Illegal Drug Use and Government Policy: Evidence from a Darknet Marketplace*

Priyanka Goonetilleke
University of Pennsylvania

Anastasia Karpova
Princeton University

Artem Kuriksha†
University of Pennsylvania

Peter Meylakhs
Sechenov University

†**Job Market Paper.** Click here for latest version.

April 26, 2024

Abstract

This paper develops a structural model of demand for illegal drug varieties and studies how consumers substitute between different types of drugs in response to government policies. We use a unique longitudinal dataset on prices, quantities, and individual decisions that we obtained by scraping a darknet marketplace that covered the majority of the retail illegal drug trade in Russia. Our estimation procedure exploits a novel set of micro-level moment conditions to identify correlations in preferences for specific drug types and the degree of attachment to them. We find that the median own-price elasticity of demand for illegal drugs is -3.6, and there is high substitution within two classes of drugs: medium-risk stimulants and cannabis. We validate our estimates using exogenous variation in the price of hashish caused by increased policing. The estimated model is used to evaluate counterfactual drug policies. We find that the legalization of cannabis has the benefit of decreasing the use of riskier drugs while increasing cannabis use. For every 4 additional doses of cannabis consumed, 1 less dose of another drug is consumed. Our estimates show that the recent introduction of a new family of synthetic drugs has increased total drug demand in the country by 40%, suggesting that governments should allocate resources to prevent the introduction of new drug products. Finally, our model helps identify the optimal drugs to target for interdiction, specifically those without close substitutes, such as α-PVP.

*We are indebted to Juan Camilo Castillo, Andrew Shephard, and Petra Todd for their continuous support of this project. We are very grateful to Juan Pablo Atal, Hanming Fang, and Jesus Fernandez-Villaverde as well as Diane Alexander, Abby Alpert, Edvard Bakhitov, Janet Currie, Ulrich Doraszelski, Renata Gaineddenova, Jeff Gortmaker, Kate Ho, Steven Tate, Eugeny Yakovlev, Cung Truong Hoang, Kathleen Hui, Margarita Khvan, and Anya Schetkina for helpful discussions. We thank Ekaterina Aleksandrova, Nicolas Christin, and Vitovt Kopytok for their help with the datasets we use. We thank George Zhang for research assistance. We thank Mikhail Golichenko, Maxim Gorbunov, Andrey Kaganskikh, Anastasia Kuzina, Aleksei Lakhov, Inessa Romanova, Ivan Zhavoronkov, the journalists of Proekt Media, and the anonymous interviewees for helping us understand the context of the drug trade in Russia. We thank the Graduate Student Government of the School of Arts & Sciences at the University of Pennsylvania for financial support. All mistakes are our own.

1. Introduction

The illegal drug trade is a global problem with implications for public health, property crime, violence, unemployment, and incarceration rates.[1] The U.S. government spent more than $40 billion in 2022 in an attempt to address this problem (National Drug Control Budget, 2023). Approximately half of these resources are dedicated to restricting supply, with the rationale that such measures decrease drug availability, raise prices, and consequently reduce drug use. However, the merits of supply-side enforcement are debatable (New York Times, 2023; The Economist, 2023). In many jurisdictions, including Canada, Thailand, and several U.S. states (such as Arizona, Illinois, and New York), policymakers are implementing the radically different policy of legalization. This has so far been enacted in relation to cannabis, a popular class of drugs that are thought to be relatively safer than other drug types.

Evaluating the effects of such policies on drug use requires an understanding of consumers' demand for various illegal drugs and, in particular, how they substitute between different drug types. Substitution is likely to decrease the efficiency of interventions targeting specific drugs, such as seizures or crop eradication. Although the use of the targeted drug decreases, these actions may also increase the use of other drug types that serve as substitutes. Conversely, substitution can yield beneficial consequences from the legalization of low-risk drugs if it results in a reduction in the use of more dangerous drugs.

This paper examines demand for illegal drugs and the impact of drug policies on drug use. We are able to study demand for a wide range of illegal drugs due to unique, high-quality panel data covering the drug market in Russia. We begin by documenting heterogeneity in preferences for illegal drugs. Next, we develop and estimate a demand model that can account for the observed patterns. Since our model allows for consumer heterogeneity, it can generate realistic predictions regarding the substitution between a wide variety of illegal drugs. We then apply the estimated model to investigate the impact of different counterfactual drug policies.

To estimate demand for drugs, we must address the significant challenge that the market for illegal drugs is usually not observed. We make use of data derived from a darknet

[1] According to the Centers for Disease Control and Prevention (CDC), drug overdose-related mortality in the U.S. has been steadily increasing and surpassed 100,000 deaths in the past two years (CDC, NCHS, 2022, 2023). Immense losses of human lives are not the only consequence of illegal drug use: drugs are associated with property crime, violence, and unemployment (Fryer et al., 2013). The "war on drugs" imposes a heavy burden on society: in 2020, almost 200,000 prisoners in the U.S. were sentenced for drug-related offenses (Carson, 2021). Harwood and Bouchery (2004) estimated that the total cost of drug abuse in the U.S. exceeds $200 billion per year. Moreover, drugs present a global issue, with the United Nations (UN) reporting that approximately 284 million people worldwide used drugs in 2020 (UN Office on Drugs and Crime, 2022).

marketplace known as Hydra, which operated from 2015 to 2022. Hydra was the largest darknet marketplace in the world and, in the later years of its existence, covered the majority of the retail drug trade in major Russian cities. Thus, Hydra presents a unique opportunity to observe the market for illegal drugs and learn about drug users' preferences. We compiled a novel micro-level panel dataset by regularly scraping data from the marketplace for over a year. A crucial advantage of our dataset is that it enables us to estimate the quantities and prices of drugs sold in each location where the market operated.

We combine data on drug listings with an individual-level panel dataset of marketplace user reviews, which allows us to infer individual consumption patterns. We study the intertemporal correlation between the drugs reviewed by particular consumers. Our findings indicate that consumers typically exhibit attachment to a specific drug, most often choosing the same drug over multiple time periods. However, the average degree of attachment varies among different drug types. For instance, consumers of cocaine tend to display a stronger intertemporal attachment than MDMA users. Furthermore, our analysis reveals that consumers may also demonstrate preferences for groups of drugs. In particular, individuals who have purchased amphetamine, MDMA, or mephedrone — three stimulants known to have similar effects — are substantially more likely to purchase these three drug types in other time periods. Consequently, we can anticipate higher substitution within this group of drugs, which holds significant implications for the effects of drug policies. For example, amphetamine-focused drug enforcement is expected to reduce amphetamine demand but also increase the demand for MDMA and mephedrone. In many cases, the substitution patterns suggested by our data do not correspond to what one would expect *ex-ante* from the basic characteristics of drugs available in the medical literature. For instance, we observe minimal substitution between mephedrone and $\alpha-$PVP, despite both being classified as "bath salts."

We develop a model that can capture these patterns, building upon the mixed logit framework (also known as BLP, see Berry, 1994; Berry et al., 1995). The model accounts for consumer heterogeneity, which is critical in our setting for making accurate predictions regarding substitution patterns. Because the observable characteristics have little power to explain the consumption patterns we observe, we allow for heterogeneity in preferences by introducing random coefficients for dummies for a set of the most popular drug types. These coefficients are consumer-specific and describe idiosyncratic attachment to these drug types.

To identify substitution patterns, we exploit a novel set of moment conditions derived from our micro-level data on consumer reviews. These moments capture how the drug types chosen by the same user are correlated over time. We develop a simulation procedure that allows us to utilize these moments even when information on purchases is partially missing, which occurs because not all orders are reviewed or only a subset of reviews is available.

Thus, we solve the usual challenge of identifying the covariance of random coefficients when estimating BLP-type models. In our case, identification is particularly demanding because it requires a large number of drug-specific price instruments. Our moments effectively identify covariances between random coefficients and facilitate estimation in a manner akin to the use of second-choice data (Berry et al., 2004). Our method may be applicable in other settings where demand is estimated using data from an online marketplace when second-choice data is unavailable, as reviews can often be scraped at a low cost.

Our estimates reveal a significant degree of price sensitivity among drug users, with a median price elasticity for drug products of -3.6. We have identified four drugs characterized by a relatively high degree of substitution: amphetamine, hashish, marijuana, and MDMA. Importantly, we find substantial heterogeneity in substitution patterns. For instance, diversion ratios indicate that there is five times more substitution from amphetamine to MDMA than to α-PVP, although the latter two drugs have roughly equal market shares. This shows that a model without consumer heterogeneity would yield inaccurate predictions of consumer responses to changes in drug prices.

We are able to validate our estimates by exploiting an exogenous supply-side shock that occurred during the period when we scraped Hydra. In the summer of 2019, the availability of hashish dramatically decreased due to a series of overseas operations targeting the trafficking of this drug. Our estimated model closely predicts the observed response of consumption to increased prices.

We then employ our model to evaluate the outcomes of several counterfactual supply-side drug policies. First, we investigate the impact of cannabis legalization on the consumption of other drugs. Substitution can serve as a significant rationale for legalization if it leads to a reduction in the use of more dangerous drugs. We assume that legalization induces the same substitution patterns as a reduction in the price of cannabis.[2] Our findings indicate that legalization is accompanied by a significant increase in cannabis consumption. For instance, the model predicts that if the price of cannabis decreases by 50%, cannabis use will increase by 320%. However, the use of other drugs will decrease by 14%, which suggests that governments can achieve a reduction in the consumption of the riskiest drugs through legalization. The most significant reductions from such a price decrease are in the consumption of amphetamine (16.2%), MDMA (16.1%), and cocaine (15.6%). The smallest reduction will occur for $\alpha-$PVP (7.8%). More broadly, we find in a series of experiments that for every four additional doses of cannabis used, approximately one less non-cannabis

[2]Price reduction was found in studies that followed previous instances of legalization (Anderson et al., 2013; Hall et al., 2023). This approach is also valid for other aspects of legalization, such as diminished risks related to purchase and the elimination of the stigma of illegality, provided that their influence on utility is uniform across consumers and thus has a monetary equivalent.

dose would be consumed. Substitution is not limited to drugs with medium risks and also occurs from high-risk ones, such as cocaine.

Second, we study the introduction of new drugs. In recent years, synthetic drugs have gained popularity in many countries. We study the impact of their introduction on overall drug use, accounting for the fact that a portion of the demand for these new drugs represents substitution from pre-existing drug types. We focus on a class of drugs known as "bath salts," which have an extremely large market share in Russia, accounting for nearly half of all drugs sold. We simulate our model with all bath salts eliminated from consumers' choice sets. We find that the introduction of bath salts led to a 40% increase in the total demand for illegal drugs. Although substitution from preexisting drug types was substantial, the effect of these new drugs on overall drug use was sizeable. This result underscores that governments should allocate resources to prevent the introduction of new drugs.

Third, we apply our estimates to study the effects of targeted enforcement on the amount of consumption of all drugs. To that end, we analyze how the demand for illegal drugs would be affected if a particular drug were eliminated. We conceptualize this scenario as an extreme case of successful supply-side interventions by the government. We observe that the impact on total consumption is the smallest for drugs that have close substitutes, namely amphetamine, MDMA, mephedrone, hashish, and marijuana. Our findings suggest that the most substantial effects occur for drugs that have no close substitutes, such as $\alpha-$PVP, cocaine, and opioids. Specifically, we find that the share of consumers who switch to a substitute is two times larger after eliminating amphetamine than after eliminating $\alpha-$PVP.

Finally, we study whether our estimates support the concern of Becker et al. (2006) that drug enforcement can increase the total revenue of the black market if demand for drugs is inelastic. We find that the effect on total revenue is always negative but varies significantly across drugs. Specifically, we observe that targeting substances with many substitutes, such as amphetamine, is more likely to increase the total revenue of drug sellers. This is because enforcement increases the revenue from the substitutes of the targeted drug.

Our analysis has several limitations. First, we do not directly observe specific transactions. Instead, our data provides several proxies for drug sales, which are valid under a set of assumptions about the dependence between sales, listings, and reviews. We provide evidence to support the validity of our proxies. Despite this limitation, data from a dominant marketplace for illegal drugs can offer higher data quality than what researchers have previously had to use for studying the demand for illegal drugs. Second, we model consumer preferences as static. This implies, in particular, that the stock of addiction is fixed in the model. Hence, our model primarily addresses short-term substitution and cannot, for instance, capture the potential effects of cannabis as a gateway drug. Measuring these effects is beyond the scope of this paper.

Our paper makes two contributions to the literature on the estimation of demand for illegal drugs. First, by utilizing data scraped from a large marketplace, we obtain high-quality information about the consumption and prices of drugs.[3] Because of the illegal nature of the drug trade, researchers have generally been unable to access transaction data, which has long been recognized as a major problem (Manski et al., 2001).[4] As a result, prior research on the demand for drugs has been forced to rely on proxy measures of consumption, such as emergency department visits (Caulkins, 2001; Dave, 2006), traffic fatalities (Anderson et al., 2013), toxicology tests of arrestees (Dave, 2008), self-reported information from surveys (DeSimone and Farrelly, 2003; Van Ours and Williams, 2007), small-scale experiments (Jofre-Bonet and Petry, 2008; Olmstead et al., 2015), and user feedback on marijuana purchases (Davis et al., 2016). To estimate prices, researchers have often relied on recorded purchases made by undercover drug enforcement agents (Saffer and Chaloupka, 1999). However, this data has a low frequency, a number of methodological shortcomings (Manski et al., 2001), and over-represents large transactions (Horowitz, 2001).

Second, we are the first to study demand for the full set of drugs popular in a particular market. To the best of our knowledge, we obtain the first estimates of price elasticities for new and highly popular synthetic drugs like mephedrone. Crucially, we study substitution between drug types that comprise nearly the whole drug market. In contrast, previous studies have predominantly considered substitution between just two drugs or sin goods (DeSimone and Farrelly, 2003; Anderson et al., 2013; Powell et al., 2018).[5] Our review data provides a unique opportunity to observe substitution between drug types. See Gallet (2014) for a meta-analysis of the literature on demand for illegal drugs.

More broadly, our paper is related to the literature on the effectiveness of supply-side drug policies. We contribute to it by studying the effects of policies on total drug use while incorporating substitution between drug types. The literature has also investigated how the risks induced by policing, punishment, and incarceration affect the profits of drug dealers and the prices of illegal drugs (Levitt and Venkatesh, 2000; Kuziemko and Levitt, 2004). Several papers have studied the effect of interventions on cartel violence (Angrist and Kugler, 2008;

[3] To the best of our knowledge, we are the first to utilize data scraped from a darknet marketplace for demand estimation. Other examples of papers in the economic literature that utilize data scraped from the dark web include Červenỳ and van Ours (2019), Bhaskar et al. (2019), and Espinosa (2019).

[4] In some specific cases, researchers could utilize data from the regulated trade of drugs that are considered illegal in other contexts. For instance, Van Ours (1995) and Liu et al. (1999) employ data from actual transactions during the regulated opium trade of the early 20th century. Hollenbeck and Uetake (2021) estimate the demand for legalized marijuana using the BLP framework.

[5] An exception is Ramful and Zhao (2009), who study the extensive margin of drug use for three drug types: marijuana, cocaine, and heroin.

Dell, 2015; Castillo et al., 2020). Becker et al. (2006) provide a seminal theoretical analysis of drug enforcement.

We present a structural model of the demand for drugs. Other structural models of the illegal drug market have focused on particular drug types and did not take potential substitution between drugs into account (Kennedy et al., 1993; Galenianos et al., 2012; Adda et al., 2014; Mejia and Restrepo, 2016; Galenianos and Gavazza, 2017). A structural model allows us to study a range of counterfactuals. This is in contrast to papers such as Dobkin and Nicosia (2009); Dobkin et al. (2014); Moore and Schnepel (2021), which used events studies to estimate the impact of isolated large-scale shocks. These shocks are, by nature, rare, while our model can be used to assess routine policy responses.

The rest of the paper is organized as follows. Section 2 describes the context of illegal drugs and the operation of the marketplace. Section 3 describes the data we use. Section 4 presents our demand model and the details of its estimation. In Section 5, we apply our model to calculate the effects of several supply-side policies. Section 6 concludes.

2. Market for Illegal Drugs

For researchers, a key feature of the illegal drug market is the difficulty of observing it. In this paper, we address this issue by leveraging a unique setting within which most of the drug trade was concentrated on a single website called Hydra. We describe this online platform and the related context in this section.

2.1. The Hydra marketplace

The Hydra marketplace operated on the Tor network, which allows for encryption and anonymization of traffic by routing it through a series of volunteer-run servers. The network can be accessed using a specialized browser that is available for free download. Because of the anonymization of traffic, the government cannot restrict access to websites on the Tor network in the same way it can with conventional websites.

As on other darknet markets, participation in the platform was anonymous. The marketplace began operation in 2015 (VICE, 2020) and primarily served the Russian market. After its predecessor RAMP was shut down by the Russian police in 2017, Hydra grew without any significant competition until its shutdown[6] in 2022. The unprecedented length of its existence allowed Hydra to become the largest darknet marketplace in the world. The U.S. Department of Justice estimated that Hydra accounted for 80% of all darknet market cryp-

[6] The shutdown was a joint operation of German and U.S. law enforcement.

tocurrency transactions in 2021 (States of America V. Dmitry Olegovich Pavlov, 2022). U.S. Department of the Treasury (2022) estimates that the yearly revenue of Hydra in 2020 was $1.3 bln. At the time of its closure, Hydra had spread to the majority of cities in Russia and is thought to have been the most popular method to purchase drugs for retail consumption in several of the largest cities (Goonetilleke et al., 2023).

Our data allows us to compute market shares of different drugs in specific locations. This is due to the fact that Hydra used a dead-drop distribution system, unlike most darknet marketplaces, which primarily deliver drugs through legitimate postal services (VICE, 2020). The system of dead-drops was first adopted by RAMP in response to a 2014 law that required the postal service to inspect packages for illegal substances (Saidashev and Meylakhs, 2021). To circumvent this, shops hired couriers to hide drugs throughout the city prior to purchase. The details of these hidden drugs would then be posted on the marketplace so that potential customers could select the listing that best suited their requirements. Appendix Figure F.2 provides an example of a page with listings. After purchasing the drug, consumers received information about the exact location of the dead-drop (see Appendix Figure F.5 for an example). While a small proportion of drugs were still delivered via mail, the majority of drugs sold for retail consumption were delivered via dead-drops.[7] Shops recruited couriers on the platform, posting job offers on the website.

Similar to other darknet marketplaces (Janetos and Tilly, 2017), there was a reputation system. Buyers could review purchases, providing a numeric rating and detailed text comments. In addition, there were marketing options for shops. One of the key forms of advertising available to shops was purchasing one of the 20 positions on the main page of the website (see Appendix Figure F.1 for an example). These positions were allocated via an auction and served to increase the visibility of the shops allocated these positions (Goonetilleke et al., 2023). These characteristics of shops appear to have been important factors in buyers' choice process and thus will be incorporated into our demand model in Section 4.

One limitation of our data is that sales of fentanyl and several other synthetic opioids were prohibited in the marketplace. Thus, our analysis is not informative about demand for this drug. Goonetilleke et al. (2023) provides a detailed discussion of the scope, structure, and rules of Hydra.

[7] Drugs delivered via mail were primarily drugs that are particularly difficult for law enforcement to detect, such as LSD.

2.2. Drug types

In Table 1, we describe the characteristics of the most popular drug types on Hydra.[8] The medical literature divides drugs according to their pharmacological effects. Stimulants, like MDMA,[9] cocaine, amphetamine, and methamphetamine, increase the activity of the central nervous system. Two other types of stimulants in our data are highly popular: mephedrone and $\alpha-$PVP.[10] These substances belong to the new family of drugs known as synthetic cathinones, which emerged in the late 2000s and are colloquially referred to as "bath salts."

Depressants, like cannabis, opioids, and GHB, decrease the activity of the central nervous system. Cannabis can be distributed in different forms, the two most popular of which are marijuana buds and hashish. Hashish is produced by compressing and processing cannabis plants. Opioids, such as heroin and methadone, produce morphine-like effects and are commonly recognized for their potent effects and significant risks.[11] GHB is a substance that can be used for medical purposes but is also a popular recreational drug. Finally, for hallucinogens, such as LSD and psilocybin, the main effect is the altered perception of reality.

Drugs also differ in other dimensions. In particular, some drugs are considered "club" (or "party") drugs. These drugs are popular among younger individuals and are typically consumed at bars, nightclubs, concerts, and parties (Bowden-Jones and Abdulrahim, 2020). Drugs can also vary in their most common mode of administration, but each drug type typically can be administered in multiple ways. The mode of administration can impact the likelihood of developing dependence (Hatsukami and Fischman, 1996).

[8] See Manski et al. (2001) for a summary of the medical literature that examines the properties of illegal drugs.

[9] MDMA is often referred to as ecstasy.

[10] Mephedrone is also known as 4-methylmethcathinone and is often referred to as "meow meow." $\alpha-$PVP is also known as $\alpha-$pyrrolidinovalerophenone and is often referred to as "flakka."

[11] However, methadone can also be used for medical purposes, in particular, as a treatment for heroin addiction. Methadone therapy is not legal in Russia (Idrisov et al., 2017).

Table 1: Drug types most present on Hydra

Psychoactive class	Drug type	Club drug	Bath salt	Production	Administration	Form	Physical harm index	Dependence index	Overdose ratio
Stimulants	α-PVP	Y	Y	Synthetisized	Smoking, nasal, oral	Powder, crystal	—	—	—
	Amphetamine	Y	N	Synthetisized	Oral, nasal	Powder, crystal	1.81	1.67	—
	Cocaine	Y	N	Synthesized from organic compounds	Nasal	Powder	2.33	2.39	15
	MDMA	Y	N	Synthetisized	Oral	Pills, crystal	1.05	1.13	16
	Methamphetamine	Y	N	Synthetisized	Oral, smoking	Crystal, powder	—	—	10
	Mephedrone	Y	Y	Synthetisized	Oral, nasal	Powder, crystal	—	—	—
Depressants	GHB	Y	N	Synthesized	Oral	Liquid	0.86	1.19	8
	Hashish	N	N	Organic	Smoking	Resin	0.99	1.51	> 1000
	Heroin	N	N	Synthesized from organic compounds	Injection	Powder	2.78	3.00	6
	Marijuana	N	N	Organic	Smoking	Flowers, leaves	0.99	1.51	> 1000
	Methadone	N	N	Synthetisized	Oral	Powder	1.86	2.08	20
Hallucinogens	LSD	Y	N	Synthetisized	Oral	Saturated paper	1.13	1.23	1000
	Psilocybin	N	N	Organic	Oral	Mushrooms	—	—	1000

Note: “Administration” refers to the most common method of drug administration. “Form” refers to the most common substance forms in which drugs are sold. Harm and dependence indexes are sourced from Nutt et al. (2007). Overdose ratios, also known as safety ratios, are defined as the ratio of the acute lethal dose to the dose most commonly used and are obtained from Gable (2004).

There are several measures of the harmfulness of different drugs. First is the "overdose potential," which Gable (2004) defines as the ratio of the acute lethal dose to the most commonly used dose. Second, we consider the physical harm and dependence indexes provided by Nutt et al. (2007).[12] Based on all of these measures, heroin emerges as the most dangerous drug in our sample. Methadone and cocaine also are characterized by high levels of dependence and harm, while cannabis and MDMA have smaller risks according to these measures. Safety ratio and harm and dependence indexes are unavailable for bath salts, which are newer drugs. However, Patocka et al. (2020) report a significant risk of overdose associated with α-PVP, and Winstock et al. (2011) document multiple harmful effects attributed to mephedrone. It should be noted that there is no universally accepted measure of total harm for drugs. Overdose potential does not take into account overdose risks related to mixing drug types or the inability of consumers to control the exact dosage. It also ignores all kinds of risks not associated with overdoses. The indexes of Nutt et al. (2007) have been widely debated, in large part because of the difficulty of separating and correctly identifying individual and societal harms (Caulkins et al., 2011).

2.3. Production

Drugs differ significantly in production technology and whether they are imported or produced locally. In the Russian, U.S., or European markets, some drug types are entirely imported. For example, the production of cocaine requires the leaves of the coca plant, which is grown almost exclusively in Colombia, Peru, and Bolivia (UN Office on Drugs and Crime, 2022). Similarly, the production of heroin requires the opium poppy, which is mostly grown in Afghanistan (UN Office on Drugs and Crime, 2016). The origin is less clear for other drug types. Cannabis plants can be grown indoors, and, thus, marijuana can be both imported or produced locally. However, the production of hashish for the European market is mainly concentrated in Morocco and Afghanistan (UN Office on Drugs and Crime, 2016). There is evidence that bath salts for the Russian market are produced locally from precursors imported from chemical suppliers in China (Baza.io, 2020).

We leverage the differences in production across various drug types in two ways. First, we consider distance to the main ports as an instrument for price, under the assumption that the cost of cocaine and hashish increases with distance from the point of entry. Second, we use a policing shock that affected the supply of Moroccan hashish to validate our demand model.

[12]These indexes are derived from averaged scores provided by surveyed experts, where a score of 0 indicates no risk and scores of 1, 2, and 3 represent some, moderate, and extreme risk, respectively.

2.4. Regulation

All drugs listed in Table 1 are classified as illegal substances in Russia. In particular, the process of picking up drugs purchased on Hydra was risky for consumers because drug possession is illegal. The range of possible punishments depends on the amount registered during the arrest. The government defines several thresholds, which vary depending on the type of drug (Government of Russia, 2012). The threshold closest to retail purchases is the "significant amount." Consumers with an amount registered above this threshold can receive a prison sentence.[13] This can potentially affect the preferences of consumers, decreasing demand for dead-drops of large amounts. We address this possibility by including the size of the dead-drop in the set of product characteristics.

The high risks associated with consumption imply that a legalization policy will increase demand for the legalized drug. In Section 5.1, we study the effects of legalization, assuming that a reduction in the non-monetary costs of consumption generates the same substitution patterns as a reduction in the price of the legalized drug.

3. Data

3.1. Datasets

3.1.1. Listings

Our first source of data contains information on listings on Hydra. We collect this data by scraping the Hydra website.[14] We describe the details of the scraping process in Appendix Section A.1. Our dataset covers all listings on the platform for 31 days between June 2019 and September 2020. See Appendix Table A.1 for the list of available dates.

Each shop could offer several products. For each product, the platform displays the listings offered by the shop. The platform allowed two types of listings: pre-order and instantaneous listings. Instantaneous listings provided the details about dead-drops that were already hidden in the city and could be purchased immediately. Pre-order listings allowed consumers to contact the shop to buy the drug, which was then deposited after the

[13] The Russian law distinguishes possession without the intent to supply (Article 281), where possession includes purchasing, storing, producing, or transporting drugs, and drug offenses with the intent to supply (Article 281.1). Among the total of 531,998 people convicted in 2020, 56,556 people were convicted for Article 281 (Judicial Department, 2021a). Among them, 25% received a prison sentence (Judicial Department, 2021b). In the same year, 14,223 people were convicted for Article 281.1, and among them, 92% received a prison sentence, with 65% receiving a sentence longer than five years. Risks can be very high for all consumers, as the legal system tends to over-classify cases as offenses with the intent to supply.

[14] We are indebted to Ekaterina Aleksandrova for her invaluable contribution to data collection efforts.

purchase was made. As noted by Goonetilleke et al. (2023), pre-order listings are more likely to be used for wholesale transactions or for more "exotic" drugs and, thus, constitute only a small fraction of all purchases. Furthermore, they are only loosely connected to sales, as posting pre-order listings did not imply any pre-commitment on the shop's part. Therefore, we restrict our attention to instantaneous listings. Overall, we observe more than 3,410,000 instantaneous listings across 40 different drug types, 8,283 shops, and 1,337 different cities or towns.

As instantaneous listings describe the pre-hidden drugs, we observe the characteristics of each dead-drop, including the mode of hiding, amount (weight or counts), approximate location, and price. Our data also includes the information that was displayed as a part of the product's description: shop name and shop identifier, product name, drug type, and average ratings for the shop and product. Finally, we observe the approximate cumulative sales for the shop (this number was displayed by the platform, see Appendix Figure F.3 for an example). We describe the details of data cleaning in Appendix Section A.2.

Definition of dose. For two reasons, the units listed by shops can be non-comparable for some listings. First, drug potency per gram can be different for different substances. Second, drugs can be sold in different substance forms. For example, MDMA is sold both as pills and crystals, which complicates the calculation of market shares for this drug. Moreover, GHB is typically sold in liquid form, unlike most other drugs. We solve the problem of non-comparability by normalizing listed amounts using the "standard amount" for each drug-form combination. We call standard amounts "doses" for the sake of simplicity. Doses describe the smallest frequently purchased amount within the drug-form pair.[15] We discuss the details in Appendix Section A.3 and Appendix Table A.2.

3.1.2. Reviews

The second dataset we use is purchased from a private firm, which was spun out from the CyLab Security and Privacy Institute at Carnegie Mellon University.[16] This dataset allows us to see a large subset of the customer reviews posted on the platform. For each review, we observe the associated text, the purchased product, the shop, the nickname of the reviewer, the time when the review was posted, and the numerical rating the buyer has given. We end up with approximately 215,000 reviews. We observe reviews for 784 shops on Hydra, which

[15]Our definition of a dose can be larger than the typical amount consumed if dead drops were intended to contain enough drugs for multiple consumptions.

[16]The firm scrapes data from several dozen darknet marketplaces and also provided data for United Nations Office on Drugs and Crime (2021). Details about the project can be found in Soska and Christin (2015) and Christin (2022).

account for 47% of the shops in our listings data.

3.1.3. Supplementary data from Hydra

We use additional data sources to derive characteristics of products that are likely to be relevant from the consumers' perspective. We include these variables in the demand model. First, we use scrapes of the front page to identify whether a particular shop was advertised on the front page each month (see Section 2.1). Considering that consumers may prefer shops with established reputations, we incorporate into the model the duration of a shop's presence in the marketplace. To achieve this, we access historical data on the aggregate number of reviews on a shop-drug level. This data covers reviews left by Hydra users from six different cities since late 2016 and until 2021.[17] We use this information to identify the age of shops on Hydra.[18]

3.1.4. Demographic data

We use a 10% subsample of the Russian Census of 2010 to obtain city-level data on population and estimate the market size. While we observe listings from more than 1,000 cities, we restrict our attention to the country's largest cities. We exclude cities with a small Hydra presence (defined as the ratio of listings to population). We also observe an unusually high presence in satellite cities around Moscow, which we interpret as evidence that dead-drops hidden in these communities also serve consumers from Moscow. For this reason, we regard these locations as the same market as Moscow. We use 34 large cities in our data.

3.2. Descriptive statistics

Table 2 presents summary statistics for the most popular drug types traded on Hydra, where listings were restricted to the cities of interest. The four most popular drug types are stimulants: mephedrone, amphetamine, $\alpha-$PVP, and MDMA. They are followed by marijuana, hashish, and cocaine. The popularity of other drug types on Hydra was substantially lower. The most expensive drug observed is cocaine, with an average price per dose of ≈$65. All other drugs are significantly cheaper, with the price per dose about three times lower. Among drug types with large market shares, $\alpha-$PVP is the cheapest, with an average price per dose of $14.

[17] This data was purchased from an independent data collector, who also provided data for several media investigations (Knife Media, 2020; Proekt, 2019). As this data does not contain price information, we cannot use it for demand estimation.

[18] Each shop on the platform had a unique ID. We conclude that each new shop received an ID incremented by 1. Thus, we can establish the approximate date of registration for any merchant from its ID.

Table 2: Summary statistics for select drug types on Hydra

	Daily Listings	Average Price ($ per dose)	Median Quantity	Market Share (doses)	# of Sellers
Mephedrone	11,168	23	2g	28%	2,200
Amphetamine	5,345	16	2g	15%	1,738
α-PVP	4,808	14	1g	11%	1,377
Marijuana	3,041	20	2g	8%	1,906
Hashish	3,040	16	2g	9%	1,831
MDMA (pill)	2,697	18	3 counts	7%	1,428
Cocaine	2,602	65	1g	6%	1,059
MDMA (crystal)	1,368	19	1g	4%	673
LSD	1,147	26	3 counts	3%	493
GHB	804	13	100ml	1%	76
Methadone	752	22	0.5g	2%	331
Heroin	293	20	0.5g	1%	201
Other Cannabis	705	13	2g	2%	690
Other Psychedelics	739	17	3 counts	2%	399
Other Stimulants	342	19	1g	1%	217
Other Opioids	11	30	1g	$< 1\%$	12

Note. Data includes listings from 34 cities of interest. Prices are converted to USD using an exchange rate of 74 RUB per USD.

The last column of Table 2 shows that each drug type could be purchased from a large number of shops. Goonetilleke et al. (2023) present additional evidence that Hydra had a high degree of competition. This motivates our approach to restrict attention to the demand side only. Changes in demand should be close to the equilibrium change in consumption if the supply of drugs is elastic.

Table 3 describes variation in prices of drug listings on Hydra. Several patterns are noticeable. First, there is considerable variation in price across shops, which suggests that the quality of products might vary across shops. To account for this, we include a set of proxies for quality in our demand model. In particular, we include a set of shop characteristics, such as the shop's age and average rating and whether the shop advertises itself on the marketplace's front page. Appendix Table A.3 presents summary statistics of the variables that describe shops on the platform. Second, the last column of Table 3 shows that the proportion of price variation that occurred over time was the largest for hashish. This is explained by the policing shock to the supply of this drug that occurred during the sample period. In Section 4.3, we use this variation to validate our model estimates.

Table 3: Variation in prices

	Price per gram			Variation		
Drug type	Mean	Median	Std	Shop	City	Time
α-PVP	24	23	5	29.7%	18.3%	29.3%
Amphetamine	15	14	4	68.1%	6.7%	0.3%
Cocaine	124	122	23	63.7%	7.4%	2.3%
Hashish	15	12	6	20.2%	3.7%	61.0%
Marijuana	19	18	5	34.9%	5.2%	38.1%
MDMA	35	34	8	65.9%	11.4%	3.1%
Mephedrone	21	21	5	45.3%	18.3%	0.8%

Note. The variation due to each factor represents the ratio of the total variance explained by corresponding fixed effects in the regression of price on shop-level, city-level, and date-level fixed effects. Only the median quantity for each drug type is considered.

3.3. Proxies for sales

Because we cannot directly observe sales on Hydra, we use instantaneous listings as our proxy for sales. To calculate market shares of different products, we assume that the number of listings with specific characteristics is proportional to the number of transactions with the same characteristics. Our assumption is motivated by the fact that depositing each dead-drop was expensive for shops and required a substantial payment to the courier (Goonetilleke et al., 2023). Therefore, posted instantaneous listings represent a strong commitment to sell.

The credibility of this assumption is supported by the strong correlation between listings and several proxies for sales on Hydra that are available on aggregate levels. First, because we observe rounded total shop deals for each shop, we can calculate the difference between the total deals at the end and the beginning of the observed period. Table 4 demonstrates that the correlation between the observed change in total shop deals and the number of listings is 0.7.

Second, the number of reviews in our data can serve as another proxy for shop sales. As can be seen in Table 4, the shop-level correlation between the number of reviews and the number of listings is 0.62. It is important to note that because of the inherent noise in these proxies, the observed correlations between listings and the two proxies for sales should be lower than the actual correlation between listings and sales. In particular, the difference in deals is susceptible to large rounding errors[19] and the total number of reviews in our dataset suffers from incomplete scraping coverage.

[19] For example, the platform displayed the same number of total deals for shops with 149,000 and 100,000 actual sales.

Table 4: Correlations between different proxies for sales

	Listings observed	Cumulative deals	ΔCumulative deals	Reviews observed
Listings observed	-	0.74	0.70	0.62
Cumulative deals	0.74	-	0.93	0.65
ΔCumulative deals	0.70	0.93	-	0.70
Reviews observed	0.62	0.65	0.70	-

Note. Correlations are reported on the shop level. "Listings observed" stands for the total number of listings in the data for the shop. "Cumulative deals" stands for the approximate total number of sales by the shop as displayed by the platform. "ΔCumulative deals" stands for the change between the first and the last days when the shop is observed in the data.

Finally, listings exhibit a strong time correlation with cryptocurrency inflows to Hydra. Flashpoint, Chainanalysis (2021) provide estimates of the monthly revenue of Hydra over time. These estimates are based on counting transactions in the Bitcoin blockchain to wallets that analysts identified as belonging to Hydra. Appendix Figure A.1 demonstrates the similarity between the dynamics for listings and cryptocurrency inflows.

The assumption that listings are proportional to transactions comes with several caveats because on any given day, several potential mechanisms could cause the number of listings to differ from the number of sales. First, it is highly possible that it took several days for a particular listing to be sold. Second, each courier likely deposited several dead-drops in the same neighborhood. If a shop had multiple dead-drops of the same drug type, amount, hiding mode, and price in the same approximate location, these dead-drops appeared on the marketplace under the same listing.[20] Therefore, one listing could potentially correspond to several transactions. Third, a particular dead-drop could remain unsold, resulting in no transactions. Fourth, shops could list one dead-drop under several adjacent neighborhoods to maximize their presence in search results. Finally, some transactions on the platform were conducted via pre-orders.

Because our estimation procedure is based on market shares of different drug types, it would be affected if the ratio of listings to dead drops was different across drug types. For example, it could be the case that, first, several dead-drops of the same weight and hiding type are hidden in the same neighborhood and, second, that this is most common for the more popular drug types. In this case, listings could disproportionately underestimate

[20]This feature of the platform was described in the instructions published on the website.

transactions for more popular drugs. To examine this possibility, we employ reviews as an alternative proxy for sales. Appendix Table A.5 presents the market shares of different drug types, as defined using both listings and reviews. While we observe some minor differences, overall, we find that the shares calculated using both methods are very close to each other. This supports our assumption that listings are proportional to sales.

3.4. Proxies for quality

Potential heterogeneity in unobserved quality creates an identification problem in demand estimation. Quality is likely to be positively correlated with price and demand, so not including quality in the model can lead to an underestimation of the sensitivity of demand to price. In the context considered here, "quality" is likely to be related to the purity and potency of the drug or the ease of recovering the dead drop. To mitigate the issue of confounding quality, we incorporate quality into the analysis by using user ratings as a proxy. However, there are three major concerns about using ratings as a proxy for quality. First, ratings left by users on online platforms are likely to be a function of both quality *and price*.[21] Luca and Reshef (2021) find that ratings can be highly responsive to prices: they estimated that a price increase of 1% leads to a 3%-5% decrease in the average rating. Second, Filippas et al. (2022) show that ratings on online platforms are subject to inflation. Finally, ratings on Hydra had very low granularity: 94% of all reviews we observe had ratings 10/10, with an average rating of 9.6 (see Appendix Table A.4). Average ratings of shops and products displayed by the platform were even higher because the platform automatically assigned the highest rating to an order if the consumer did not rate it. Thus, ratings can contain little information about underlying product quality.

To address the issue, we construct another proxy for quality by employing natural language processing techniques. We determine the sentiment of each review in our dataset. Filippas et al. (2022) suggest that written feedback can provide more information about the fundamentals of obtained quality. We find evidence that supports this idea in the context of Hydra. Appendix Table A.4 shows that an average review consists of 16 words, indicating the potential informativeness of the reviews. Moreover, Appendix Table B.1 provides examples of reviews with negative sentiment despite the consumer awarding the highest possible rating. We utilize the average sentiment score of reviews for each shop as an additional proxy for quality. Further details about the text analysis are provided in Appendix B.

[21] This is possible under two feasible models of consumer behavior. First, users can rate not the quality of goods purchased but their net utility, which decreases with price. Second, users can rate based on their satisfaction relative to the reference point of expected quality, which, in turn, is likely to be positively correlated with price.

3.5. Reviewer behavior

Data on reviews provides us with a unique opportunity to elicit information about the behavior of individual drug consumers. We identify reviews made by particular consumers using usernames displayed on the platform. Appendix Table A.4 presents summary statistics for our review data. Because we only observe a subsample of reviews on the platform, most reviews are not included in our sample, and some users appear in our data only once. We have identified 43,381 users who have left more than one review in our dataset, almost half of whom reviewed different types of drugs, suggesting the significance of substitution for illegal drugs. For users with more than one review, we have a total of 132,855 reviews.

We use these cases to estimate the correlation between choices made by the same consumer over time. Figure 1 presents a matrix of drugs' market shares conditional on reviewing a particular drug. Each element P_{jk} of this matrix represents the probability that for a random review for drug j, a random different review by the same user is for drug k.[22]

Several patterns are noticeable in Figure 1. First, the diagonal elements of the matrix are the largest. This means that conditional on having consumed a drug, most consumers purchase the same drug in other periods. However, the degree of attachment differs between drugs. For example, the share of opioids is very large conditional on reviewing an opioid, even though opioids have a small market share. Cocaine consumers also have a high probability of consuming cocaine in other periods. This is in line with the dependence index (DI) provided in Table 1, which is the largest for opioids and cocaine (heroin has a DI equal to 3.0, methadone has a DI of 2.08, and cocaine has a DI of 2.33). Moreover, the ranking of drugs by dependence index coincides with the ranking by diagonal elements for the most popular drugs: amphetamine (a DI of 1.67), cannabis (a DI of 1.51), and MDMA (a DI of 1.13). At the same time, while a DI is not available for the two bath salts, $\alpha-$PVP and mephedrone, our estimates show that consumers of these drugs have a high degree of attachment. This suggests that dependence on bath salts is comparable to that of cocaine.

Second, the matrix provides evidence of the substitution between drug types. For four popular drug types (amphetamine, hashish, marijuana, and MDMA), the diagonal elements are around 50%, indicating that consumers are almost as likely to buy another drug in other

[22] For $j \neq k$, the elements P_{jk} of this matrix are given by the share of k in all other reviews by the same user weighted by the number of reviews this user left for j:

$$P_{jk} = \sum_{i:R_i>1} \frac{R_{ij}}{\sum\limits_{i':R_{i'}>1} R_{i'j}} \frac{R_{ik}}{R_i - 1} = \frac{1}{\mathbb{E}\big[R_{ij} \mid R_i > 1\big]} \mathbb{E}\Big[\frac{R_{ij}R_{ik}}{R_i - 1} \mid R_i > 1\Big]. \tag{1}$$

For $j = k$, the elements are $P_{jj} = \big(\sum_{i:R_i>1} R_{ij}(R_{ij} - 1)/(R_i - 1)\big)/\big(\sum\limits_{i':R_{i'}>1} R_{i'j}\big)$. Here, R_{ij} is the total number of observed reviews for product j by user i and $R_i = \sum_j R_{ij}$ is the total number reviews by user i.

Figure 1. Conditional shares of different drug types

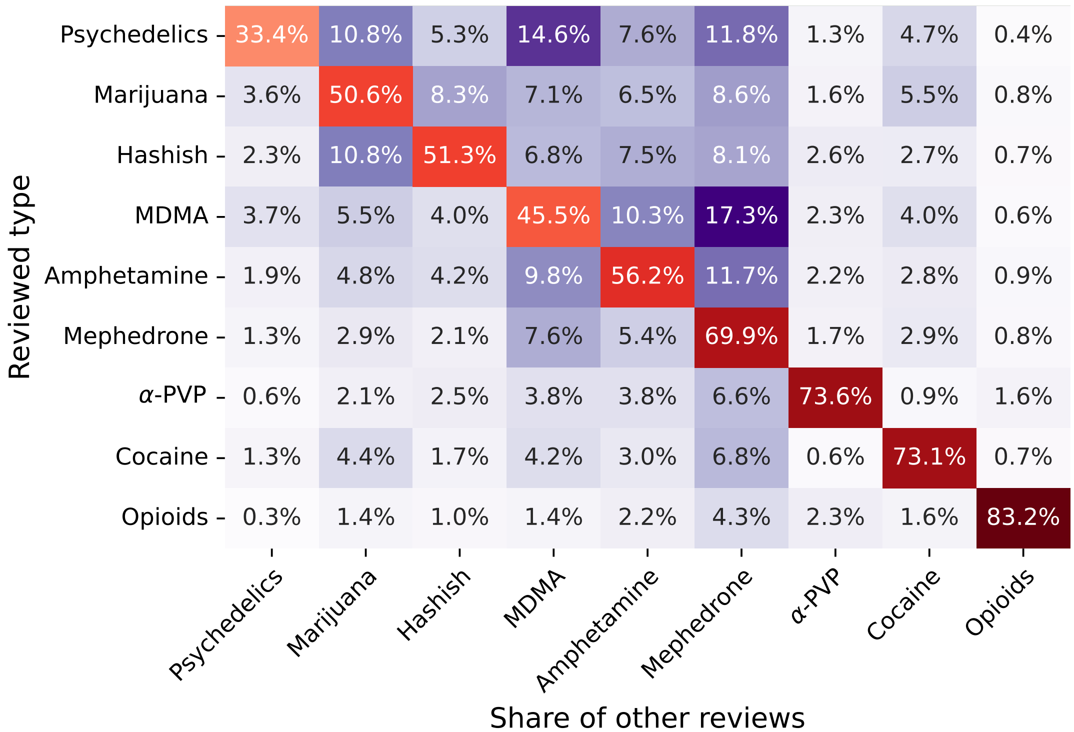

Note: Opioids include heroin and methadone. Psychedelics include mushrooms and LSD.

periods as they are to buy the same drug type.

Third, we see that substitution is more likely to happen within certain groups of products. The following group of drugs is particularly noticeable: amphetamine, MDMA, and mephedrone. For example, for a consumer who purchased MDMA, the two other most likely choices are mephedrone (17.3%) and amphetamine (10.3%). Similarly, for a consumer who purchased amphetamine, the two other most likely choices are mephedrone (11.7%) and MDMA (9.8%). Finally, for a consumer who purchased mephedrone, the two other most likely choices are MDMA (7.6%) and amphetamine (5.4%). The similarity of their effects might explain this pattern. In Section 2.2, we demonstrate that all these drugs are stimulants, have the same mode of administration, and fall into the category of club drugs. Medical studies also found strong similarities in the effects induced by these drugs (Poyatos et al., 2022). However, it would be difficult to predict the observed substitution patterns *ex-ante* only on the basis of drug characteristics. For example, $\alpha-$PVP is another popular bath salt with similar characteristics, but consumers who purchased any of the three drugs switched to it much less often.

Another cluster in Figure 1 is hashish and marijuana. These two drugs are produced from cannabis plants, belong to the same psychoactive class, and have the same administration. Marijuana is the second choice for hashish consumers in our data.[23] At the same time, these two drugs have substantial flow to and from the three stimulants mentioned above.

In Figure 1, we provide novel evidence of large taste heterogeneity between drug consumers. The observed patterns can be interpreted in terms of drug properties, and we can expect them to affect how consumers substitute between drug types. A model without consumer heterogeneity – for instance, the standard multinomial logit model – would fail to generate realistic predictions about substitution. In the next section, we develop a demand model that can reproduce the observed patterns.

4. Demand Model

To account for the taste heterogeneity documented in Section 3.5, we use the BLP approach, which was introduced in Berry (1994) and Berry et al. (1995) and has become the workhorse method for demand estimation. Because we are particularly interested in substitution between different drug types, we define products as individual drug types. The indirect utility that a consumer i can obtain from buying a drug of type j in city c in period t is given by

$$U_{ijct} = -\alpha p_{jct} + x_{jct}\beta + \sum_{g \in G} \lambda_i^g I(j \in g) + \xi_{jct} + \varepsilon_{ijct}, \tag{2}$$

where p_{jct} is the average price per dose of drug j in city c in period t, x_{jct} is a vector of observed product characteristics, ξ_{jct} are unobserved product characteristics, and ε_{ijct} are taste shocks independent from other random variables.

Product characteristics x_{jct} include dummies for each drug type, number of doses, hiding method, substance form, and proxies for quality and marketing activities by shops.[24] We also include date-level fixed effects to account for the growth of the platform and potential differences in drug consumption across seasons and days of the week. We include time trends for mephedrone and amphetamine because these drugs exhibited growth in popularity over our sample period. Preferences for product characteristics p_{jct} and x_{jct} are given by α and β and are the same for all consumers.

We include random coefficients for dummies for several of the most popular drugs or broader drug categories $g \in G$. This is the key component of the model that allows us to

[23]The same is not true for marijuana, for which the second choice is mephedrone. This can be explained by the fact that hashish was disappearing from the market during the period we observe (see Section 4.3).

[24]Due to aggregation, categorical characteristics (e.g., hiding type) are converted to the proportion of listings that have this characteristic. Continuous characteristics are converted to the average across all listings within the given product.

incorporate idiosyncratic attachment to particular drugs and correlation in preferences for different drug types. Random coefficients vary across consumers and are given by

$$\begin{pmatrix} \lambda_i^1 \\ \vdots \\ \lambda_i^K \end{pmatrix} = \Sigma \begin{pmatrix} \nu_i^1 \\ \vdots \\ \nu_i^K \end{pmatrix}, \tag{3}$$

where we assume that ν_i are drawn from the multivariate standard normal distribution, and, thus, the covariance between random coefficients is equal to $\Sigma^T\Sigma$. Because including multiple random coefficients is computationally expensive, we limit their number to $K = 6$. Therefore, we include them only for the drug types with the largest market shares. Because of their similarity, we use a common dummy variable for the two types of cannabis (hashish and marijuana). Otherwise, we choose not to impose any *ex-ante* restrictions and define g to be individual drug types for $\alpha-$PVP, amphetamine, cocaine, MDMA, and mephedrone.

Our interpretation of type-specific random coefficients is the following. A large value of $\mathbb{V}ar(\lambda^j)$ implies that a substantial fraction of consumers will have a large draw of λ^j and is likely to purchase drug j in many periods, being unwilling to substitute other drug types for it. However, a high value of $cov(\lambda^j, \lambda^k)$, where $k \neq j$, implies that consumers who have a large draw of λ^j are likely to also have a large draw of λ^k. In this case, a large proportion of consumers of j would be willing to substitute drug j for drug k. These type-specific random coefficients allow our specification to account for the heterogeneity in consumption patterns that we observe in Section 3.5. Heterogeneity in drug consumption is also widely discussed in addiction studies.[25]

The unobserved product characteristics ξ_{jct} are common across all agents for each market-drug combination and can be correlated with p_{jct}. This accounts for potential endogeneity, which is possible, for example, because of unobserved quality or because drug sellers strategically respond to aggregate-level demand shocks by adjusting prices. Consumers can also choose the outside option of not purchasing any drug, for which the indirect utility is normalized to be mean-zero: $U_{i0ct} = \varepsilon_{i0ct}$. In each period, the consumer chooses the option that provides the highest indirect utility.

Discussion. We include dummies for the most popular drug types to account for consumer heterogeneity. This approach allows for flexible substitution patterns between drugs and is feasible because we have a relatively small number of products in the model. An alternative approach could be to model drugs in characteristic space. However, in Section 3.5, we show

[25]See Reuter (2010). In particular, one can expect that males consume more drugs (Pacula, 1997) or that younger people may prefer party drugs.

that the characteristics provided in Table 1 fail to explain all relevant substitution patterns. It is inherently challenging to identify a small set of characteristics that would adequately capture the relevant differences between drugs. This is due to a combination of issues. To begin with, there is no consensus on how to measure attributes such as "pleasure," which likely play a significant role in determining drug consumption. Moreover, in cases where a well-defined measure does exist, such as the overdose ratio, it is not available for all types of drugs. Finally, the characteristics provided in the medical and chemistry literature are typically categorical and describe the grouping of substances into broader categories. The example of $\alpha-$PVP highlights that a simple categorization by psychoactive class cannot adequately capture the relevant substitution patterns. At the same time, introducing dummies for additional categorizations would rapidly inflate the dimensionality of the characteristic space.

4.1. Estimation

Identification of the non-linear parameters in the mixed logit model is a well-known challenge because it requires a large number of IVs, and aggregate data often does not have enough variation. To solve this we use micro-moment conditions that describe the reviewing behavior of consumers. In our demand model, two products, j and k, are close substitutes if their random coefficients are correlated. In this case, people who review drug j would also often review drug k. Thus, the panel structure of our review data can help identify the covariances between random coefficients.

4.1.1. Reviews

We capture correlation in tastes for particular drug types by matching co-movements of purchases for drugs j and k across consumers in the data for different pairs (j, k). We infer purchases from reviews. Specifically, our micro moments are averages of $R_{ij}R_{ik}$, where R_{ij} is the total number of observed reviews for product j by user i:

$$R_{ij} = \sum_t R_{ijt},\ R_{ijt} = I(\text{review by } i \text{ for drug type } j \text{ in period } t). \tag{4}$$

Intuitively, if consumers who like drug j usually like drug k, then $R_{ij}R_{ik}$ should be larger on average, all else being equal. These moments are also related to conditional market shares, as can be seen from equation 1. Our sample is subject to selection because it only has users with at least one review observed. Thus, the appropriate model counterpart of these quantities in the data is

$$\mathbb{E}\Big[R_{ij}R_{ik} \mid R_i > 0\Big] = \frac{\mathbb{E}R_{ij}R_{ik}}{\mathbb{P}\big(R_i > 0\big)}, \tag{5}$$

where $R_i = \sum_j R_{ij}$ is the total number of observed reviews left by consumer i.

To generate these values using our demand model, we must account for two possibilities. First, not every purchase was reviewed. Second, not every review was scraped by the data provider. We do this using the following framework. There are T periods over which consumers can make purchases and leave reviews. We assume that each purchase is reviewed randomly. The probability of leaving a review conditional on a purchase is allowed to depend on drug type and equals π_j^{review}.[26] We assume that the scraping process is also random, with the conditional probability of a given review being scraped equal to π_t^{scrape}. This probability is allowed to depend on time to account for potential imbalances in scraping over time.[27] The product of these numbers $\pi_{jt} = \pi_j^{review}\pi_t^{scrape}$ is the probability that a purchase is converted into a scraped review. Therefore, the probability of observing a review by user i for drug j at period t equals $\mathbb{P}(R_{ijt} = 1) = \pi_{jt}s_{ijt}$, where s_{ijt} is the predicted probability that consumer i purchases j in period t.

We can use this to find the expectations for particular agents. First, the expected product of total reviews for consumer i is

$$\begin{aligned}\mathbb{E}\Big[R_{ij}R_{ik} \mid i\Big] &= \mathbb{E}\Big[\big(\sum_{t=1}^{T} R_{ijt}\big)\big(\sum_{t=1}^{T} R_{ikt}\big) \mid i\Big] = \sum_{t_1,t_2}\mathbb{E}\Big[R_{ijt_1}R_{ikt_2} \mid i\Big] \\ &= \sum_{t_1\neq t_2}\pi_{jt_1}\pi_{kt_2}s_{ijt_1}s_{ikt_2} + I(j=k)\sum_{t}\pi_{jt}s_{ijt}.\end{aligned} \tag{6}$$

Second, the probability of selection into the observed sample equals

$$\mathbb{P}\big(R_i > 0 \mid i\big) = 1 - \prod_{t=1}^{T}\Big(1 - \sum_{j=1}^{J}\pi_{jt}s_{ijt}\Big). \tag{7}$$

The moments defined by equation 5 can be approximated using averages over N simulated

[26]The probability of reviewing can be different for different drug types – for example, if people make the effort to review a purchase more often for more expensive drug types.

[27]Scraping coverage fluctuated primarily because the data provider used a varying number of active scraping bots. See Section 2.4 for a discussion of our review data.

consumers:[28]

$$\mathbb{E}\Big[R_{ij}R_{ik} \mid R_i > 0\Big] \approx \frac{\frac{1}{N}\sum_{i=1}^{N}\mathbb{E}\big[R_{ij}R_{ik} \mid i\big]}{\frac{1}{N}\sum_{i=1}^{N}\mathbb{P}\big(R_i > 0 \mid i\big)}. \tag{8}$$

Finally, the probability of conversion into a review can be estimated as the ratio of reviews to total sales:

$$\hat{\pi}_{jt} = \frac{R_{jt}}{N \times s_{jt}}, \tag{9}$$

where R_{jt} is the total number of observed reviews for day t and type j, N is the total market size, and s_{jt} is the market share of drug j in period t respectively.[29]

4.1.2. Procedure

We split time into $T = 31$ discrete periods, where each period corresponds to one of the days when listings were scraped. Because scraping happened at varying frequencies, our time periods are of varying lengths. In Appendix C.2, we present the details of how we apply equations 6 and 7 to this case.

We estimate our model in two stages. In the first stage, we find non-linear parameters (Σ) using review data and the aggregate price-quantity data from listings. The non-linear parameters will be identified by matching the observed micro-level moments. In the second stage, which coincides with the standard BLP procedure, we estimate the linear parameters of the model (α, β) using IV restrictions and the aggregate price-quantity data.

It is convenient to express indirect utility as $U_{ijct} = \delta_{jct} + \mu_{ijct} + \varepsilon_{ijct}$, where $\delta_{jt} = -\alpha p_{jct} + x_{jct}\beta + \xi_{jct}$ is the mean utility, which is the component that is common for all consumers in a particular market, and $\mu_{ijt} = \sum_{g\in G}\lambda_i^g I(j \in g)$, which is the consumer-specific component. We make the conventional assumption that ε_{ijct} are from the standard Gumbel distribution. For a fixed pair of δ_{jct} and μ_{ijct}, the probability that a consumer i purchases product j equals

$$s_{ijct} = \frac{\exp(\delta_{jct} + \mu_{ijct})}{1 + \sum_{k=1}^{J}\exp(\delta_{kct} + \mu_{ikct})}. \tag{10}$$

Each value of Σ defines a distribution $F(\mu|\Sigma)$ of idiosyncratic utilities. The predicted share

[28] For simplicity, we omit the city index c here. In practice, we sample agents from each market proportionally to the market size N_c and assume that each agent faces prices from the same city across all periods.

[29] By choosing π_{jt} this way, we guarantee that simpler moment conditions like $\mathbb{E}_i R_{ijt} = R_{jt}/N$ are satisfied, and our micro moments target not levels but *comovements* of observed reviews.

of consumers in city c who purchase product j in period t equals

$$s_{jct}(\Sigma) = \int \frac{\exp(\delta_{jct} + \mu_{ijct})}{1 + \sum_{k=1}^{J} \exp(\delta_{kct} + \mu_{ikct})} dF(\mu|\Sigma) \tag{11}$$

and can be approximated by numerical integration.[30] The system defined by equation 11 can be inverted (Berry et al., 1995); that is, values $\delta_{jct}(\Sigma)$ can be found such that predicted market shares match the observed market shares. We then can find the choice probabilities $s_{ijct}(\Sigma)$ for each simulated agent and calculate predicted moments given by Equation 5. Using gradient descent, we find non-linear parameters Σ such that predicted moments match the moments observed in the data. Our demand estimation procedure, outlined in Algorithm 1, is similar to the procedure described in Conlon and Gortmaker (2023) for survey-type micro moments. Our code is based on the PyBLP package (Conlon and Gortmaker, 2020).[31]

Algorithm 1 Estimation of nested logit with micro moments.

Sample N agents with nodes ν_i. Iterate until convergence in Σ:

1. Calculate $\mu_{ijct}(\Sigma)$.
2. Find $\delta_{jct}(\Sigma)$ such that predicted market shares equal observed market shares.
3. Using $\delta_{jct}(\Sigma)$ and $\mu_{ijct}(\Sigma)$, compute predicted values $\mathbb{E}\big[R_{ij}R_{ik} \mid i\big]$ and $\mathbb{P}\big(R_i > 0 \mid i\big)$ for each agent i.
4. Estimate micro moments $g^M(\theta) = \Big(\big(\frac{1}{N}\sum_i \mathbb{E}\big[R_{ij}R_{ik} \mid i\big]\big)/\big(\frac{1}{N}\sum_i \mathbb{P}\big(R_i > 0 \mid i\big)\big) - \overline{R_{ij}R_{jk}}\Big)_{(j,k)\in P}$ for a set of product pairs P.
5. Update Σ by minimizing error function $g(\Sigma)'Wg(\Sigma)$.

Recover linear parameters (α, β) from regression of $\delta_{jct}(\Sigma)$ on x_{jct}, p_{jct} using a collection of instrumental variables Z.

In Appendix Section C.1, we show that, compared to the standard BLP estimation procedure, our micro moments substantially improve estimation precision in test simulations. We use the optimal weighting matrix W_B for BLP moments. We use diagonal weighting (Altonji and Segal, 1996), with a matrix W_M that scales each moment by the variance predicted by the logit model. We simulate $N = 100,000$ agents; the number of agents from each city is proportional to the corresponding market size. In Appendix Section C.3, we provide analytical expressions for gradients that enable a substantial reduction in estimation time.

In the set of product pairs P for constructing our micro moments we include all pairs

[30] We use 1,000 Halton draws for numerical integration.

[31] We are extremely grateful to Jeff Gortmaker, who provided helpful suggestions about using PyBLP for our study.

of drugs for which we have random coefficients. Because we use a common dummy for hashish and marijuana, we structure our moments for cannabis in the same way, considering products of reviews for cannabis and reviews for other drugs in the moments. Thus, we have $K(K+1)/2$ moments, and the parameters of the model are exactly identified.

We include several different variables in the set of instruments Z. First, we use prices in other geographic markets in the same period (Hausman et al., 1994; Nevo, 2001). Second, we use the "differentiation IVs" of Gandhi and Houde (2019), which measure the extent to which observed product characteristics distinguish each product from others in the market. Third, we use several instruments that measure the degree of competition in each market, such as the number of listings and the number of shops. Fourth, we use distance to the nearest port, motivated by the fact that some drugs are available only from abroad and the cost of within-country transportation increases with distance from the point of entry (see the discussion in Section 2.3).[32]

Because we allow for an outside option, we need to define the total market size to calculate market shares. In Section 3.3, we discuss our assumption that the number of listings is proportional to the number of transactions, and we provide supporting evidence. To define the market size in terms of transactions, we need to calculate the coefficient of proportionality between listings and transactions. Using the change in shop-level cumulative deals over the observed period, we find that the ratio of daily transactions to listings is approximately equal to 0.7. We assume that each person between 18 and 45 can consume drugs 1 time per month and that 1 standard purchase is enough to consume drugs 3 times.[33] This is motivated by the data on mortality causes in Russia, in which we find that the majority of deaths associated with drug consumption are of individuals aged between 18 and 45. We present the details of our definition of the market size in Appendix Section D.

Discussion. Our identification of Σ is based on micro moments. Theoretically, the mixed logit model can be estimated using only aggregate data. However, in our case, Σ has a particularly large dimensionality and $K(K+1)/2 = 21$ parameters to estimate. This implies two restrictive requirements for the estimation of the model from aggregate data. First, we would need considerable variation in the data to have enough statistical power. Second, and perhaps more restrictive, we would need a large number of excluded instruments that shift the prices of individual drug types.[34] One way to simplify the estimation of Σ is to impose

[32] Hansen et al. (2023) find that legal seaborne trade flows increase availability and deaths from fentanyl.

[33] This is similar to the definition used by Hollenbeck and Uetake (2021), who assume that each resident of a market can purchase 4 grams of legal marijuana per month.

[34] This is particularly challenging in the context of illegal drugs, where some of the traditionally used instruments, such as taxes, tariffs, and firm-level costs, are not available because of the illegal nature of the market.

additional restrictions – for example, by assuming that random coefficients are uncorrelated. This is undesirable because our empirical analysis of reviews suggests that allowing for a correlation between preferences for particular drug types is crucial, and the model is likely to predict incorrect substitution patterns if it does not account for it.

Our method is related to other studies that have addressed the identification of non-linear parameters in BLP-type models by utilizing additional micro-level data when they are available (Chintagunta and Dubé, 2005; Bayer et al., 2007). In particular, micro-level data can be incorporated as additional moment conditions – for example, when describing the relationship between the choices and demographics of consumers (Petrin, 2002). Our approach is closest to estimation using second-choice data (Berry et al., 2004; Conlon et al., 2021; Conlon and Gortmaker, 2023). In this approach, two products, j and k, are inferred to be close substitutes if consumers purchase k when j is not available. This is similar to our identification procedure, which infers that j and k are close substitutes if the same consumers review them in different time periods.

We propose a new method to estimate non-linear parameters in BLP-type models. Our method identifies non-linear parameters using correlations in choices across time, where choices are inferred from irregularly observed reviews. This approach can be particularly useful in the study of online marketplaces because review data can often be collected from these platforms at a small cost. It can be an alternative to second-choice data, which is often unavailable.

4.2. Estimates

Table 5 presents point estimates and standard errors for the linear parameters.[35] For purposes of comparison, the first column provides the estimates from the standard logit model. We find a negative relationship between demand and price. The implied price elasticities are discussed below. The estimates for other linear coefficients have interpretable parameters. In particular, we find that consumers prefer the hidden and magnet delivery methods over the third hiding method, which is burying the drugs in the soil. This preference can be explained by the fact that retrieving dead-drops from soil is riskier and less convenient. It is also consistent with how shops advertised their goods on Hydra, as can be seen in Appendix Figure F.4. We find that consumers had a preference for drugs listed as "very high quality."[36] However, we also find that reputation measures, such as review sentiment

[35]The current estimates for standard errors do not account for randomness from the first stage of estimation. In the next version of the paper, we intend to provide estimates of standard errors obtained through bootstrapping.

[36]The platform introduced quality labels, VHQ (very high quality) and HQ (high quality), which represented substance purity levels above 98% and 95%, respectively. Although shops self-reported these labels,

Table 5: Coefficient estimates with standard errors

	Logit		Mixed logit	
	Coef.	S.E.	Coef.	S.E.
Price	-0.219	0.008	-0.243	0.032
Magnet	1.240	0.097	0.775	0.251
Hidden	0.548	0.072	0.799	0.360
Crystal form	2.417	0.122	0.148	1.143
Very high quality	0.445	0.087	0.600	0.176
High quality	-0.492	0.085	-0.160	0.260
Product rating	0.083	0.029	0.149	0.045
Reviews sentiment	0.156	0.027	0.049	0.056
Shop age	-0.024	0.003	-0.023	0.010
Shop rating	0.036	0.015	0.034	0.027
"Trusted seller"	0.710	0.065	0.680	0.232
2 doses	-1.013	0.138	-1.433	0.347
3 doses	-1.783	0.164	-2.082	0.530
4 doses	-1.776	0.141	-2.178	0.474
FE	Date		Date	
Markets	1,054		1,054	
Obs.	12,203		12,203	

and product and shop ratings, had a relatively small effect on utility. At the same time, the label of a trusted shop, which could be purchased by any shop that met a set of criteria, had a substantial positive effect on utility. Finally, because our demand model considers price per gram, we find that consumers preferred smaller dead-drops, other things being equal. This can be rationalized by buyers facing budget constraints or incurring inventory holding costs. It can also be explained by the risks associated with purchasing larger quantities of drugs, as discussed in Section 2.4.

Figure 2 shows the matrix of covariances between random coefficients. Our findings are consistent with the evidence from reviewing behavior, which is discussed in Section 3.5. First, we find substantive variances for each random coefficient, which corresponds to our finding that consumers most often review the same drug in different periods. Second, we find relatively higher variance for amphetamine, mephedrone, $\alpha-$PVP, and cocaine. We find lower variance for MDMA and cannabis, which is consistent with both smaller attachment found in review data and lower dependence indexes for these drugs.

Figure 3 shows estimated correlations between random coefficients. Consistent with our discussion in Section 3.5, we find that taste shocks are positively correlated for three

Hydra supposedly conducted random tests to ensure that the drugs sold met these standards (Goonetilleke et al., 2023).

Figure 2. Estimated covariances of random coefficients: $\Sigma^T \Sigma$

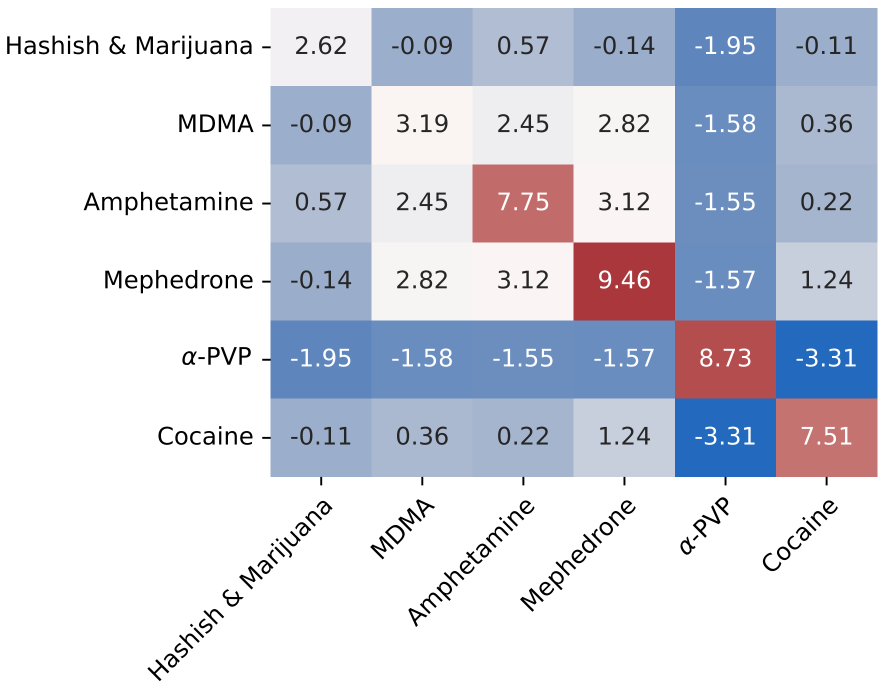

Figure 3. Estimated correlations of random coefficients: $corr(\lambda^j, \lambda^k)$

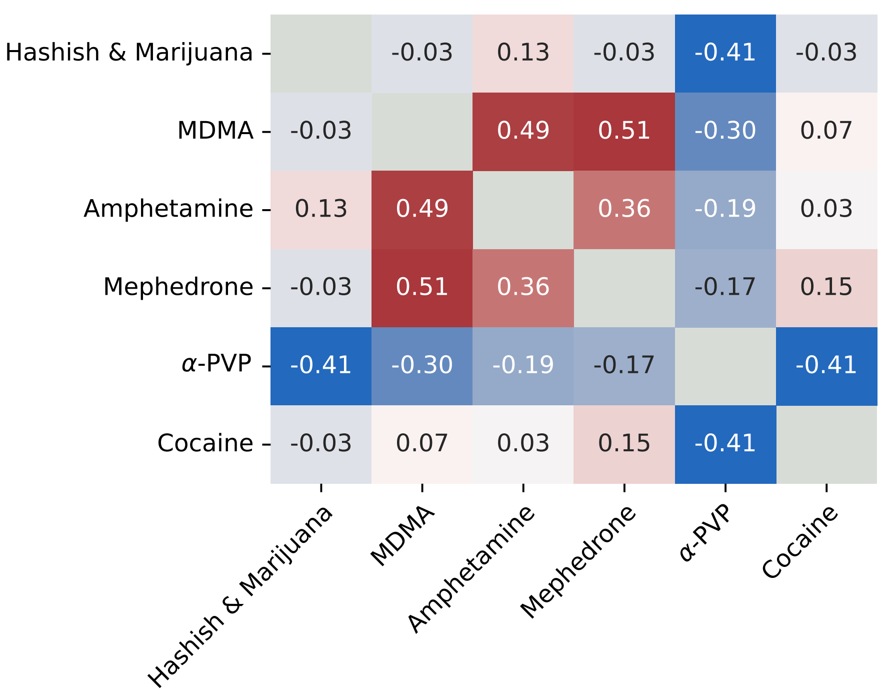

Figure 4. Median cross-price elasticities of demand

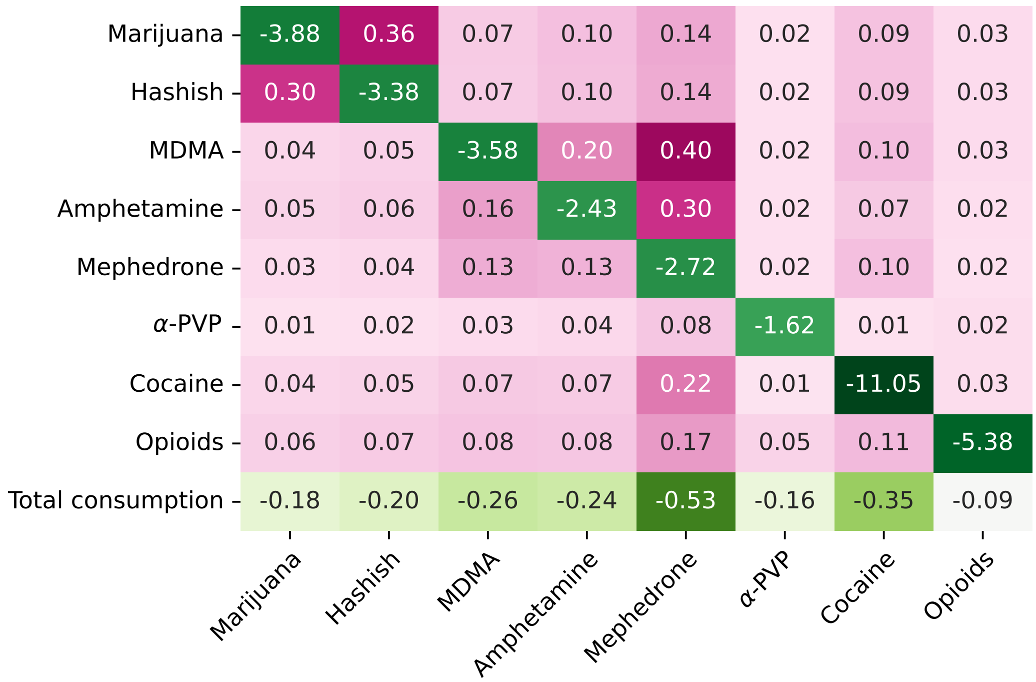

Note: Element $\mathcal{E}_{j,k}$ of this matrix presents the median of cross-price elasticities $\varepsilon_{jkct} = \frac{p_{jct}}{s_{jct}} \frac{\partial s_{jct}}{\partial p_{kct}}$ across all markets (c, t).

stimulants: amphetamine, MDMA, and mephedrone. At the same time, we see a negative correlation between the random coefficient for $\alpha-$PVP and the random coefficients for all other drugs. This is in line with the observation that consumers who reviewed $\alpha-$PVP rarely review any other drug type.

Figure 4 presents median cross-price elasticities in all markets, where a market is defined as a city-period combination. Appendix Figure E.1 shows the distributions of own-price elasticities for the eight most popular drug types. Several factors determine the scale of elasticity. First, products with many close substitutes should have more elastic demand. Second, products with high attachment (variance of the corresponding random coefficient) should have less elastic demand. Finally, models built on the logit framework tend to predict higher elasticity for more expensive products. We find the lowest own-price elasticity for $\alpha-$PVP, which can be explained by a combination of its high attachment, the relatively low price for this drug, and the fact that it has no close substitutes. We find the highest own-price elasticity for cocaine, which probably can be explained by its high price.

From their meta-analysis, Gallet (2014) conclude that the median price elasticity obtained in studies on demand for drugs is -0.33, which is lower than the own-price elasticities we

Figure 5. Actual and predicted consumption of hashish

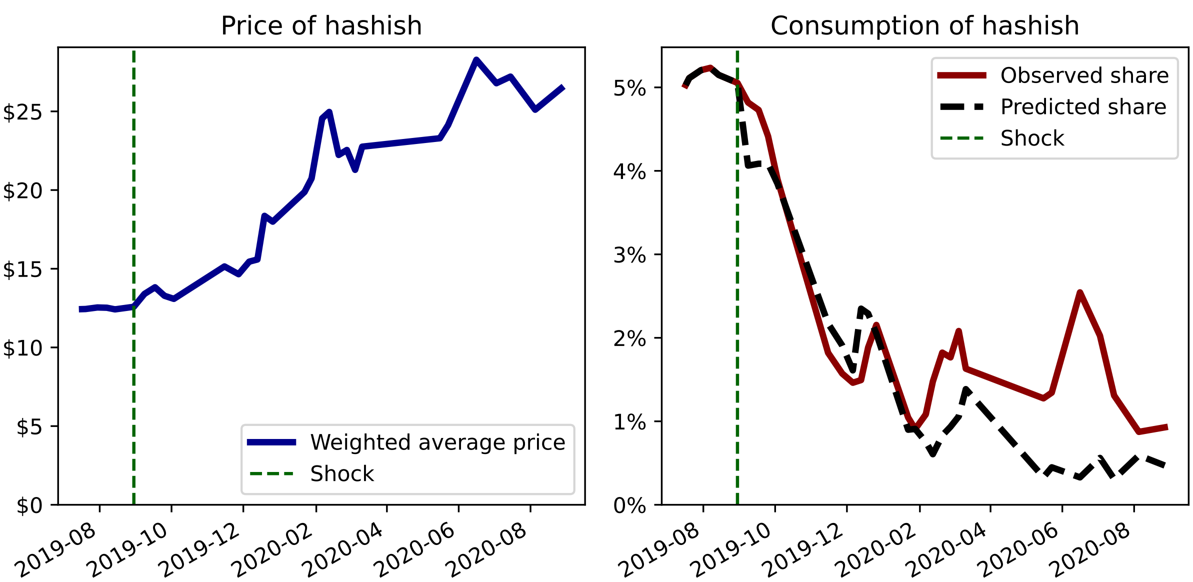

Note: Price is the average across all cities in the sample, where weights are proportional to the market size. Prices are converted to USD using an exchange rate of 74 RUB per USD.

obtain. However, as discussed in our literature review, previous studies typically relied on crude proxies for drug consumption and lower-quality price data. Our estimates are close to those of Miravete et al. (2018), who find an average elasticity of demand for hard liquor of −2.8. Moreover, we are able to validate our estimates for price sensitivity by examining the effects of an exogenous supply shock, as discussed in Section 4.3.

Estimates for cross-price elasticities reflect our findings in Sections 3.5 and 4.2. In particular, an increase in the price of hashish is predicted to have the greatest impact on the demand for marijuana, and vice versa. A similar substitution pattern is characteristic of the triad of related stimulants: amphetamine, MDMA, and mephedrone. An increase in the price of one of them has the largest effect on demand for the other two. For example, an increase in the price of mephedrone has the largest effect on demand for MDMA and amphetamine.

4.3. Validation: hashish shock

In 2019, increased enforcement targeting the production and trafficking of Moroccan hashish substantially decreased the supply of this drug in the European markets (EMCDDA and Europol, 2020). This coincided in time with a major hashish seizure within Russia

(TASS, 2019). These shocks were followed by a significant increase in the price of hashish on the Russian drug market (FilterMag, 2020). Because this price change can be attributed to particular shocks of supply and, therefore, is not likely to result from a shift in demand, we use it to validate our model estimates.

Figure 5 shows predicted demand for hashish given the observed prices on the market, where values of ξ_{jct} are fixed at the average over the period preceding the shock for each product-market combination. Our model closely predicts the decline in hashish consumption for the first four months of the price increase. Over time, the quality of the prediction declines, which can be explained by the effect of demand-side shocks accumulating over time, which is not accounted for in our exercise. However, even over the long run, the fit of our prediction is reasonably good.

5. Effects of Drug Policies

We use our model of the demand for illegal drugs to assess the effects of drug policies and account for substitutions between drug types.

5.1. Cannabis legalization

Cannabis legalization, one of the most discussed drug policies, recently has been adopted by various jurisdictions in the U.S. and around the world.[37] While legalization stems from various motivations, including reducing incarceration and policing costs, one widely discussed aspect is its potential impact on the use of other drugs. Previous studies, driven by high opioid mortality and the potential for marijuana to serve as a substitute for prescription opioids in chronic pain treatment, have primarily focused on the effect of legalization on opioid consumption.[38] State-level event studies provide mixed evidence of the impact of legalization on opioid use. Bachhuber et al. (2014), Powell et al. (2018), and Chan et al. (2020) have reported that legalization has led to a reduction in opioid overdoses. However, Drake et al. (2021) found only a short-term effect, while Shover et al. (2019) argue that the association between legalization and opioid overdoses became positive over time.

This aligns with our analysis of drug reviews in Section 3.5, where we find that consumers who purchase opioids rarely review cannabis and other drugs. Thus, cannabis is unlikely to

[37]The extent of legalization can vary; some jurisdictions have legalized only the medical use of marijuana, while others have also legalized recreational use. In the U.S., the first instance of medical use legalization was in California in 1996, while Colorado and Washington were the first states to legalize recreational marijuana in 2012. As of October 2023, recreational cannabis has been legalized in 23 U.S. states and the federal District of Columbia.

[38]Hurd et al. (2019) suggest that cannabis also can alleviate the symptoms of opioid use disorder.

Figure 6. Cannabis legalization and predicted drug use

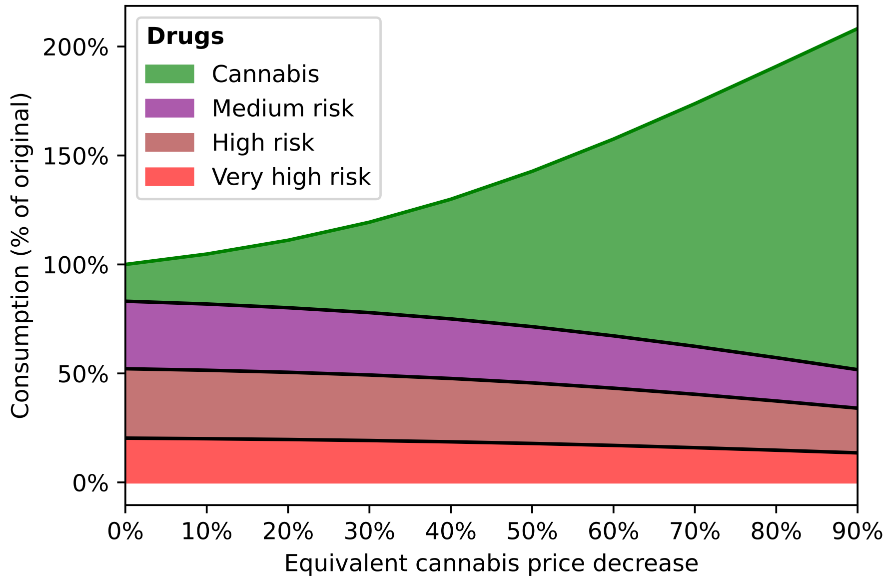

Note: This plot shows predicted consumption of different drug groups when the price of cannabis (hashish and marijuana) decreases by $x\%$. Very high-risk drugs include $\alpha-$PVP, cocaine, and opioids; high-risk drugs include mephedrone and stimulants not listed as medium-risk drugs; medium-risk drugs include amphetamine, GHB, MDMA, LSD, and other psychedelics.

function as a substitute for opioids, and legalization has a low potential to decrease their use. However, our analysis suggests a larger substitutability between cannabis and other drugs, particularly amphetamine, cocaine, and MDMA. Consequentially, legalization might have the potential to reduce the consumption of these drug types. Given that the risks associated with cannabis are likely to be smaller than those associated with other drugs, this relationship, if it can be demonstrated, could serve as a significant motivation for legalization.

However, this benefit should be weighed against the potential increase in cannabis consumption. We apply our model to quantify the trade-off between the consumption of cannabis and the consumption of other drugs. We also ask whether, after legalization, users substitute cannabis for drug types that cause large or small harm.

We assume that legalization induces the same substitution patterns as a reduction in the price of cannabis. This has two complementary interpretations.[39] First, legalization was found to decrease cannabis prices in the U.S. (Anderson et al., 2013) and Canada (Hall et al., 2023). The production costs of legal marijuana are known to be very low, and a price decrease can occur after legalization because sellers do not incur costs and risks associated

[39]This exercise ignores the effects of legalization that are not directly related to the demand for drugs. For example, Adda et al. (2014) and Gavrilova et al. (2019) study the effect of legalization on crime.

Table 6: Cannabis legalization and predicted consumption change

	10.0% reduction	25.0% reduction	50.0% reduction	75.0% reduction	90.0% reduction
Panel A: Change in use					
α-PVP	-0.8%	-2.5%	-7.8%	-17.4%	-25.0%
Amphetamine	-1.8%	-5.7%	-16.2%	-31.5%	-42.0%
Cocaine	-1.8%	-5.5%	-15.6%	-30.5%	-40.8%
Hashish	31.6%	99.6%	289.1%	555.0%	720.9%
Marijuana	40.3%	126.5%	357.5%	699.3%	940.1%
MDMA	-1.8%	-5.6%	-16.1%	-31.6%	-42.2%
Mephedrone	-1.2%	-3.8%	-11.7%	-24.5%	-33.9%
Non-cannabis use	-1.5%	-4.9%	-14.0%	-28.0%	-37.8%
Cannabis use	35.8%	112.7%	322.3%	625.0%	827.2%
Total use	4.8%	15.0%	42.7%	82.2%	108.2%
Panel B: Change per 1 dose decrease in use of non-cannabis drugs					
Cannabis use	4.7 doses	4.7 doses	4.7 doses	4.5 doses	4.4 doses
Total use	3.7 doses	3.7 doses	3.7 doses	3.5 doses	3.4 doses

with illegal production, transportation, and sale (Caulkins, 2010). Second, many of the effects of legalization on consumers, such as diminished risks related to purchase and the elimination of the stigma of illegality, can be considered equivalent to a reduction in price. If the effect of these factors on indirect utility from marijuana consumption is positive and homogeneous across consumers, then it is equivalent to a price reduction.

This assumption can be violated if, for example, consumers have a heterogeneous distaste for illegality. If agents who did not consume drugs before legalization have a higher distaste for illegality, the utility from consumption of cannabis will increase disproportionally more for them than for current drug consumers. In this case, our model can underestimate the increase in demand for cannabis after legalization.

In practice, policymakers often choose policies that limit the extent of the price reduction after legalization.[40] Governments could achieve further price decreases by setting lower taxes and increasing the number of licenses granted. For this reason, we consider a range of counterfactual price reductions from the current price levels of marijuana and hashish. Figure 6 shows predicted consumption of cannabis and other drugs, where we group drugs by estimated risk using the harm index developed by Nutt et al. (2007).[41] Table 6 shows

[40] Hollenbeck and Uetake (2021) show that in Washington state, the small number of retail licenses resulted in high retail margins.

[41] We include $\alpha-$PVP in the list of high-risk drugs because Patocka et al. (2020) reports substantial risk

predicted consumption for individual drug types.

We find that the government can achieve significant success in reducing the consumption of more dangerous drugs. However, this will be accompanied by a substantial increase in cannabis consumption. For instance, if the price of cannabis falls by 50%, the use of other drugs will decrease by 14%, while cannabis use will increase by 322%.

To predict the absolute effect of legalization, we would need to estimate the associated price reduction. This number depends on the government's choices and consumers' utility costs associated with illegality and thus is hard to determine. In other words, determining the number is difficult. However, Panel B of Table 6 shows that the relative change in use remains approximately consistent across a range of experiments. To achieve a one-dose reduction in the consumption of other drugs, society would need to accept 4.5 additional doses of cannabis. Therefore, the relevant policy consideration is whether the average social cost of the use of one dose of other drugs exceeds the social cost associated with the use of 4.5 doses of cannabis.

We can also observe that gains from substitution can be limited due to the fact that the primary substitutes for marijuana and hashish typically are drugs considered to be lower- to medium-risk. Figure 6 illustrates that substitution primarily occurs with the least dangerous drugs. Panel A of Table 6 indicates that an increase in the availability of cannabis should have the most significant impact on the demand for MDMA and amphetamine, while its effect on $\alpha-$PVP would be relatively smaller. Thus, the potential benefits of legalizing cannabis for the purpose of substitution are constrained because the reduction in consumption is more pronounced for drugs with medium risks rather than those with high or very high risks.

5.2. Introduction of new drugs

As discussed in Section 3.2, the two new drugs, mephedrone and α-PVP, had market shares on Hydra of 28% and 11%, respectively. These drugs fall under the category of synthetic cathinones, commonly referred to as "bath salts" (Soares et al., 2021). They can be considered part of the broader phenomenon of "legal highs," which are newly synthesized substances that mimic the effects of conventional drugs.[42] These substances typically maintained legal status for several years before governments adjusted legislation to ban them. Another prominent category of legal highs is "synthetic cannabinoids," which gained popularity in the U.K., U.S., New Zealand, and several European countries Peacock et al. (2019).

The introduction of these products to the market potentially increased total drug con-

of overdoses from this drug.

[42]While mephedrone had been known to scientists since the late 1920s, it was rediscovered by underground chemists in the late 2000s and found its way to the black market soon thereafter.

Table 7: Effect of bath salts' introduction on drug demand

Drug type	Estimated effect
Amphetamine	-17.0%
Cocaine	-11.9%
Hashish	-8.0%
LSD	-13.1%
Marijuana	-9.3%
MDMA	-22.7%
Opioids	-15.7%
Other cannabis	-15.6%
Other stimulants	-15.2%
Other psychedelics	-14.5%
Total (with bath salts)	40.8%

sumption for two reasons. First, the price of these substances might be lower than that of traditional drugs. Second, these new substances might possess characteristics different from those of "established" illegal drugs, making them attractive to some of the consumers who previously had not purchased any drugs. We are unaware of any attempts to estimate the effect of emerging drugs on drug use. Given that governments can allocate resources to prevent the discovery or adoption of new illegal drugs, this question has significant policy implications. These include faster legislative responses to ban new products and more stringent regulations governing research into new substances. Our estimates could provide insights into the potential benefits of these interventions.[43]

With a sizable 39% share, bath salts dominate the market in Russia. The causal effect of their introduction lies between two hypothetical extreme cases. In the first scenario, there is no substitution from other drug types: all people who consume bath salts in the sample period would consume no drugs if bath salts had never appeared on the market. The second scenario is perfect substitution: all people who consume bath salts in the sample period would consume some other drugs if bath salts had never appeared on the market. We apply our estimated model and simulate it under the assumption that all bath salts were eliminated from the market when our dataset was collected. The results are presented in Table 7.

We estimate that the introduction of bath salts has increased the total demand for illegal drugs by 41%. A naive calculation that ignores substitution from other drug types would

[43] Additionally, our estimates can allow other researchers to disentangle the effect of the introduction of bath salts from other factors that affect the drug market, helping them evaluate relevant policies, regulations, and other market shifts.

suggest that their introduction had a larger hypothetical effect of $1/(1-0.39)-1 \approx 64\%$. Thus, the substitution from the types that previously had existed was substantial but does not affect the main conclusion. The introduction had a large effect on total drug use and brought many new consumers to the market. This effect results from two mechanisms. First, the attachment to specific drug types, discussed above in Sections 3.5 and 4.2, limits the scope of potential substitutions between drugs. Second, we find that $\alpha-$PVP, one of the bath salts, lacks close substitutes.

Our model also enables us to estimate how the introduction of bath salts has affected the demand for specific preexisting drugs. As is shown in Table 7, the drugs most significantly impacted are MDMA and amphetamine, which are the closest substitutes of mephedrone. Their consumption is 23% and 17% lower relative to what is predicted in the counterfactual scenario when there is no competition from bath salts. Conversely, the drugs least affected are hashish and marijuana. For other drugs, the effect is approximately -15%, but our ability to estimate substitution from them is limited because we do not include random coefficients for these drug types.

It would be difficult to forecast changes in the consumption of illegal drugs if new synthetic drugs were introduced in the future. The effect of such an introduction would depend on the price of new drugs and the substitutability between them and existing drug types. However, our estimates suggest that the effect of new products can be dramatic, and governments should allocate resources to prevent the emergence in the future of new drugs.

5.3. Drug elimination

Supply-side interventions can increase the consumption of other drugs if different types of illegal drugs are substitutes. Manski et al. (2001) hypothesized that this could offset the benefits of reducing the availability of the targeted drug. Moreover, such substitution may be towards drug types more dangerous than the targeted one. For example, Alpert et al. (2018) and Evans et al. (2019) found that mortality from heroin drastically increased after a supply-side intervention changed the formulation of Oxycontin, as individuals addicted to Oxycontin switched to heroin. This raises the question: What are the effects of drug reduction policies given potential substitutions between drugs? We examine how the demand for illegal drugs would be impacted if a particular drug were to be eliminated. We conceptualize this scenario as an extreme case of a successful targeted intervention by the government. Specific interventions may have effects that are close to complete elimination, at least in the short run, as seen in examples such as crackdowns on heroin in Australia (Moore and Schnepel, 2021) and on methamphetamine in the U.S. (Dobkin and Nicosia, 2009).

Figure 7. Diversion ratios for most popular drug types

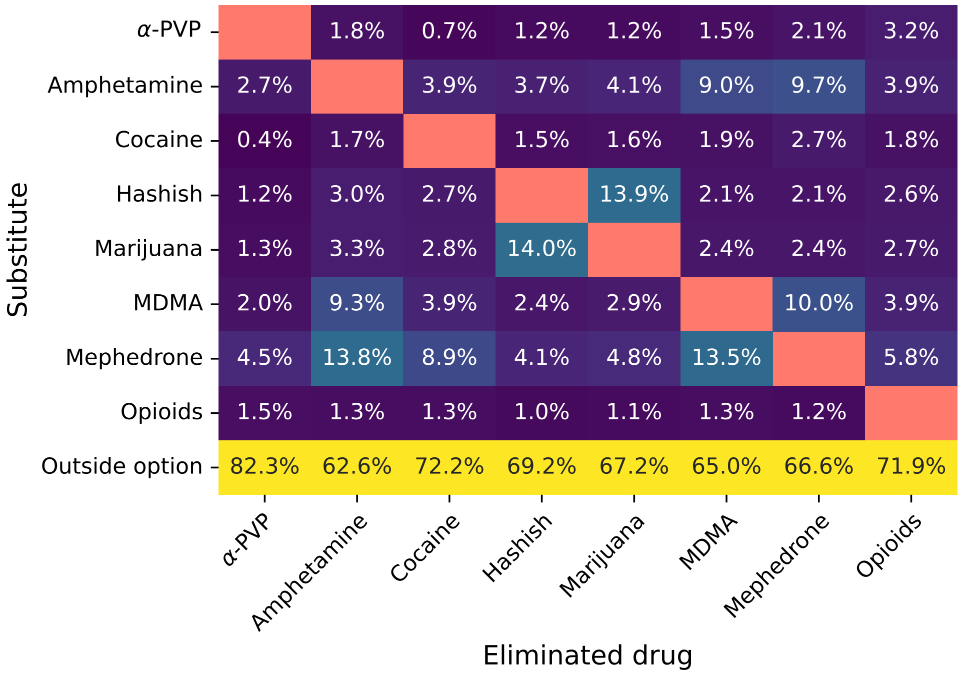

Note: Diversion ratios are defined as $D_{jk} = (s_j^{-k} - s_j)/s_k$, where s_j^{-k} is the counterfactual market share of product j when product k is deleted from all choice sets.

Figure 7 presents diversion ratios resulting from elimination, which indicate the fraction of drug k's consumption that would shift to each of the potential substitutes (including the outside option). Our findings here largely align with our Section 3.5 and Section 4.2 discussions. For instance, following the elimination of hashish, consumers will switch to marijuana more than to any other drug, and vice versa. Similarly, if amphetamine, MDMA, or mephedrone were eliminated, consumers would disproportionately transition to the remaining two drugs. Consequently, our findings suggest that the impact on total consumption is least significant for drugs that have close substitutes, namely amphetamine, MDMA, mephedrone, hashish, and marijuana.

We also observe the most substantial effects for drugs with no close substitutes, such as $\alpha-$PVP and cocaine. For example, our model predicts that after the elimination of α-PVP, only 18% of its consumers will switch to another drug type. By the same measure, enforcement would be half as effective for amphetamine because 38% of consumers would find another drug to which to switch. Our findings suggest that all else being equal, the government should prioritize the enforcement of drugs with few close substitutes. Finally, while we do not have a random coefficient for opioids in our demand model, we find a highly

Table 8: Elasticity of revenue with respect to drug prices

	Own revenue	Total revenue (w/o substitution)	Total revenue (with substitution)
α-PVP	-0.577	-0.036	-0.006
Amphetamine	-1.317	-0.160	-0.038
Cocaine	-8.547	-1.522	-1.382
Hashish	-2.693	-0.185	-0.104
Marijuana	-2.959	-0.230	-0.136
MDMA	-2.417	-0.224	-0.105
Mephedrone	-1.418	-0.410	-0.206
Opioids	-3.980	-0.121	-0.085

strong attachment to these drugs in Section 3.5. We can reasonably assume that our model overestimates substitution from opioids, and therefore, our conclusion also applies to them.

5.4. Enforcement and revenue

The model developed in Becker et al. (2006) highlights the key role of price elasticities in determining the effects of drug enforcement.[44] Supply-side interventions, such as seizures or crop eradication, increase the price of the targeted drug type. However, if the demand for this drug is inelastic, enforcement will lead to a decrease in consumption that is relatively small compared to the price increase. This poses a difficult dilemma for the government: the black market's total revenue can increase due to drug enforcement. Higher revenue in drug markets can lead to increased resources devoted to drug smuggling or greater incentives to fight for control over the drug trade (Becker et al., 2006; Castillo et al., 2020). Even if the goal of reducing consumption is achieved, society will face a larger scope of associated illegal activities, including gang wars, officials' corruption, and attacks on journalists and civil activists.

The elasticity of revenue for product j with respect to p_j equals $1 + \varepsilon_{jj}$, where ε_{jj} is its own-price elasticity. This motivates an approach popular in the literature on the demand for illegal drugs, where the elasticity of demand for a particular drug is estimated and compared with -1 (see Gallet, 2014 for a review). However, this approach ignores the possibility of substitution. The revenue of the black market from other drugs increases if people who stop consuming the targeted drug do not leave the market but, instead, consume a substitute

[44] See White and Luksetich (1983) for an earlier discussion of this idea. They also suggest that demand elasticity is crucial for determining the effect of enforcement targeting drug users.

drug. The elasticity of the total revenue of the black market for drugs with respect to the price of drug j equals

$$\frac{p_j}{\sum_{k=1}^{J} p_k s_k} \frac{\partial \sum_{k=1}^{J} p_k s_k}{\partial p_j} = s_j^r + \sum_{k=1}^{J} s_k^r \varepsilon_{kj} = s_j^r \underbrace{\left(1 + \varepsilon_{jj}\right)}_{\text{Own revenue}} + \underbrace{\sum_{k \neq j}^{J} s_k^r \varepsilon_{kj}}_{\text{Substitution}}, \tag{12}$$

where $s_j^r = p_j s_j / \left(\sum_{k=1}^{J} p_k s_k\right)$ is the revenue share of product j. Thus, even if the own-price elasticity ε_{jj} is below -1, total revenue can increase from enforcement. The correct assessment of this possibility requires estimates not only of the own-price elasticity but also of the cross-price elasticities of demand for different drug types.

We apply our model estimates to evaluate how drug-specific enforcement affects total revenue. Table 8 reports the revenue elasticities for the most popular drug types.[45] In the first column, we report the effect on revenue from sales of the targeted drug. This is the number typically used in papers that estimate demand for a single drug type. In the second column, we apply the formula for total revenue without the second term to highlight the effect of substitution in our estimates, which would be the effect on total revenue if all $\varepsilon_{jk} = 0$ for $j \neq k$. In the third column, we report the elasticity of total revenue.

Our estimates suggest that enforcement does not increase revenue for any major drug type. However, our findings indicate that enforcement actions against $\alpha-$PVP cause only a minimal decrease in revenue for drug sellers due to its low own-price elasticity. At the same time, our estimates show that the effect on total revenue can be miscalculated for types with close substitutes if researchers ignore substitution. For example, when considering potential substitution, we find that enforcement of amphetamine, which has two close substitutes (MDMA and mephedrone), is almost revenue-neutral.

5.5. External validity

The external validity of our analysis might have several limitations due to the differences between the market for drugs in Russia and those in other countries. First, the composition of drugs in the consumption bundle might be different. In particular, in the U.S. and Europe, bath salts are significantly less popular than they are in Russia. Moreover, sales of fentanyl were prohibited on the Hydra marketplace, and thus, our analysis is not informative about demand for this drug. Second, darknet markets have a relatively small market share in the U.S., and most of the trade happens through offline dealers. Search costs are likely to be much larger for consumers buying drugs on the street. Therefore, sellers in the U.S.

[45] We use a version of this formula for many markets and report the elasticity of total revenue.

might have more market power than those who operated on Hydra. Moreover, an online marketplace can lead to more substitution because buying drugs of different types is easier there. Third, the U.S. population generally has higher incomes, potentially lowering price elasticity among American consumers.

6. Conclusion

We analyze the market for illegal drugs utilizing data scraped from Hydra, the world's largest darknet marketplace to date. This dataset enabled us to estimate a model of demand for a wide range of illegal drugs and study substitution between them. To identify consumer preferences, we employed a novel approach based on micro-level moment conditions that capture inter-temporal correlations in individual choices. Our findings reveal significant variation in the level of attachment to different drugs and substitutability between them. Several substances demonstrate close substitutability: the three medium-risk stimulants and the two types of cannabis. We employ our model to evaluate the effects of key drug policies that affect the supply of illegal drugs. The legalization of cannabis can achieve a decrease in the use of riskier drugs, but such a decrease will be accompanied by a substantial increase in cannabis consumption. Governments should proactively seek to prevent the introduction of new drugs into the market because recent introductions of new drugs, such as bath salts, have had pronounced effects on overall drug consumption. Drug enforcement is likely to be more successful when it targets drugs with few substitutes.

We foresee several important directions for future research. First, our paper models consumer preferences as static. A valuable extension of this framework might involve a demand model in which preferences can change over time, particularly in the case of accumulating addiction. In particular, such a model would allow us to separately study the short-term and long-term effects of drug policies.[46] Second, our discussion focuses on the demand for drugs and abstracts away from the supply side, effectively assuming that the market is competitive and the supply of drugs is perfectly elastic. For instance, this assumption might be violated if some drug sellers possess market power. While we find a high degree of competition between retail sellers on Hydra, there may be less competition between upstream suppliers. This assumption is also less realistic in the context of the traditional drug trade, where search costs should be higher than in an online platform. A model that incorporated endogenous supply would allow us to relax this assumption or study interventions that target particular sellers.

[46] See, in particular, Becker and Murphy (1988) and Gruber and Köszegi (2001). See Hui (2023) for a review of the recent economic studies incorporating addiction.

References

Adda, J., B. McConnell, and I. Rasul (2014) "Crime and the depenalization of cannabis possession: Evidence from a policing experiment," *Journal of Political Economy*, Vol. 122, No. 5, pp. 1130–1202.

Alpert, A., D. Powell, and R. L. Pacula (2018) "Supply-side drug policy in the presence of substitutes: Evidence from the introduction of abuse-deterrent opioids," *American Economic Journal: Economic Policy*, Vol. 10, No. 4, pp. 1–35.

Altonji, J. G. and L. M. Segal (1996) "Small-sample bias in GMM estimation of covariance structures," *Journal of Business & Economic Statistics*, Vol. 14, No. 3, pp. 353–366.

Anderson, M., B. Hansen, and D. I. Rees (2013) "Medical marijuana laws, traffic fatalities, and alcohol consumption," *The Journal of Law and Economics*, Vol. 56, No. 2, pp. 333–369.

Angrist, J. D. and A. D. Kugler (2008) "Rural windfall or a new resource curse? Coca, income, and civil conflict in Colombia," *The Review of Economics and Statistics*, Vol. 90, No. 2, pp. 191–215.

Bachhuber, M. A., B. Saloner, C. O. Cunningham, and C. L. Barry (2014) "Medical Cannabis Laws and Opioid Analgesic Overdose Mortality in the United States, 1999-2010," *JAMA Internal Medicine*, Vol. 174, No. 10, pp. 1668–1673.

Bayer, P., F. Ferreira, and R. McMillan (2007) "A unified framework for measuring preferences for schools and neighborhoods," *Journal of Political Economy*, Vol. 115, No. 4, pp. 588–638.

Baza.io (2020) "Euphoria Cargo," https://cargo.baza.io, [Accessed 18-Jun-2023].

Becker, G. S. and K. M. Murphy (1988) "A theory of rational addiction," *Journal of Political Economy*, Vol. 96, No. 4, pp. 675–700.

Becker, G. S., K. M. Murphy, and M. Grossman (2006) "The market for illegal goods: the case of drugs," *Journal of Political Economy*, Vol. 114, No. 1, pp. 38–60.

Berry, S., J. Levinsohn, and A. Pakes (1995) "Automobile prices in market equilibrium," *Econometrica: Journal of the Econometric Society*, pp. 841–890.

——— (2004) "Differentiated products demand systems from a combination of micro and macro data: The new car market," *Journal of Political Economy*, Vol. 112, No. 1, pp. 68–105.

Berry, S. T. (1994) "Estimating discrete-choice models of product differentiation," *The RAND Journal of Economics*, pp. 242–262.

BHASKAR, V., R. LINACRE, AND S. MACHIN (2019) "The economic functioning of online drugs markets," *Journal of Economic Behavior & Organization*, Vol. 159, pp. 426–441.

BOWDEN-JONES, O. AND D. ABDULRAHIM (2020) *What Are Club Drugs and NPS and Why Are They Important?*, p. 3–4: Cambridge University Press, DOI: 10.1017/9781911623106.002.

CARSON, E. A. (2021) "Prisoners in 2020–Statistical tables," *National Criminal Justice Reference Service*, Vol. 302776, pp. 1–50.

CASTILLO, J. C., D. MEJÍA, AND P. RESTREPO (2020) "Scarcity without leviathan: The violent effects of cocaine supply shortages in the mexican drug war," *Review of Economics and Statistics*, Vol. 102, No. 2, pp. 269–286.

CAULKINS, J. P. (2001) "Drug prices and emergency department mentions for cocaine and heroin," *American Journal of Public Health*, Vol. 91, No. 9, pp. 1446–1448.

——— (2010) "Estimated cost of production for legalized cannabis," *RAND Drug Policy Research Centre: Santa Monica, CA, USA*.

CAULKINS, J. P., P. REUTER, AND C. COULSON (2011) "Basing drug scheduling decisions on scientific ranking of harmfulness: false promise from false premises," *Addiction*, Vol. 106, No. 11, pp. 1886–1890.

CDC, NCHS (2022) "U.S. Overdose Deaths In 2021 Increased Half as Much as in 2020 – But Are Still Up 15%," https://www.cdc.gov/nchs/pressroom/nchs_press_releases/2022/202205.htm.

——— (2023) "Provisional Drug Overdose Death Counts," https://www.cdc.gov/nchs/nvss/vsrr/drug-overdose-data.htm.

ČERVENỲ, J. AND J. C. VAN OURS (2019) "Cannabis prices on the dark web," *European Economic Review*, Vol. 120, p. 103306.

CHAN, N. W., J. BURKHARDT, AND M. FLYR (2020) "The effects of recreational marijuana legalization and dispensing on opioid mortality," *Economic Inquiry*, Vol. 58, No. 2, pp. 589–606.

CHINTAGUNTA, P. K. AND J.-P. DUBÉ (2005) "Estimating a stockkeeping-unit-level brand choice model that combines household panel data and store data," *Journal of Marketing Research*, Vol. 42, No. 3, pp. 368–379.

CHRISTIN, N. (2022) "Measuring And Analyzing Online Anonymous ('DARKNET') Marketplaces,"Technical report, Carnegie Mellon University.

Conlon, C., J. Mortimer, and P. Sarkis (2021) "Estimating preferences and substitution patterns from second choice data alone,"Technical report, Working Paper at IIOC.

Conlon, C. and J. Gortmaker (2020) "Best practices for differentiated products demand estimation with PyBLP," *The RAND Journal of Economics*, Vol. 51, No. 4, pp. 1108–1161.

——— (2023) "Incorporating micro data into differentiated products demand estimation with PyBLP," Working paper.

Culotta, A., G. Z. Jin, Y. Sun, and L. Wagman (2022) "Safety Reviews on Airbnb: An Information Tale."

Dave, D. (2006) "The effects of cocaine and heroin price on drug-related emergency department visits," *Journal of Health Economics*, Vol. 25, No. 2, pp. 311–333.

——— (2008) "Illicit drug use among arrestees, prices and policy," *Journal of Urban Economics*, Vol. 63, No. 2, pp. 694–714.

Davis, A. J., K. R. Geisler, and M. W. Nichols (2016) "The price elasticity of marijuana demand: Evidence from crowd-sourced transaction data," *Empirical Economics*, Vol. 50, No. 4, pp. 1171–1192.

Dell, M. (2015) "Trafficking networks and the Mexican drug war," *American Economic Review*, Vol. 105, No. 6, pp. 1738–1779.

DeSimone, J. and M. C. Farrelly (2003) "Price and enforcement effects on cocaine and marijuana demand," *Economic Inquiry*, Vol. 41, No. 1, pp. 98–115.

Dobkin, C. and N. Nicosia (2009) "The war on drugs: methamphetamine, public health, and crime," *American Economic Review*, Vol. 99, No. 1, pp. 324–349.

Dobkin, C., N. Nicosia, and M. Weinberg (2014) "Are supply-side drug control efforts effective? Evaluating OTC regulations targeting methamphetamine precursors," *Journal of Public Economics*, Vol. 120, pp. 48–61.

Drake, C., J. Wen, J. Hinde, and H. Wen (2021) "Recreational cannabis laws and opioid-related emergency department visit rates," *Health Economics*, Vol. 30, No. 10, pp. 2595–2605.

EMCDDA and Europol (2020) "Impact of COVID-19," `https://www.emcdda.europa.eu/publications/joint-publications/eu-drug-markets-impact-of-covid-19_en`, [Accessed 18-Jun-2023].

Espinosa, R. (2019) "Scamming and the reputation of drug dealers on darknet markets," *International Journal of Industrial Organization*, Vol. 67, p. 102523.

Evans, W. N., E. M. Lieber, and P. Power (2019) "How the reformulation of OxyContin ignited the heroin epidemic," *Review of Economics and Statistics*, Vol. 101, No. 1, pp. 1–15.

Filippas, A., J. J. Horton, and J. M. Golden (2022) "Reputation Inflation," *Marketing Science*, Vol. 41, No. 4, pp. 733–745.

FilterMag (2020) "Revealing Trends in Russia's Dark-Web Drug Markets," https://filtermag.org/russia-dark-web-drugs/, [Accessed 18-Jun-2023].

Flashpoint, Chainanalysis (2021) "Cybercrime market "Hydra": Where the crypto money laundering trail goes dark," https://go.flashpoint-intel.com/docs/chainalysis-hydra-cryptocurrency-cybercrime.

Fryer, J., Roland G, P. S. Heaton, S. D. Levitt, and K. M. Murphy (2013) "Measuring Crack Cocaine and Its Impact," *Economic Inquiry*, Vol. 51, No. 3, pp. 1651–1681.

Gable, R. S. (2004) "Comparison of acute lethal toxicity of commonly abused psychoactive substances," *Addiction*, Vol. 99, No. 6, pp. 686–696.

Galenianos, M. and A. Gavazza (2017) "A structural model of the retail market for illicit drugs," *American Economic Review*, Vol. 107, No. 3, pp. 858–896.

Galenianos, M., R. L. Pacula, and N. Persico (2012) "A search-theoretic model of the retail market for illicit drugs," *The Review of Economic Studies*, Vol. 79, No. 3, pp. 1239–1269.

Gallet, C. A. (2014) "Can price get the monkey off our back? A meta-analysis of illicit drug demand," *Health Economics*, Vol. 23, No. 1, pp. 55–68.

Gandhi, A. and J.-F. Houde (2019) "Measuring substitution patterns in differentiated-products industries," *NBER Working paper*, No. w26375.

Gavrilova, E., T. Kamada, and F. Zoutman (2019) "Is legal pot crippling Mexican drug trafficking organisations? The effect of medical marijuana laws on US crime," *The Economic Journal*, Vol. 129, No. 617, pp. 375–407.

Goonetilleke, P., A. Knorre, and A. Kuriksha (2023) "Hydra: Lessons from the world's largest darknet market," *Criminology & Public Policy*, Vol. 22, No. 4, pp. 735–777.

Government of Russia (2012) "Postanovlenie ob utverzhdenii perechnya znachitel'nogo, krupnogo i osobo krupnogo razmerov narkoticheskih sredstv i psihotropnyh veshchestv (Decree on approval of the list of significant, large and especially large quantities of narcotic drugs and psychotropic substances)," https://www.consultant.ru/cons/cgi/online.cgi?req=doc&rnd=qdYQKA&base=LAW&n=407909, [Accessed 07-July-2023].

Gruber, J. and B. Köszegi (2001) "Is addiction "rational"? Theory and evidence," *The Quarterly Journal of Economics*, Vol. 116, No. 4, pp. 1261–1303.

Hall, W., D. Stjepanović, D. Dawson, and J. Leung (2023) "The implementation and public health impacts of cannabis legalization in Canada: a systematic review," *Addiction*.

Hansen, B., T. Moore, and W. Olney (2023) "Importing the opioid crisis? International trade and drug overdoses,"Technical report, Working paper.

Harwood, H. J. and E. Bouchery (2004) *The economic costs of drug abuse in the United States, 1992-2002*: Executive Office of the President, Office of National Drug Control Policy.

Hatsukami, D. K. and M. W. Fischman (1996) "Crack cocaine and cocaine hydrochloride: Are the differences myth or reality?," *Journal of the American Medical Association*, Vol. 276, No. 19, pp. 1580–1588, DOI: 10.1001/jama.1996.03540190052029.

Hausman, J., G. Leonard, and J. D. Zona (1994) "Competitive analysis with differenciated products," *Annales d'Economie et de Statistique*, pp. 159–180.

Hollenbeck, B. and K. Uetake (2021) "Taxation and market power in the legal marijuana industry," *The RAND Journal of Economics*, Vol. 52, No. 3, pp. 559–595.

Horowitz, J. L. (2001) "Should the DEA's STRIDE data be used for economic analyses of markets for illegal drugs?" *Journal of the American Statistical Association*, Vol. 96, No. 456, pp. 1254–1271.

Hui, K. (2023) "The Impact of a Vape Ban on Cigarette Smoking and Life Expectancy."

Hurd, Y. L., S. Spriggs, J. Alishayev, G. Winkel, K. Gurgov, C. Kudrich, A. M. Oprescu, and E. Salsitz (2019) "Cannabidiol for the reduction of cue-induced craving and anxiety in drug-abstinent individuals with heroin use disorder: a double-blind randomized placebo-controlled trial," *American Journal of Psychiatry*, Vol. 176, No. 11, pp. 911–922.

Idrisov, B., S. M. Murphy, T. Morrill, M. Saadoun, K. Lunze, and D. Shepard (2017) "Implementation of methadone therapy for opioid use disorder in Russia–a modeled cost-effectiveness analysis," *Substance Abuse Treatment, Prevention, and Policy*, Vol. 12, pp. 1–6.

Janetos, N. and J. Tilly (2017) "Reputation dynamics in a market for illicit drugs," *arXiv preprint arXiv:1703.01937*.

Jofre-Bonet, M. and N. M. Petry (2008) "Trading apples for oranges? Results of an experiment on the effects of heroin and cocaine price changes on addicts' polydrug use," *Journal of Economic Behavior & Organization*, Vol. 66, No. 2, pp. 281–311.

Judicial Department (2021a) "Otchet o chisle privlechennyh k ugolovnoj otvetstvennosti i vidah ugolovnogo nakazaniya (Report on the number of persons prosecuted and types of criminal punishment)," http://www.cdep.ru/userimages/Statistika_zameni/10.1-svod-2020.xls, [Accessed 07-July-2023].

——— (2021b) "Svedeniya o licah, osuzhdennyh za prestupleniya, svyazannye s nezakonnym oborotom narkoticheskih sredstv (Information about persons convicted of crimes related to drug trafficking)," http://www.cdep.ru/userimages/Statistika_zameni/6-mvnon-svod-2020.xls, [Accessed 07-July-2023].

Kennedy, M., P. Reuter, and K. J. Riley (1993) "A simple economic model of cocaine production," *Mathematical and Computer Modelling*, Vol. 17, No. 2, pp. 19–36.

Knife Media (2020) "Chto proishodit s rossiskoi narkotorgovlei iz-za koronavirusa? (What happens with Russian drug trade because of the Covid-19 pandemic?)," https://pandemic-research.github.io/coronavirus/, [Accessed 18-Jun-2023].

Kuziemko, I. and S. D. Levitt (2004) "An empirical analysis of imprisoning drug offenders," *Journal of Public Economics*, Vol. 88, No. 9-10, pp. 2043–2066.

Levitt, S. D. and S. A. Venkatesh (2000) "An economic analysis of a drug-selling gang's finances," *The Quarterly Journal of Economics*, Vol. 115, No. 3, pp. 755–789.

Liu, J.-L., J.-T. Liu, J. K. Hammitt, and S.-Y. Chou (1999) "The price elasticity of opium in Taiwan, 1914–1942," *Journal of Health Economics*, Vol. 18, No. 6, pp. 795–810.

Luca, M. and O. Reshef (2021) "The effect of price on firm reputation," *Management Science*, Vol. 67, No. 7, pp. 4408–4419.

Manski, C. F., J. V. Pepper, and C. V. Petrie (2001) *Informing America's policy on illegal drugs: What we don't know keeps hurting us*: National Academies Press.

Mejia, D. and P. Restrepo (2016) "The economics of the war on illegal drug production and trafficking," *Journal of Economic Behavior & Organization*, Vol. 126, pp. 255–275.

Miravete, E. J., K. Seim, and J. Thurk (2018) "Market power and the Laffer curve," *Econometrica*, Vol. 86, No. 5, pp. 1651–1687.

Moore, T. J. and K. T. Schnepel (2021) "Opioid use, health and crime: Insights from a rapid reduction in heroin supply,"Technical report, National Bureau of Economic Research.

National Drug Control Budget (2023) "National Drug Control Budget Funding Highlights," https://www.whitehouse.gov/wp-content/uploads/2023/03/FY-2024-Budget-Highlights.pdf.

Neelakantan, A., T. Xu, R. Puri, A. Radford, J. M. Han, J. Tworek, Q. Yuan, N. Tezak, J. W. Kim, C. Hallacy et al. (2022) "Text and code embeddings by contrastive pre-training," *arXiv preprint arXiv:2201.10005*.

Nevo, A. (2001) "Measuring market power in the ready-to-eat cereal industry," *Econometrica*, Vol. 69, No. 2, pp. 307–342.

New York Times (2023) "America's New Drug Policy," https://www.nytimes.com/2023/05/04/briefing/us-drug-policy-reducing-harm.html, [Accessed 08-Jun-2023].

Nutt, D., L. A. King, W. Saulsbury, and C. Blakemore (2007) "Development of a rational scale to assess the harm of drugs of potential misuse," *The Lancet*, Vol. 369, No. 9566, pp. 1047–1053.

Olmstead, T. A., S. M. Alessi, B. Kline, R. L. Pacula, and N. M. Petry (2015) "The price elasticity of demand for heroin: Matched longitudinal and experimental evidence," *Journal of Health Economics*, Vol. 41, pp. 59–71.

Pacula, R. L. (1997) "Women and substance use: are women less susceptible to addiction?" *The American Economic Review*, Vol. 87, No. 2, pp. 454–459.

Patocka, J., B. Zhao, W. Wu, B. Klimova, M. Valis, E. Nepovimova, and K. Kuca (2020) "Flakka: new dangerous synthetic cathinone on the drug scene," *International Journal of Molecular Sciences*, Vol. 21, No. 21, p. 8185.

States of America V. Dmitry Olegovich Pavlov, U. (2022) "United States District Court," https://www.justice.gov/opa/press-release/file/1490906/download, [Accessed 18-Jun-2023].

Peacock, A., R. Bruno, N. Gisev, L. Degenhardt, W. Hall, R. Sedefov, J. White, K. V. Thomas, M. Farrell, and P. Griffiths (2019) "New psychoactive substances: challenges for drug surveillance, control, and public health responses," *The Lancet*, Vol. 394, No. 10209, pp. 1668–1684.

Petrin, A. (2002) "Quantifying the benefits of new products: The case of the minivan," *Journal of Political Economy*, Vol. 110, No. 4, pp. 705–729.

Powell, D., R. L. Pacula, and M. Jacobson (2018) "Do medical marijuana laws reduce addictions and deaths related to pain killers?" *Journal of Health Economics*, Vol. 58, pp. 29–42.

Poyatos, L., A. Torres, E. Papaseit, C. Pérez-Mañá, O. Hladun, M. Núñez-Montero, G. de la Rosa, M. Torrens, D. Fuster, R. Muga et al. (2022) "Abuse Potential of Cathinones in Humans: A Systematic Review," *Journal of Clinical Medicine*, Vol. 11, No. 4, p. 1004.

Proekt (2019) "Vsya eta dur. Issledovanie o tom, na chem sidit Rossiya (All those drugs. A study on narcotics that Russia is hooked on)."

Ramful, P. and X. Zhao (2009) "Participation in marijuana, cocaine and heroin consumption in Australia: a multivariate probit approach," *Applied Economics*, Vol. 41, No. 4, pp. 481–496.

Reuter, P. (2010) *Understanding the demand for illegal drugs*: National Academies Press.

Saffer, H. and F. Chaloupka (1999) "The demand for illicit drugs," *Economic Inquiry*, Vol. 37, No. 3, pp. 401–411.

Saidashev, R. and A. Meylakhs (2021) "A qualitative analysis of the Russian cryptomarket Hydra," *Kriminologisches Journal*, Vol. 53, pp. 169 – 185.

Shover, C. L., C. S. Davis, S. C. Gordon, and K. Humphreys (2019) "Association between medical cannabis laws and opioid overdose mortality has reversed over time," *Proceedings of the National Academy of Sciences*, Vol. 116, No. 26, pp. 12624–12626.

Soares, J., V. M. Costa, M. d. L. Bastos, F. Carvalho, and J. P. Capela (2021) "An updated review on synthetic cathinones," *Archives of Toxicology*, Vol. 95, No. 9, pp. 2895–2940.

Soska, K. and N. Christin (2015) "Measuring the longitudinal evolution of the online anonymous marketplace ecosystem," in *24th USENIX Security Symposium (USENIX Security 15)*, pp. 33–48.

TASS (2019) "V Pskovskoj oblasti tamozhenniki obnaruzhili 450 kg gashisha v toplivnom bake gruzovika (In the Pskov region, customs officers found 450 kg of hashish in the fuel tank of a truck)," https://tass.ru/proisshestviya/6973691, [Accessed 18-Jun-2023].

The Economist (2023) "Oregon botches the decriminalisation of drugs," https://www.economist.com/leaders/2023/04/13/oregon-botches-the-decriminalisation-of-drugsl, [Accessed 08-Jun-2023].

UN Office on Drugs and Crime (2016) "World Drug Report 2016," https://www.unodc.org/doc/wdr2016/WORLD_DRUG_REPORT_2016_web.pdf.

——— (2022) "World Drug Report 2022," https://www.unodc.org/unodc/en/data-and-analysis/world-drug-report-2022.html.

United Nations Office on Drugs and Crime (2021) "Global overview of drug demand and drug supply," https://www.unodc.org/unodc/en/data-and-analysis/wdr-2021_booklet-2.html, [Accessed 18-Jun-2023].

U.S. Department of the Treasury (2022) "Treasury Sanctions Russia-Based Hydra, World's Largest Darknet Market, and Ransomware-Enabling Virtual Currency Exchange Garantex," https://home.treasury.gov/news/press-releases/jy0701, [Accessed 18-Jun-2023].

Van Ours, J. C. (1995) "The price elasticity of hard drugs: The case of opium in the Dutch East Indies, 1923-1938," *Journal of Political Economy*, Vol. 103, No. 2, pp. 261–279.

Van Ours, J. C. and J. Williams (2007) "Cannabis prices and dynamics of cannabis use," *Journal of Health Economics*, Vol. 26, No. 3, pp. 578–596.

VICE (2020) "A New Breed of Drug Dealer Has Turned Buying Drugs into a Treasure Hunt," https://www.vice.com/en/article/g5x3zj/hydra-russia-drug-cartel-dark-web, [Accessed 18-Jun-2023].

White, M. D. and W. A. Luksetich (1983) "Heroin: price elasticity and enforcement strategies," *Economic Inquiry*, Vol. 21, No. 4, p. 557.

Winstock, A., L. Mitcheson, J. Ramsey, S. Davies, M. Puchnarewicz, and J. Marsden (2011) "Mephedrone: use, subjective effects and health risks," *Addiction*, Vol. 106, No. 11, pp. 1991–1996.

Appendices

A. Data

A.1. Scraping of Hydra

The scraping process was done by running a program on a personal computer.[47] The computer operated on OS Windows 10 and had an AMD processor and 4GB of RAM. The process was organized in two stages. In the first stage, the program scraped all pages with the output of the search in 62 categories of the Hydra website to obtain URLs of all products within each category.[48] After that, the program iterated over all obtained product URLs and scraped each product-specific page to collect information on the listings available for the product.

Table A.1: List of dates when scrapes of Hydra are available.

Date	Day of week	Week #	Listings	Date	Day of week	Week #	Listings
Jul 17, 2019	Wed	29	73,313	Jan 22, 2020	Wed	4	94,859
Jul 20, 2019	Sat	29	73,448	Jan 28, 2020	Tue	5	95,634
Jul 30, 2019	Tue	31	77,652	Feb 06, 2020	Thu	6	106,351
Aug 07, 2019	Wed	32	77,898	Feb 12, 2020	Wed	7	107,449
Aug 14, 2019	Wed	33	80,911	Feb 20, 2020	Thu	8	112,504
Aug 30, 2019	Fri	35	87,357	Feb 27, 2020	Thu	9	110,864
Sep 08, 2019	Sun	36	84,934	Mar 05, 2020	Thu	10	118,579
Sep 17, 2019	Tue	38	84,511	Mar 11, 2020	Wed	11	114,769
Sep 25, 2019	Wed	39	88,750	May 16, 2020	Sat	20	119,769
Oct 03, 2019	Thu	40	91,293	May 23, 2020	Sat	21	126,192
Nov 15, 2019	Fri	46	89,510	Jun 16, 2020	Tue	25	138,312
Nov 27, 2019	Wed	48	93,188	Jul 03, 2020	Fri	27	153,464
Dec 06, 2019	Fri	49	96,817	Jul 15, 2020	Wed	29	156,465
Dec 13, 2019	Fri	50	103,550	Aug 05, 2020	Wed	32	167,312
Dec 19, 2019	Thu	51	101,335	Aug 27, 2020	Thu	35	168,508
Dec 26, 2019	Thu	52	105,832				

Available dates. The script was run on 33 days from July 17, 2019 to Aug 27, 2020. On two days, November 21 of 2019 and September 9 of 2020, the program failed to complete

[47] The code is available upon request.

[48] For example, ⋆.onion/catalog/3?page=1 was the first page listing marijuana products.

scraping due to a technical error. We exclude these days from the sample. The list of days and the total number of listings scraped are provided in Table A.1.

A.2. Data cleaning

We drop all listings with a price per gram greater than three times the median price of the same drug type and amount. This is necessary because shops on Hydra sometimes set prohibitively high prices instead of deleting a listing when they were out of stock.[49] We also observe several shops with many thousands of listings and a much smaller cumulative number of fulfilled orders. A common feature of such shops is that they operated in more cities than even the largest shops on the platform. This seems to be inconsistent with the normal operation of shops on Hydra. One potential explanation is that this is a form of drop shipping: these shops copy the listings of other shops and sell those listings at a premium. Because drop-shipping merchants copy the listings of other shops, including their listings would lead to a double-counting of dead drops. We drop from our data all shops for which there are more listings than total sales.

We exclude all reviews of job postings or non-drug products sold on Hydra. We also remove duplicates if we observe several reviews with the exact same text left for the same product by the same user. We only use reviews for the period when we have listings from the marketplace.

A.3. Dose definition

Table A.2 displays the distribution of different quantities for each drug type. To account for potential differences in potency between various drug types and substance forms, we normalize the listed amounts by dividing them by a drug-specific quantity, which we refer to as the "standard amount" or "dose." Intuitively, the standard amount represents the first frequently used quantity in the distribution of listed amounts. We define a dose for each drug type as an amount with a frequency of at least 15% and at least 40% of the frequency of any other higher quantity.

Our interpretation of this definition is that the first popular amount in the distribution of quantities corresponds to the "minimal suitable quantity." While the distribution of purchased amounts could be endogenous and dependent on the price and other factors, it is plausible that for each drug type, there exists a subset of constrained consumers who will only purchase this minimal suitable quantity. Our strategy aims to identify this specific quantity and use it for normalization.

[49] Sellers use such strategies on legal online platforms as well, e.g., on Airbnb (Culotta et al., 2022).

Table A.2: Shares of different amounts for each drug type

	0.1	0.25	0.5	1	2	3	Other	Total
2c	**63.5**	7.7	0.8	0.0	0.0	0.0	28.0	100.0
α-PVP	0.1	1.2	**33.1**	34.5	15.8	0.0	15.3	100.0
Amphetamine	0.0	0.0	5.6	**32.0**	27.1	18.4	16.8	100.0
Marijuana (buds)	0.0	0.0	4.3	**39.6**	25.3	16.0	14.7	100.0
Cannabinoids	0.0	0.0	**20.5**	43.2	35.6	0.0	0.8	100.0
Cannafood	0.0	0.0	0.7	**24.2**	46.2	10.9	18.0	100.0
Cocaine	0.3	3.3	**31.3**	40.6	13.9	0.0	10.6	100.0
Dissociatives	0.0	2.7	**33.6**	37.3	9.9	0.0	16.4	100.0
DMT	**16.0**	23.8	39.6	0.0	0.0	0.0	20.7	100.0
GHB	0.7	16.3	**64.5**	7.8	0.0	0.0	10.6	100.0
Hashish oil	1.3	**13.7**	23.3	46.6	0.0	0.0	15.1	100.0
Hashish	0.0	0.1	5.1	**32.9**	27.4	17.5	17.1	100.0
Heroin	1.4	**14.5**	34.0	37.7	0.0	0.0	12.4	100.0
Marijuana (leaves)	0.0	0.0	0.8	**17.2**	25.1	26.1	30.9	100.0
LSD	0.0	0.0	0.0	10.0	**34.8**	18.4	36.9	100.0
MDMA (pill)	0.0	0.0	0.0	8.2	**27.2**	24.0	40.6	100.0
MDMA (crystal)	0.0	2.2	**27.9**	41.5	21.0	0.0	7.4	100.0
MDPV	0.0	1.3	**28.7**	43.2	25.7	0.0	1.0	100.0
Mephedrone	0.0	0.3	9.2	**36.0**	24.5	15.9	14.0	100.0
Methadone	4.6	**20.8**	33.8	23.8	0.0	0.0	17.1	100.0
Methamphetamine	0.0	1.6	**19.4**	42.5	29.7	0.0	6.7	100.0
Mushrooms	0.0	0.0	0.0	**19.3**	9.1	40.2	31.4	100.0
NBOME	0.0	0.0	0.0	11.6	**25.1**	23.3	39.9	100.0
Opioids	1.0	5.1	**36.4**	28.6	17.8	0.0	11.1	100.0
Psychedelics	0.0	0.0	0.9	3.5	**20.8**	29.0	45.8	100.0
Synthetic cannabinoids	0.0	0.0	2.8	**27.2**	27.4	20.9	21.7	100.0
Total	0.2	1.2	13.0	32.3	22.8	12.9	17.6	100.0

Note: Each column shows shares of a corresponding amount for each drug type. For MDMA (pills) and LSD, the amount is in counts. For all other drug types, the amount is in grams. The standard amount for each drug type is highlighted in bold and defined as an amount with a frequency of at least 15% and a frequency that is at least 40% of the frequency of any other higher amount.

A.4. Descriptive statistics

Shops characteristics. Table A.3 presents descriptive statistics for the characteristics of shops on Hydra.

Reviews. Table A.4 presents summary statistics for our data on reviews.

Table A.3: Summary statistics: properties of shops on Hydra

	Mean	Median	Std	Min	Max
Products offered	7.61	6	5.83	1	60
Drug types offered	3.35	3	2.25	1	14
Cities present	3.36	2	4.61	1	33
Daily listings	37.53	20	77.69	1	2,315
Age (months)	17.71	15.10	11.42	0	44.50
Total sales	13,830	4,500	39,333	3	800,000
Rating	4.90	4.93	0.11	2.65	5
Trusted Seller	0.17	–	–	–	–

Table A.4: Summary statistics for reviews data

Panel A: Available data					
Reviews	Total				215,733
	By users with 1 review				82,878
	By users with > 1 reviews				132,855
Users	Total				126,259
	With 1 review				82,878
	With > 1 reviews				43,381
	With > 1 purchased types				21,864
Shops	Total				784
Panel B: Descriptive statistics					
	Mean	Median	Std	Min	Max
Rating	9.57	10	1.88	0	10
Number of words	16.39	8	26.76	0	1596
Reviews per nickname	1.71	1	1.54	1	63
Days between reviews	58.11	14	97.98	0	441

Note: Days between reviews are calculated for reviews such that a review left by the same user at a later day is present in the sample.

Table A.5: Shares of drug types in all listings and reviews

Drug type	Share of listings	Share of reviews	Drug type	Share of listings	Share of reviews
Mephedrone	29.6%	26.3%	Heroin	0.8%	1.4%
Amphetamine	13.9%	12.8%	Marijuana (leaves)	0.7%	0.6%
α-PVP	12.5%	8.6%	NBOME	0.7%	0.7%
MDMA	10.8%	12.4%	Synthetic cannabinoids	0.7%	1.1%
Marijuana (buds)	7.9%	10.4%	Mushrooms	0.4%	0.9%
Hashish	7.8%	7.7%	Hashish oil	0.3%	0.5%
Cocaine	7.1%	9.0%	Dissociatives	0.3%	0.5%
LSD	3.1%	2.6%	DMT	0.2%	0.2%
Methadone	2.0%	2.7%	Cannafood	0.1%	0.3%
Methamphetamine	0.9%	0.9%	2C-B	0.1%	0.3%

A.5. Proxies for sales

Shares of drugs. Table A.5 shows market shares of different drug types defined through listings and reviews. While we do not observe actual transactions, we can compare the two proxies for transactions on the aggregate level to test their validity. We find that these two numbers generally are close to each other. The largest absolute discrepancies are observed for mephedrone, $\alpha-$PVP, MDMA, and marijuana.

Listings and cryptocurrency inflows. Figure A.1 shows the monthly estimates of Hydra revenue provided in Flashpoint, Chainanalysis (2021)[p. 5] and the average daily number of listings in our data during the same period. Our results indicate a strong correlation between cryptocurrency inflows to Hydra and the number of listings across time.

Figure A.1. Estimated revenue and listings on Hydra over time

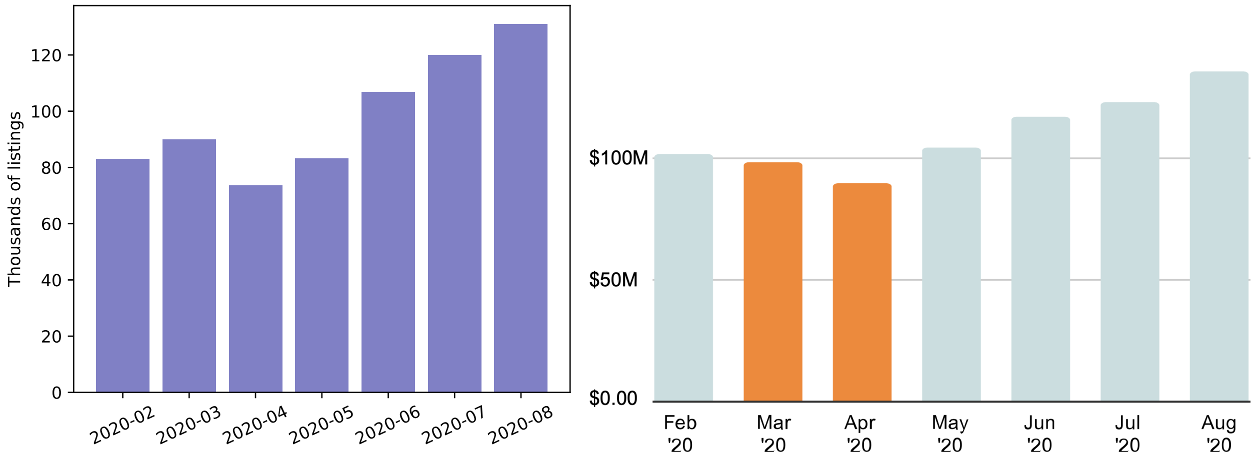

(a) Mean daily number of listings

(b) Estimated monthly revenue

Note: Estimates and the plot for monthly revenue are obtained from Flashpoint, Chainanalysis (2021)[p. 5]. To calculate mean listings for April 2020, we use another dataset, which was purchased from an independent data collector for this particular month and was also used by Goonetilleke et al. (2023). This data is not used for demand estimation.

B. Sentiment of Reviews

We apply two approaches to extract the sentiment of reviews: the lexicon approach and LLM embeddings. In the first approach, we start by constructing a balanced sample of reviews with positive and negative ratings. We label a review as positive if it has a rating of 10/10. We label a review as negative if it has a rating below 6/10. We obtain a total of 14,000 negative reviews. We randomly select the same number of positive reviews. We then apply lemmatization to words and drop all "stopwords," for example, prepositions and pronouns. We find in the corpus of processed reviews the 200 most frequently used terms. We use the frequencies of these terms to generate a vector of 200 elements for each review. We standardize these variables and include the length of a review as an additional predictor. This gives us a vector representation $\mathbb{X}_i$ for each review i in the sample. We then run a logistic regression to estimate the model

$$\mathbb{P}(\text{review } i \text{ is positive}) = \text{logit}(\mathbb{X}_i'\beta).$$

We obtain an out-of-sample accuracy of 85% with this algorithm. We use $\mathbb{X}_i'\hat{\beta}$ as our measure of the (positive) sentiment of the review.

The intuition behind this method is the following. From rating-based labels of reviews, the algorithm learns which words have good and bad connotations. Then, weighting these words allows us to distinguish differences in the sentiment even within the 96% of reviews that have the highest possible rating. Figure B.1 describes the words that have the largest power for predicting positive or negative labels. The most predictive signal for positive feedback is "10": reviewers type numerical ratings to express satisfaction. The length of the review is a strong predictor of negative feedback. Most of the words we find predictive for positive feedback describe general satisfaction with the purchase, for example, "thank you" or "super." Many words are related to the delivery process. For example, "not-found" describes the common problem of not being able to find the purchased dead-drop; "touch" is a colloquial way to explain that the drug was picked up quickly; and "photo" is often used to complain about the quality of the photo of the dead-drop location. Some words seem to be used to describe the substance, e.g., "quality" and "stuff." Others describe the process of disputes, e.g., "favor," "dispute," or "coupon."

However, this method does not take into account many properties of language – for example, the difference between "recommend" and "not recommend." For this reason, we also use modern developments in large language models for our sentiment analysis. For each review R, we obtain a vector embedding $e(R) \in \mathbb{R}^{1536}$ using the API from OpenAI.[50]. We

[50] See Neelakantan et al. (2022) for more details.

Figure B.1. Most predictive words for positive and negative reviews

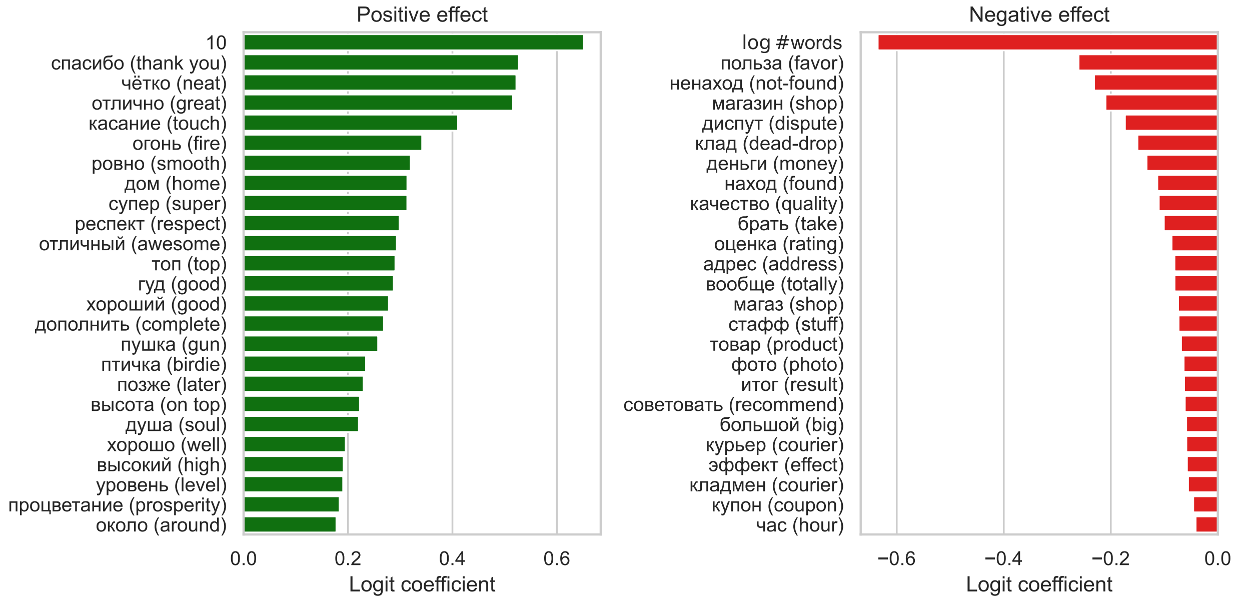

then manually choose a small sample[51] of positive and negative reviews, with 25 reviews in each group. We define our measure of positive sentiment as the difference between average cosine similarity to good reviews and average cosine similarity to bad reviews. That is,

$$sentiment(R) = \frac{1}{\#G} \sum_{X \in G} D_C\big(e(R), e(X)\big) - \frac{1}{\#B} \sum_{X \in B} D_C\big(e(R), e(X)\big),$$

where G is the set of good reviews, B is the set of bad reviews, and cosine similarity is given by

$$D_C(X, Y) = 1 - \frac{X \cdot Y}{\|X\| \|Y\|}.$$

The two obtained measures are highly correlated, with a correlation coefficient of 0.67. We use the first principal component of these two measures to obtain our final estimate of review sentiment. Given the varying number of reviews across shops in our sample, we employ empirical Bayes to regularize the obtained shop-level average sentiment.

[51]To minimize the computational costs associated with including every additional review in this sample, we manually selected reviews that encompass a variety of scenarios reflecting both satisfaction and dissatisfaction.

Table B.1: Examples of reviews with highest rating and negative sentiment

Date	Drug	Translation	Original	Rating
2020-06-04	Amphetamine	Fast collection, respect to the courier. But the quality is below average, I expected a lot more. Did not get any pleasure or feelings from it.	В касание, минеру респект. Но качество ниже среднего, я ожидал на много большего, а не получил от этого ни удовольствия ни ощущений	10
2020-06-09	Heroin	Fast collection, the product is damp. Brothers, do not even think to buy heroin from here, the dead-drops are from the winter, and the product does not work well.	Забрал в касание, товар отсырел, братчанин, не вздумай тут покупать хмурый, зимние адреса, товар прёт ну точно не 777	10
2019-06-06	Mephedrone	Did not find the dead-drop, but I can only blame myself. Also, do not want to start a dispute because of just 0.5 grams. I guess I will not buy dug dead-drops for a while.	Сокровище не нашёл, но тут скорее могу винить только себя, да и из-за 0.5 диспут открывать не хочется.. Пожалуй, не буду больше пока брать прикопы)	10
2020-08-01	Amphetamine	Fast collection, also an interesting experience. But the quality is quite bad... No offense guys. Rating 10/10/10, will not lower it.	В касание! Интересный опыт по касашке... Но качество чёт подводит.... Без обид, пацаны. Оценка 10.10.10 понижать не буду	10
2020-05-26	Mephedrone	We found everything but with lots of complications. The product was around the specified location.	Все нашли но с большими трудностями товар был рядом с указанным местом	10
2019-02-15	MDMA	In general it was good, but some of the pills were broken. The courier confuses left and right. Liked the quality.	В целом всё в порядке, но таблы оказались поломанные. И кладмен путает лево и право. Качество понравилось	10
2020-04-03	Marijuana buds	Good buds but not dried enough. Thus, the quantity actually is smaller than specified	Хорошие шишки, только недосушены, соответсвенно количество меньше чем заявлено	10
2020-05-20	Hashish	Did not find the dead-drop, it was hidden badly and the location was marked badly. When you pay 2800 rubles per 1 gram you expect a good dead-drop. The support responses slower than once per day. In the end, they gave me a coupon. Overall, not satisfied with the shop.	Был ненаход, откровенно говоря плохо спрятали и плохо метку поставили, когда 2800 за 1г. отдаешь рассчитываешь на нормальную закладку, поддержка у магазина отвечат даже не раз в сутки, в итоге разошлись купоном, всем магазином в целом не доволен.	10

Note: Table sourced from Goonetilleke et al. (2023).

C. Micro Moments

Here we provide a more detailed discussion of our micro-level moment conditions. We start by showing how micro moments can facilitate demand estimation in a simulated dataset.

C.1. Simulated example

We generate simulated data in which consumers have correlated taste shocks for two products. This is a simplified version of the demand model presented in Section 4. Specifically, there are products $j = 1, \ldots, 5$ sold by 5 different firms. Consumer preferences are given by

$$U_{ijt} = -\alpha p_{jt} + \beta^0 + \sum_{n=1}^{3} \beta^n x_{jt}^n + \lambda_i^0 + \lambda_i^1 I(j = 1) + \lambda_i^2 I(j = 2) + \xi_{jt} + \varepsilon_{ijt}. \tag{13}$$

The linear parameters of demand are $\alpha = -5$, $\beta = (5, 1, 1, 1)$. Consumers have correlated random coefficients λ for the constant term and the dummies for products 1 and 2:

$$\begin{pmatrix} \lambda_i^0 \\ \lambda_i^1 \\ \lambda_i^2 \end{pmatrix} = \Sigma \nu_i,\ \nu_i \sim \mathcal{N}(0, I_3),\ \Sigma = \begin{pmatrix} 2 & 0 & 0 \\ 1 & 2 & 0 \\ 1 & 2 & 2 \end{pmatrix}. \tag{14}$$

Intuitively, consumers who like product 1 (2) are more inclined to like product 2 (1) and less inclined to choose the outside option. Prices are given by the Bertrand-Nash equilibrium where producers maximize total profits $(p_{jt} - MC_{jt}) s_{jt}$ and face marginal costs that are given by

$$MC_{jt} = 1 + 0.1 \sum_{n=1}^{3} x_{jt}^n + 0.1 \sum_{n=1}^{7} z_{jt}^n + \omega_{jt}, \tag{15}$$

where z^n are observed cost shifters, x^n are observed product characteristics, and ω is an unobserved product characteristic. We simulate $T = 100$ markets with 500 Monte Carlo agent draws. Variables x, z are all iid from $\mathcal{N}(0, 1)$, and $\omega, \xi \sim \mathcal{N}(0, 1)$ with $corr(\omega, \xi) = 0.5$.

We then obtain an analog of our review moments. We simulate 100,000 agents in this economy who keep the same draws ν_i across all markets. To make our setting closer to the empirical setting in the paper, we consider a theoretical counterpart of reviews: for each agent-period pair, if at period t agent i chooses product $j = j(i, t)$, this choice is observed with probability $\pi = 0.1$ and added to R_{ij}. Because the econometrician can only observe consumers with at least one review, we calculate averages $\overline{R_{ij} R_{ik}}$ over agents i such that $R_i > 0$. Table C.1 shows that our inter-temporal micro moments reflect the assumptions about correlations for products 1 and 2: consumers are more likely to purchase 1 and 2 together. The assumed correlation between λ_i^1 and λ_i^2 also implies that reviews for products 3 to 5 are correlated as well. Intuitively, the agents who buy these products are the agents who do not like products 1 and 2.

Table C.1: Moment values $\overline{R_{ij}R_{ik}}$ in simulated data

	R_1	R_2	R_3	R_4	R_5
R_1	3.17	2.56	1.38	1.52	1.52
R_2	2.56	7.82	2.00	2.15	2.10
R_3	1.38	2.00	3.23	2.16	2.24
R_4	1.52	2.15	2.16	3.74	2.44
R_5	1.52	2.10	2.24	2.44	3.97

We then try to estimate the model using the simulated data. To make the exercise closer to the setting of the paper, we only estimate the demand parameters of the model and do not rely on supply-side moment conditions. We use the observable cost shifters z^n and the differentiation IVs of Gandhi and Houde (2019) as the instrumental variables. First, we apply the standard BLP procedure, which uses the aggregate price-quantity data only, and obtain estimates $\hat{\Sigma}^{BLP}, \hat{\alpha}^{BLP}, \hat{\beta}^{BLP}$. Then, we estimate the parameters by fitting the predicted micro moments to estimated micro moments, as described in Sections 4.1.1 and 4.1, and obtain estimates $\hat{\Sigma}^{Micro}, \hat{\alpha}^{Micro}, \hat{\beta}^{Micro}$. Our results are provided below:

$$\hat{\Sigma}^{BLP} = \begin{pmatrix} 1.35 & 0.00 & 0.00 \\ 1.39 & 1.26 & 0.00 \\ -0.49 & 2.54 & 2.87 \end{pmatrix}, \ \hat{\Sigma}^{Micro} = \begin{pmatrix} 1.97 & 0.00 & 0.00 \\ 1.27 & 1.96 & 0.00 \\ 1.16 & 1.97 & 1.97 \end{pmatrix},$$

$$\hat{\alpha}^{BLP} = -4.70, \ \hat{\alpha}^{Micro} = -4.44,$$

$$\hat{\beta}^{BLP} = (3.83, 1.07, 1.09, 0.96), \ \hat{\beta}^{Micro} = (4.02, 1.02, 0.96, 1.01).$$

As can be seen, micro moments substantially improved precision of estimates for Σ.

C.2. Definition of periods

In this section, we describe how we apply our micro moments from Section 4.1.1 to the case when price-quantity data is only available for a subset of days. Suppose that reviews can be observed over days $t = 1, \dots, T$. However, quantities and prices can only be observed for several specific days $\tau_1, \dots, \tau_n$, where $1 \leqslant \tau_k \leqslant T$. In our case, reviews can be observed for $T = 423$ days, but we only have listings data for $n = 31$ days. In principle, we could keep reviews for days τ_k only and use the expressions from Section 4.1.1 directly. However, that would imply not utilizing most of the review data.

The expected value of $R_{ij}R_{ik}$ among all agents who left at least one observed review is

$$\mathbb{E}\Big[R_{ij}R_{ik} \mid R_i > 0\Big] = \frac{\mathbb{E}R_{ij}R_{ik}}{\mathbb{P}\big(R_i > 0\big)}, \tag{5}$$

where we approximate the denominator and the numerator by averages $\frac{1}{N}\sum_i \mathbb{E}\big[R_{ij}R_{ik} \mid i\big]$ and $\frac{1}{N}\sum_i \mathbb{P}\big(R_i > 0 \mid i\big)$, respectively. Including reviews for all days $t = 1, \ldots, T$, the expected value of the product term for consumer i is

$$\mathbb{E}\Big[R_{ij}R_{ik} \mid i\Big] = \sum_{t_1 \neq t_2} \pi_{jt_1}\pi_{kt_2}s_{ijt_1}s_{ikt_2} + I(j=k)\sum_{t=1}^{T}\pi_{jt}s_{ijt}. \tag{16}$$

Because we do not observe prices and quantities for other days, we approximate s_{ikt} by finding the closest day $\tau(t)$ when we observe listings for each t and using $s_{ik\tau}$ instead:

$$\begin{aligned}\mathbb{E}\Big[R_{ij}R_{ik} \mid i\Big] &\approx \sum_{t_1 \neq t_2} \pi_{jt_1}\pi_{kt_2}s_{ij\tau(t_1)}s_{ik\tau(t_2)} + I(j=k)\sum_{t}\pi_{jt}s_{ij\tau(t)} \\ &= \sum_{\tau_1,\tau_2}\big(\sum_{\tau(t)=\tau_1}\pi_{jt}\big)\big(\sum_{\tau(t)=\tau_2}\pi_{kt}\big)s_{ij\tau_1}s_{ik\tau_2} \\ &+ I(j=k)\sum_{\tau}\big(\sum_{\tau(t)=\tau}\pi_{jt}\big)s_{ij\tau} \\ &- \sum_{\tau}\Big(\sum_{\tau(t)=\tau}\pi_{jt}\pi_{kt}\Big)s_{ij\tau}s_{ik\tau}.\end{aligned} \tag{17}$$

We find that terms $\sum_{\tau(t)=\tau}\pi_{jt}\pi_{kt}$ are two orders of magnitude smaller compared to terms $\sum_{\tau(t)=\tau}\pi_{jt}$ and are one order of magnitude smaller than the terms $\big(\sum_{\tau(t)=\tau_1}\pi_{jt}\big)\big(\sum_{\tau(t)=\tau_2}\pi_{kt}\big)$.[52] Therefore, we can further approximate

$$\mathbb{E}\Big[R_{ij}R_{ik} \mid i\Big] \approx \sum_{\tau_1,\tau_2}\tilde{\pi}_{j\tau_1}\tilde{\pi}_{j\tau_2}s_{ij\tau_1}s_{ik\tau_2} + I(j=k)\sum_{\tau}\tilde{\pi}_{j\tau}s_{ij\tau}, \tag{18}$$

where $\tilde{\pi}_{j\tau} = \sum_{\tau(t)=\tau}\pi_{jt}$ is the sum of probabilities of conversion into observed review over all days t attributed to τ.

If we apply the approximation by the closest observed day to the equation describing the

[52]Intuitively, the last term in equation 17 corrects for the fact that consumers cannot purchase j and k on the same day. This possibility has a relatively negligible role in $R_{ij}R_{ik}$ if reviews are observed rarely (π_{jt} are small) or T is large and cross-period combinations dominate. Both apply in our setting.

connection between reviews and sales, we obtain

$$\sum_{\tau(t)=\tau} R_{jt} = N \sum_{\tau(t)=\tau} \pi_{jt} s_{jt} \approx N \sum_{\tau(t)=\tau} \pi_{jt} s_{j\tau} = N \tilde{\pi}_{j\tau} s_{j\tau}. \tag{19}$$

Thus, we can estimate $\tilde{\pi}_{j\tau}$ as the ratio of reviews over the larger period $\{t : \tau(t) = \tau\}$ to $N s_{j\tau}$, as is done in Equation 9.

Finally, we also can approximate the selection probability in a similar way:

$$\begin{aligned} \mathbb{P}\big(R_i > 0 \mid i\big) &= 1 - \prod_{t=1}^{T} \Big(1 - \pi_{jt} \sum_{j=1}^{J} s_{ijt}\Big) \\ &\approx 1 - \prod_{t=1}^{T} \Big(1 - \pi_{jt} \sum_{j=1}^{J} s_{ij\tau(t)}\Big) \\ &\approx 1 - \prod_{\tau} \Big(1 - \tilde{\pi}_{j\tau} \sum_{j=1}^{J} s_{ij\tau(t)}\Big). \end{aligned} \tag{20}$$

Table C.2 shows the assignment of dates in our review data to dates $\tau(t)$ in our listings data and the number of reviews for each τ.

Scrape date	Period start	Period end	Length	Reviews	Scrape date	Period start	Period end	Length	Reviews
Jul 17, 2019	Jul 01, 2019	Jul 18, 2019	17	1,815	Jan 22, 2020	Jan 09, 2020	Jan 25, 2020	30	2,127
Jul 20, 2019	Jul 19, 2019	Jul 25, 2019	8	4,991	Jan 28, 2020	Jan 26, 2020	Feb 01, 2020	10	1,798
Jul 30, 2019	Jul 26, 2019	Aug 03, 2019	14	6,566	Feb 06, 2020	Feb 02, 2020	Feb 09, 2020	12	8,688
Aug 07, 2019	Aug 04, 2019	Aug 10, 2019	11	1,097	Feb 12, 2020	Feb 10, 2020	Feb 16, 2020	10	5,158
Aug 14, 2019	Aug 11, 2019	Aug 22, 2019	15	5,784	Feb 20, 2020	Feb 17, 2020	Feb 23, 2020	11	2,933
Aug 30, 2019	Aug 23, 2019	Sep 03, 2019	20	2,015	Feb 27, 2020	Feb 24, 2020	Mar 01, 2020	10	536
Sep 08, 2019	Sep 04, 2019	Sep 12, 2019	13	1,041	Mar 05, 2020	Mar 02, 2020	Mar 08, 2020	10	686
Sep 17, 2019	Sep 13, 2019	Sep 21, 2019	13	1,254	Mar 11, 2020	Mar 09, 2020	Apr 13, 2020	39	15,180
Sep 25, 2019	Sep 22, 2019	Sep 29, 2019	12	110	May 16, 2020	Apr 14, 2020	May 19, 2020	69	25,365
Oct 03, 2019	Sep 30, 2019	Oct 24, 2019	29	642	May 23, 2020	May 20, 2020	Jun 04, 2020	19	34,198
Nov 15, 2019	Oct 25, 2019	Nov 21, 2019	49	761	Jun 16, 2020	Jun 05, 2020	Jun 24, 2020	32	18,457
Nov 27, 2019	Nov 22, 2019	Dec 01, 2019	16	380	Jul 03, 2020	Jun 25, 2020	Jul 09, 2020	23	11,005
Dec 06, 2019	Dec 02, 2019	Dec 09, 2019	12	297	Jul 15, 2020	Jul 10, 2020	Jul 25, 2020	22	15,420
Dec 13, 2019	Dec 10, 2019	Dec 16, 2019	10	316	Aug 05, 2020	Jul 26, 2020	Aug 16, 2020	32	34,561
Dec 19, 2019	Dec 17, 2019	Dec 22, 2019	9	334	Aug 27, 2020	Aug 17, 2020	Sep 15, 2020	41	24,705
Dec 26, 2019	Dec 23, 2019	Jan 08, 2020	20	1,354					

Table C.2: Attribution of reviews to dates where listings were scraped

C.3. Gradients

To facilitate stability and speed of convergence, we use analytical gradients for the estimation procedure that is outlined in Section 4.1. We provide our derivations here. To simplify notation, we omit the city index because all expressions stay the same. We consider the more general case with demographic variables D in random coefficients, where idiosyncratic utilities are $\mu_{ijt} = X_{ijt}\big(\Pi D_i + \Sigma\nu_i\big)$, and the non-linear parameters of the model are $\theta = \big(\Sigma, \Pi\big)$. The choice probabilities for consumer i are given by

$$s_{ijt} = \frac{\exp(\delta_{jt} + \mu_{ijt})}{1 + \sum_{k=1}^{J}(\delta_{jkt} + \mu_{ikt})}, \tag{21}$$

and the standard multinomial logit derivatives are

$$\frac{\partial}{\partial \delta_{kt}} s_{ijt} = \frac{\partial}{\partial \mu_{ikt}} s_{ijt} = \begin{cases} s_{ijt}\big(1 - s_{ijt}\big), \ j = k \\ -s_{ijt}s_{ikt}, \ j \neq k. \end{cases} \tag{22}$$

For δ_{jt} and μ_{ijt} defined by θ, we have

$$\begin{aligned} \frac{\partial}{\partial \theta} s_{ijt}(\theta) &= \sum_{k=1}^{J} \Big[\frac{\partial s_{ijt}}{\partial \delta_{kt}} \frac{\partial \delta_{kt}}{\partial \theta} + \frac{\partial s_{ijt}}{\partial \mu_{ikt}} \frac{\partial \mu_{ikt}}{\partial \theta}\Big] \\ &= -s_{ijt} \sum_{k} s_{ikt} \Big[\frac{\partial \delta_{kt}}{\partial \theta} + \frac{\partial \mu_{ikt}}{\partial \theta}\Big] + s_{ijt}\Big[\frac{\partial \delta_{jt}}{\partial \theta} + \frac{\partial \mu_{ijt}}{\partial \theta}\Big]. \end{aligned} \tag{23}$$

As

$$\frac{\partial}{\partial \Pi} \mu_{ijt} = X_{ijt} D_i', \ \frac{\partial}{\partial \Sigma} \mu_{ijt} = X_{ijt} \nu_i', \tag{24}$$

we obtain

$$\frac{\partial}{\partial \Pi} s_{ijt}(\theta) = -s_{ijt} \sum_{k} s_{ikt} \Big(\frac{\partial \delta_{kt}}{\partial \Pi} + X_{ikt} D_i'\Big) + s_{ijt}\Big(\frac{\partial \delta_{jt}}{\partial \Pi} + X_{ijt} D_i'\Big), \tag{25}$$

and

$$\frac{\partial}{\partial \Sigma} s_{ijt}(\theta) = -s_{ijt} \sum_{k} s_{ikt} \Big(\frac{\partial \delta_{kt}}{\partial \Sigma} + X_{ikt} \nu_i'\Big) + s_{ijt}\Big(\frac{\partial \delta_{jt}}{\partial \Sigma} + X_{ijt} \nu_i'\Big). \tag{26}$$

We can apply these expressions[53] to calculate the gradient for our moments

$$m_{jk}(\theta) = \mathbb{E}\Big[R_{ij}R_{ik} \mid R_i > 0\Big] = \frac{\mathbb{E} R_{ij} R_{ik}}{\mathbb{P}\big(R_i > 0\big)}, \tag{27}$$

[53]PyBLP package reports $\frac{\partial \delta_{kt}}{\partial \theta}$.

which equals

$$\frac{\partial}{\partial\theta}m_{jk}(\theta) = \frac{1}{\mathbb{P}\big(R_i > 0\big)^2}\left(\mathbb{P}\big(R_i > 0\big)\frac{\partial\mathbb{E}\big[R_{ij}R_{ik}\big]}{\partial\theta} - \mathbb{E}\big[R_{ij}R_{ik}\big]\frac{\partial\mathbb{P}\big(R_i > 0\big)}{\partial\theta}\right). \tag{28}$$

The terms in this expression can be approximated with averages taken over random draws of agents. In particular,

$$\frac{\partial\mathbb{E}\big[R_{ij}R_{ik}\big]}{\partial\theta} \approx \frac{1}{N}\sum_{i=1}^{N}\frac{\partial}{\partial\theta}\mathbb{E}\Big[R_{ij}R_{ik} \mid i\Big], \tag{29}$$

$$\frac{\partial\mathbb{P}\big(R_i > 0\big)}{\partial\theta} \approx \frac{1}{N}\sum_{i=1}^{N}\frac{\partial}{\partial\theta}\mathbb{P}\big(R_i > 0 \mid i\big). \tag{30}$$

The gradients for the individual product terms are given by

$$\begin{aligned}\frac{\partial}{\partial\theta}\mathbb{E}\Big[R_{ij}R_{ik} \mid i\Big] &= \frac{\partial}{\partial\theta}\Big[\sum_{t_1\neq t_2}\pi_{jt_1}\pi_{kt_2}s_{ijt_1}(\theta)s_{ikt_2}(\theta) + I(j=k)\sum_{t}\pi_{jt}s_{ijt}(\theta)\Big]\\ &= \sum_{t_1\neq t_2}\pi_{jt_1}\pi_{kt_2}\big[s_{ijt_1}(\theta)\frac{\partial}{\partial\theta}s_{ikt_2}(\theta) + s_{ikt_2}(\theta)\frac{\partial}{\partial\theta}s_{ijt_1}(\theta)\big]\\ &+ I(j=k)\sum_{t}\pi_{jt}\frac{\partial}{\partial\theta}s_{ijt}(\theta).\end{aligned} \tag{31}$$

For the probability of observing at least one review by consumer i in the data, which equals

$$\mathbb{P}\big(R_i > 0 \mid i\big) = 1 - \prod_{t=1}^{T}\Big(1 - \sum_{j=1}^{J}\pi_{jt}s_{ijt}(\theta)\Big), \tag{7}$$

we obtain

$$\frac{\partial}{\partial\theta}\mathbb{P}\big(R_i > 0 \mid i\big) = \sum_{t=1}^{T}\Big[\prod_{t'\neq t}^{T}\big(1 - \sum_{j=1}^{J}\pi_{jt'}s_{ijt'}(\theta)\big)\Big]\sum_{j=1}^{J}\pi_{jt}\frac{\partial}{\partial\theta}s_{ijt}(\theta). \tag{32}$$

D. Market Size

As discussed in Section 3.3, we assume that the number of transactions is proportional to the number of listings with the same characteristics. To estimate the market size and the share of the outside option, we need to estimate the corresponding multiplier. There are two important mechanisms that can make the multiplier not equal to 1. First, because deposited dead-drops can stay unsold for several days, a listing observed on a particular day does not necessarily correspond to a transaction on that day. Second, there can be several dead-drops behind one listing. We address this using the following simple framework. Suppose there are L_t listings on the website on day t, among which L_t^{new} are added on that day. Suppose that S_t is the number of sales made on day t. We assume that each listing exists for ω days and there are κ dead-drops behind each listing. For a large T, we can approximate

$$\sum_{t=1}^{T} L_t \approx \omega \sum_{t=1}^{T} L_t^{new}, \tag{33}$$

$$\sum_{t=1}^{T} S_t \approx \kappa \sum_{t=1}^{T} L_t^{new}. \tag{34}$$

We do not observe L_t^{new}, but we can express

$$\frac{\kappa}{\omega} \approx \frac{\sum_{t=1}^{T} S_t}{\sum_{t=1}^{T} L_t} \approx \frac{\sum_{t=1}^{T} S_t}{\sum_{t=1}^{T} L_{\tau(t)}},$$

where we approximate listings on day t by listings on the closest day where scraped data is available. We approximate the numerator by the sum of differences in total sales across all shops over the observed period and obtain $\kappa/\omega \approx 0.7$.

In the Russian mortality data, the majority of deaths associated with drug use occur among individuals aged between 18 and 45. Motivated by this fact, we assume that each person between 18 and 45 can consume drugs 1 time per month. We assume that 1 standard amount is enough to consume drugs 3 times. Thus,

$$N_c = \frac{\omega}{\kappa} \frac{\text{Population between 18 and 45 in } c}{30 \times 3} \approx \frac{\text{Population between 18 and 45 in } c}{65}.$$

Under this assumption, the median market share of the outside option across markets is around 70%.

E. Estimates

Figure E.1. Distribution of own-price elasticities of demand

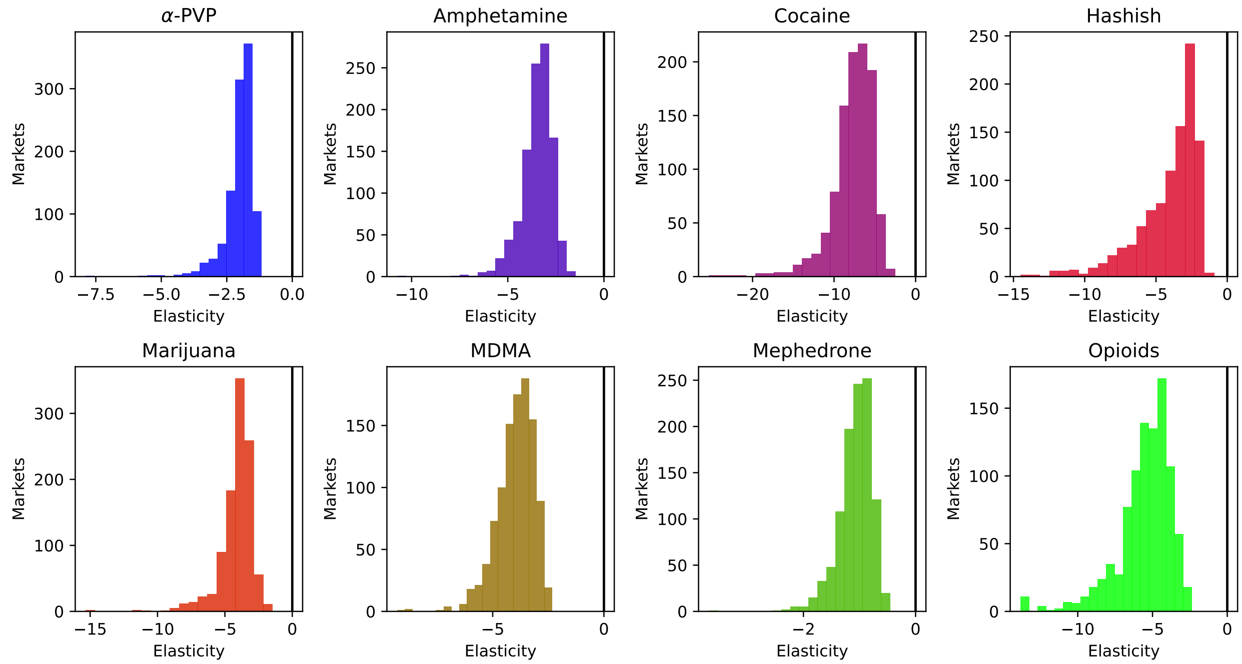

F. Screenshots and Other Materials

To illustrate and support some of the points we make in the paper, we provide several screenshots from the marketplace.

Figure F.1. Front page of Hydra

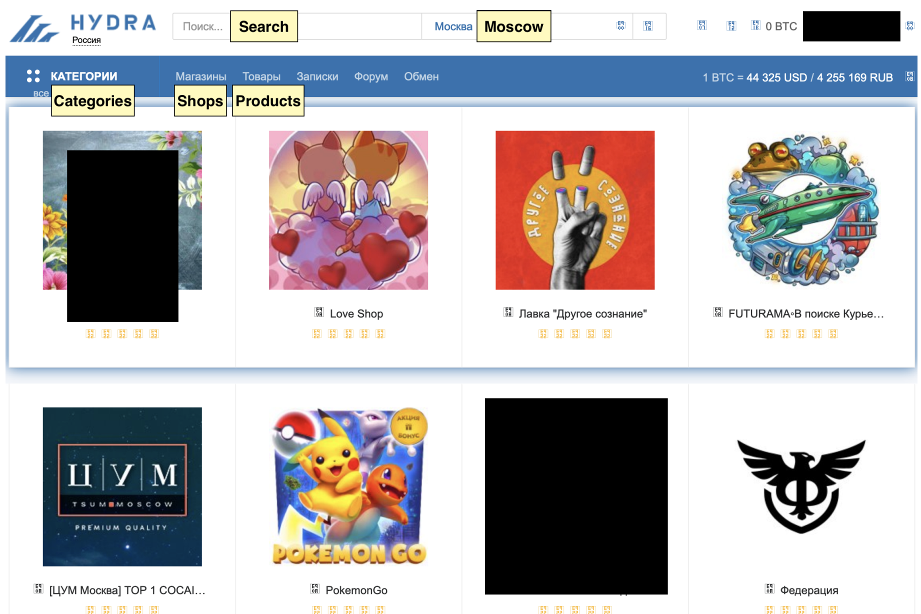

Note: Screenshot from March 26, 2022.

Figure F.2. Example of a product page with cocaine listings

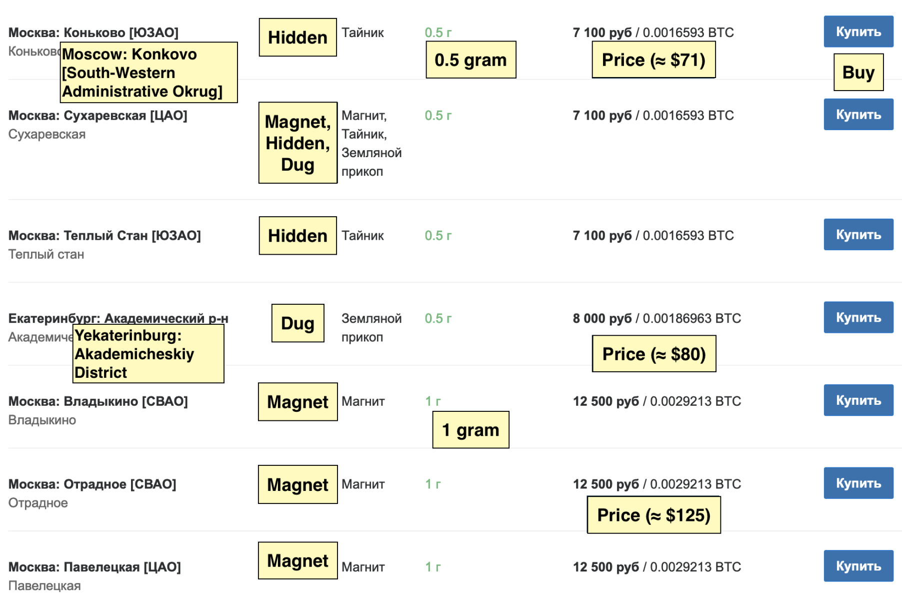

Note: Screenshot from March 26, 2022.

Figure F.3. Example of shop's cumulative number of deals displayed by the platform

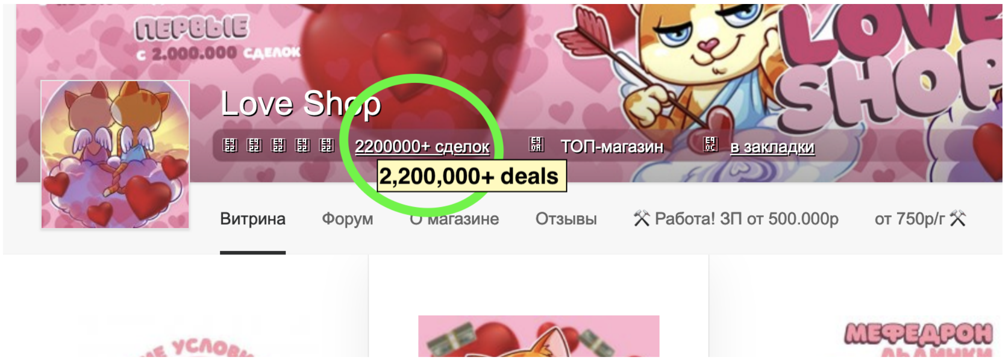

Note: Screenshot from March 17, 2022.

Figure F.4. Example of advertising of a "premium" shop.

Figure F.5. Examples of information provided to buyers

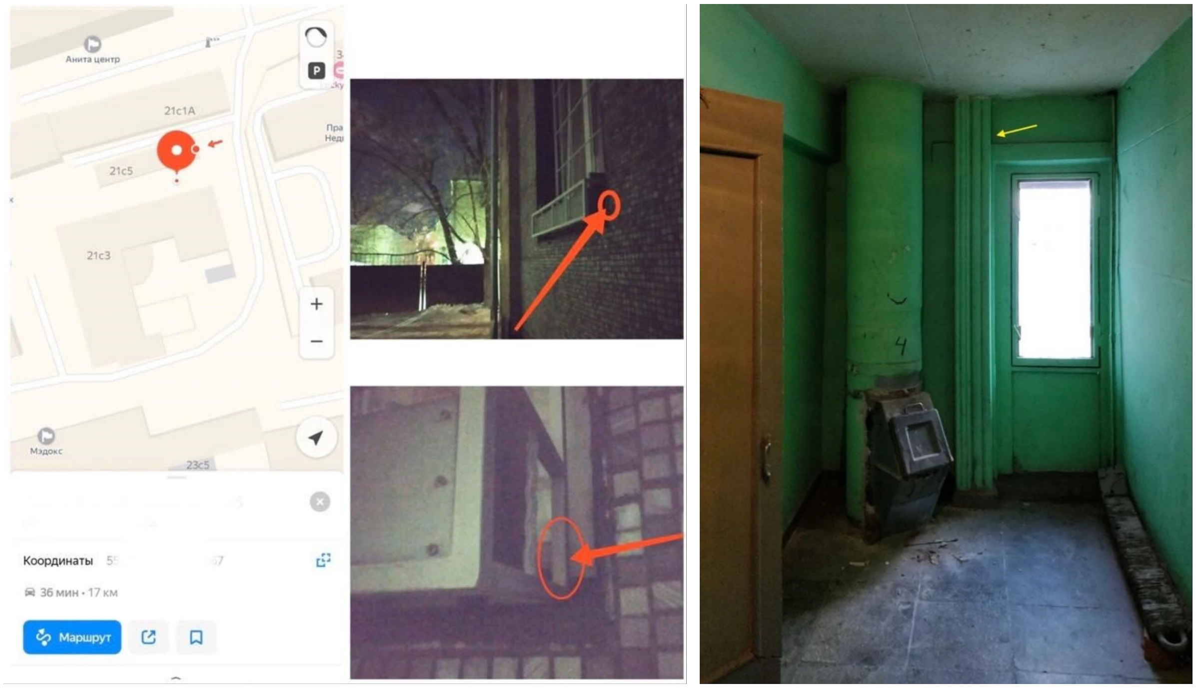

(a) Coordinates and photo of hiding place

(b) Photo of hiding place

Source: VICE.com