Black finance: the economics of money laundering*. Edward Elgar Publishing.

McMillan, Robert (2018). "The Inside Story of Mt. Gox, Bitcoin's $460 Million Disaster". In: *Wired*. URL: https://www.wired.com/2014/03/bitcoin-exchange.

Meiklejohn, Sarah et al. (2013). "A fistful of bitcoins: characterizing payments among men with no names". In: *Proceedings of the 2013 conference on Internet measurement conference*. ACM, pp. 127–140.

Mnuchin, Steven (2019). *Remarks of Secretary Steven T. Mnuchin FATF Plenary Session Orlando, Florida | U.S. Department of the Treasury*. [Online; accessed 24. Nov. 2019]. URL: https://home.treasury.gov/news/press-releases/sm713.

Morse, Julia (2019). "Blacklists, market enforcement, and the global regime to combat terrorist financing". In: *International Organization, Forthcoming*.

Ott, Thomas (2019). *Testimony of Thomas P. Ott, Associate Director, Enforcement Division, before the House Committee on Financial Services | FinCEN.gov*. [Online; accessed 21. Nov. 2019]. URL: https://www.fincen.gov/index.php/news/testimony/testimony-thomas-p-ott-associate-director-enforcement-division-house-committee.

Pope, Devin and Uri Simonsohn (2011). "Round numbers as goals: Evidence from baseball, SAT takers, and the lab". In: *Psychological science* 22.1, pp. 71–79.

Pope, Devin G, Jaren C Pope, and Justin R Sydnor (2015). "Focal points and bargaining in housing markets". In: *Games and Economic Behavior* 93, pp. 89–107.

Reuter, Peter (2005). *Chasing dirty money: The fight against money laundering*. Peterson Institute.

— (2013). "Are estimates of the volume of money laundering either feasible or useful?" In: *Research handbook on money laundering*. Edward Elgar Publishing.

Saez, Emmanuel (2010). "Do taxpayers bunch at kink points?" In: *American economic Journal: economic policy* 2.3, pp. 180–212.

Schwarz, Peter (2011). "Money launderers and tax havens: Two sides of the same coin?" In: *International Review of Law and Economics* 31.1, pp. 37–47.

Sharman, Jason C (2010). "Shopping for anonymous shell companies: An audit study of anonymity and crime in the international financial system". In: *Journal of Economic Perspectives* 24.4, pp. 127–40.

Soska, Kyle and Nicolas Christin (2015). "Measuring the longitudinal evolution of the online anonymous marketplace ecosystem". In: *24th {USENIX} Security Symposium ({USENIX} Security 15)*, pp. 33–48.

Swanstrom, Niklas (2007). "The narcotics trade: a threat to security? National and transnational implications". In: *Global Crime* 8.1, pp. 1–25.

Takáts, Előd (2011). "A theory of "Crying Wolf": The economics of money laundering enforcement". In: *The Journal of Law, Economics, & Organization* 27.1, pp. 32–78.

The Anatomy Of A Fake Cryptocurrency Trade: How Exchanges Create Phony Transactions (2019). [Online; accessed 8. Sep. 2020]. URL: https://www.forbesafrica.com/technology/2019/07/03/the-anatomy-of-a-fake-cryptocurrency-trade-how-exchanges-create-phony-transactions.

Unger, Brigitte and Joras Ferwerda (2008). "Regulating money laundering and tax havens: The role of blacklisting". In: *Discussion Paper Series/Tjalling C. Koopmans Research Institute* 8.12.

United Nations (2019). *World Drug Report 2019*. [Online; accessed 21. Nov. 2019]. URL: https://wdr.unodc.org/wdr2019/prelaunch/WDR19_Booklet_2_DRUG_DEMAND.pdf.

Williams, Phil (1994). "Transnational criminal organisations and international security". In: *Survival* 36.1, pp. 96–113.