MEASURING THE MARKET SIZE FOR CANNABIS: A NEW APPROACH USING FORENSIC ECONOMICS*

MATTHIAS PAREY† IMRAN RASUL‡

MAY 2017

Abstract

Quantifying the market size for cannabis is important given vigorous policy debates about how to intervene in this market. We develop a new approach to measuring the size of the cannabis market using forensic economics. The key insight is that cannabis consumption often requires the use of complementary legal inputs: roll-your-own tobacco and rolling papers. The forensic approach specifies how legal and illegal inputs are combined in the production of hand-rolled cigarettes and cannabis joints. These input relationships, along with market adding-up conditions, can then be used to infer the size of the cannabis market. We prove proof-of-concept that this approach can be readily calibrated using: (i) point-of-sales data on the legal inputs of roll-your-own tobacco and rolling papers; (ii) input parameter estimates drawn from a wide-ranging interdisciplinary evidence base. We then implement the approach using data from 2008-9. For those years, the forensic estimates for the UK cannabis market are near double those derived from standard demand-side approaches. We make precise what drives the measurement gap between methods by establishing: (i) the parameter adjustments needed in demand-side approaches to match the forensic measure; (ii) the changes in methodology to the forensic approach needed to match the demand-side estimate. Our analysis develops an agenda on measurement and data collection that allows for credible cost-benefit analysis of policy interventions in illicit drug markets.

JEL Classification: C80, K42.

*We gratefully acknowledge financial support from the Schoeller Foundation. We thank Oriana Bandiera, Jordi Blanes-i-Vidal, Guido Friebel, Giacomo de Giorgi, Christian Dustmann, Lucie Gadenne, Sebastian Galliani, Andrej Gill, Ali Hortacsu, Michael Kosfeld, Christian Leuz, Steve Machin, Giovanni Mastrobuoni, Martin O'Connell, Olivier Marie, Hannes Mueller, Michael Mueller-Smith, Daniel Nagin, Stephen Pudney, Ernesto Schargrodsky and numerous seminar participants for valuable comments, Elaine Kelly for help with the hospital admissions statistics, Philipp Lohan, Lee Whitlock and those at *Nielsen* that helped provide the sales data. All errors remain our own.

†Department of Economics at the University of Essex, and the Institute for Fiscal Studies. Tel: +44-1206-87-3499; Fax: +44-1206-87-2724; E-mail: m.parey@essex.ac.uk.

‡Department of Economics at University College London, and the Institute for Fiscal Studies. Tel: +44-207-679-5853; Fax: +44-207-916-2775; E-mail: i.rasul@ucl.ac.uk.

1 Introduction

Since Gary Becker's [1968] seminal contribution, the economics of crime has flourished into a mainstream field for economists. Within this vast literature, much research has been devoted to the study of illicit drug markets because they are, directly or indirectly, responsible for a significant share of all crime committed.[1] A body of work has now evaluated the causal impact of various policy interventions in the market for illicit substances [Farrelly *et al.* 1999, Dobkin and Nicosia 2009, Galenianos *et al.* 2012, Adda *et al.* 2014], with particular emphasis placed on policy impacts on the total quantity of illicit drugs supplied and demanded, or the equilibrium price of illicit substances. Such studies provide vital building blocks for any cost-benefit analysis (CBA) of policy interventions. Yet this procedure remains a hugely frustrating process because many of the key *quantities* required to conduct a meaningful CBA in relation to illicit drug market interventions, are often estimated with error or almost entirely unknown.

This paper is borne out of such frustration with the large number of 'known unknowns' in relation to illicit drug markets, that ultimately weakens the ability of social scientists to translate credible policy evaluations in such markets into concrete policy advice. Our analysis begins to fill that void by proposing a novel method to measure the aggregate size of the market for cannabis, using insights from forensic economics [Zitzewitz 2012], and then demonstrating proof-of-concept for the approach using data for 2008-9.

As Figure 1A shows, in the majority of nations, cannabis is the most highly ranked illicit substance in terms of usage. Figure 1B shows prevalence rates by country: between 3 to 5% of the global population aged 15-64 has used cannabis at least once in the past year [UNODC 2011]. Prevalence is particularly high in younger cohorts: among US high school students in their senior year for example, annual prevalence was estimated to be 35% in 2010 [UNODC 2011], and the long term trends in prevalence rates among youth remains the subject of academic debate.

A body of evidence across disciplines has established significant private and social costs associated with the market for cannabis. *Private* costs borne by users take the form of both short-run impacts on educational investments [van Ours and Williams 2009, Marie and Zolitz 2017], longer-term impacts on health from prolonged and heavy use [Fergusson and Horwood 1997, Hall and Degenhardt 2009, Marshall *et al.* 2011], as well as a potentially increased propensity to use other illicit substances [van Ours 2003, Kelly and Rasul 2014]. The social costs of the cannabis market arising through the health and criminal justice systems are substantial, amounting to 1.7% of

[1] The size of illicit drug markets has been linked to crime rates in a variety of countries [Grogger and Willis 2000, Corman and Mocan 2000, Adda *et al.* 2014]. Kuziemko and Levitt [2004] find that incarcerating drug offenders is almost as effective in reducing violent and property crime as locking up other types of offenders.

GDP for the US and UK in 2002 [ONDCP 2004].

The divergence between the marginal social and marginal private costs arising from the market for cannabis implies that all policy makers face a decision over how to intervene in this market. A variety of interventions have been implemented around the world, ranging from policies emphasizing increased deterrence/penalties, to those moving towards decriminalization or depenalization of cannabis use, to legalization of the market as a whole [Donohue *et al.* 2011]. To bridge the gap between academic research and policy advice requires effective CBA to be conducted for any given intervention. The ingredients for such analyses include understanding the behavioral responses to policies of users, non-users, suppliers, law enforcement bodies and other agents, and thus ultimately measuring the overall impacts on the aggregate quantity of cannabis supplied and demanded. However, given the illicit nature of cannabis markets, there remain first order challenges related to the reliable measurement of quantities, parameters and elasticities, that are needed for any credible CBA [Miron 2010, Bryan *et al.* 2013].

This paper begins to address this challenge by developing a new approach to measuring one key statistic for any CBA related to cannabis market interventions: the *aggregate size of the market*, namely the total quantity of cannabis demanded and supplied per time period. This statistic is also of intrinsic value because it: (i) underpins knowledge of the monetary value of the market, and hence the revenues accruing to criminal organizations supplying cannabis; (ii) provides an indication of the tax revenues that might be generated from the legalization of cannabis.

Our method takes ideas from forensic economics [Zitzewitz 2012] in that it exploits data on *licit* markets to provide insights for the measurement of *illicit* activities. More precisely, we use the insight that the consumption of cannabis is often, although not always, combined with the use of legal inputs: these highly complementary legal inputs are roll-your-own loose leaf tobacco and rolling papers. We develop and implement a 'forensic' approach in which we specify how legal and illegal inputs are combined in legal markets (the production of hand-rolled cigarettes) and illegal markets (the production of cannabis joints). This input relationship, along with simple market adding up conditions, can then be combined to make inference on the aggregate size of the market for cannabis. We show our approach in that it can be readily calibrated by utilizing: (i) available sales data on the legal inputs of roll-your-own tobacco and rolling papers; (ii) input parameter estimates that we draw from a wide range of academic disciplines including epidemiology, health policy, medicine, chemistry, psychiatry and pharmacology as well from non-academic reports/expert accounts of judicial bodies, health and drug-related organizations.[2]

[2] Other closely related work in forensic economics include Fisman and Wei [2004] who identify tax evasion in china comparing Hong Kong's exports to China and China's imports from Hong Kong, and Sukhtankar [2012] who identifies embezzlement in politically controlled sugar mills in India exploiting features of the close-to-fixed

We show proof-of-concept for the forensic approach by calibrating the model to two specific years of data (2008, 2009) for which we have obtained UK point-of-sales data for legal inputs of roll-your-own tobacco and rolling papers. Our baseline estimates for the UK cannabis market size are near double those derived from established demand-side approaches, valuing the market at £3 billion. It is thus fundamental to understand what drives the measurement gap arising between methods, for those years. We tackle the issue in two ways: (i) we first ask what would need to be the adjustments made in demand-side approaches to reconcile with the forensic measure; (ii) what adjustments to how the forensic approach is implemented are needed to derive estimates closer to those from demand-side methods. Given the framework underlying the forensic approach, we derive alternative estimates by either changing underlying parameter values, or by modifying the modelling framework itself. This produces a range of estimates under these alternative scenarios. Finally, our forensic approach straightforwardly allows for monthly or regional decompositions of the market size for cannabis. We use these decompositions to validate our approach using an independent data source: administrative records on hospital admissions. These contain ICD-10 codes and record whether cannabis use is listed as a primary or secondary cause of admission. We document that our market size estimates correlate highly with the number of cannabis-related hospital admissions, across UK regions.

The paper is organized as follows. Section 2 describes established demand- and supply-side methods used to measure the size of illicit drug markets. Section 3 develops the conceptual framework underlying a forensic approach. Section 4 describes the data and parameter inputs needed to implement the proof-of-concept for this approach. Section 5 presents our baseline market size estimates for 2008 and 2009, and shows how these estimates can be reconciled to demand- and supply-side approaches. We also conduct a validation exercise to underpin our estimates. Section 6 concludes with a discussion of how best to combine demand-side and forensic approaches going forward. We set out a future research agenda on methods and data collection that ultimately aims to bridge the gap social scientists face between research into policy interventions in illicit drug markets, CBA of those interventions, and informing ongoing and vigorous policy debates.

2 Existing Approaches to Measuring Market Size

There are two long-established methods for measuring the market size for cannabis: demand- and supply-side approaches. We describe each in turn.

proportions technology used in that sector.

2.1 Demand-side

The demand-side approach to measuring the market size for cannabis is best exemplified in Kilmer and Pacula [2009] (henceforth KP). This method starts by establishing prevalence rates from individual survey-based responses. Prevalence rates are calculated separately for 'regular' and 'occasional' users, where such types are distinguished in survey data by whether they report consuming cannabis in the last month or the last year, respectively. These rates then need to be combined with information on the number of days cannabis is typically consumed, and the amount of cannabis consumed per use-day, by user type. KP review the available evidence from a range of sources and argue that while there is relative consensus on the number of days consumed, "The lack of information about typical quantities consumed on a use day (for cannabis and other drugs) severely limits the accuracy of demand-side estimates" (p12). KP's evidence review suggests baseline figures of 2.5 (1.25) joints per day for regular (occasional) users, and $0.4g$ of cannabis per joint for both user types. Combining these figures with prevalence data, the annual consumption per user can be constructed. At a final stage this figure is then combined with population statistics for those aged 15-64, to produce an estimate of the aggregate market size for cannabis.[3]

At least three concerns arise with such demand-side estimates [Kilmer *et al.* 2011]: (i) difficulties in reaching the relevant sub-population of illicit drug users in the original survey data, especially if the primary purpose of the survey is not directly related to the consumption of illicit substances; (ii) selective non-response correlated to the use of illicit substances; (iii) systematic misreporting on the intensive margin of usage.

On the first concern, the British Crime Survey (BCS) remains the main UK-survey eliciting information on drug use. It is based on a nationally representative cross-section of the non-institutionalized population aged 16 and over. However, Mott [1985] argues that under-reporting rates in the BCS are likely to be high given that cannabis consumption is surveyed "in the context of offending rather than of drug use" (p. 37). KP similarly argue that "general population surveys often miss heavy drug users who are in treatment, in jail/prison, in an unstable housing situation, hard to locate, or unwilling to talk about their substance use" (p. 5).[4]

On concerns related to misreporting, attempts to validate the extent of misreporting typically find substantial discrepancies between survey responses and other markers of illicit drug use. For

[3]Other demand-side estimates, such as Pudney *et al.* [2006] do not rely on quantity estimates per joint, but rather focus on quantity consumed per day. They assume intensive users consume $1.2g$ per day of use ($\pm$.3g) and non-intensive users consume $.5g$ per day of use ($\pm$.4g).

[4]Some demand-side estimates for the UK have addressed this concern directly, but these come at the cost of not being so easy to implement across years. For example, Pudney *et al.* [2006] derive an estimate of the UK market size for cannabis by augmenting the standard survey based demand-size method with surveys of juveniles, and adult arrestees, to improve coverage of the relevant population.

example: (i) Harrison *et al.* [2007] compare self-reported consumption over 30 days to biological evidence from urine samples, for a sample of youth in the US–they find among those who test positively in the urine test, 39% do not self-report usage; (ii) Fendrich *et al.* [2004] report 22% of those with a positive biological test do not self-report consumption over the past year; (iii) KP compute a similar degree of under-reporting based on a sample of arrestees from the US Arrestee Drug Abuse Monitoring data. Although such validation exercises leave no doubt on the need to correct survey responses in some way, it remains unclear how to practically implement such corrections when misreporting varies across samples and the perceived confidentiality of responses, the framing of questions, and by the illicit substance in question. In practice, non-response and misreporting is jointly taken into account by re-scaling demands-side estimates by a given proportion. KP argue the norm has become for demand-side estimates to assume a re-scaling factor of 20% to offset these concerns.[5]

Figure 2A shows recent demand-side estimates for the UK cannabis market. The time series shown is based on the application of the standard KP demand-side approach to UK data, where annual prevalence rates are derived from the BCS, and annual population figures are taken from ONS statistics. This time-series shows how the size of the UK cannabis market has fluctuated between 300 and 500 metric tonnes (t) over the period 1994-2012, with there being a slight downward trend in market size between 2003 and 2009, with stable estimates since then.[6] Figure 2A also shows demand side estimates from some notable other studies, that are also listed in Panel A of Table 1. Wherever possible, we also indicate the range of market size estimates provided in these studies. These estimates cluster between $412t$ and $486t$ [Bramley-Harker 2001, Pudney 2006], and match the KP estimates quite closely, despite using variants of the basic demand-side approach. Finally, Figure 2A also shows two market size estimates derived using our forensic approach: these are at least double the demand-side estimates for the same year. We discuss this difference in detail in Section 5, where we show: (i) what adjustments need to be made in demand-side approaches to reconcile with the forensic measure; (ii) what adjustments can be made to the forensic approach to derive estimates closer to those from demand-side methods.

[5] There are alternative approaches to deal with non-response. Pudney *et al.* [2006] impute consumption of illicit drugs by non-respondents based on a set of observables, requiring assumptions about the similarity in behavior between respondents and non-respondents. Jacobi and Sovinksy [2016] present estimates from a structural model that account for non-random selection into consumption.

[6] Estimated prevalence rates in the UK in 2014-15 coincide with those for 2010-11 [Home Office 2015].

2.2 Supply-side

The supply-side approach to estimating the market size for cannabis uses as its core input data on official drug seizures from border authorities and the police. The aggregate market size is calculated by dividing the seized cannabis quantities by some *seizure rate*. In practice, seizure rates are entirely unknown and even hard to narrow down within some bounds because they are endogenously determined by the simultaneous efforts of law enforcement and drug traffickers. Given the lack of direct evidence on seizure rates, the academic literature has been wary of supply-side estimates of the market size for illicit substances [Pudney *et al.* 2006, Kilmer *et al.* 2011]. Indeed, as Kilmer *et al.* [2011, p154] state, "[w]hile this approach is easy to implement, it is unsettling, because no one has a systematic basis for estimating the seizure rate."[7]

Panel B of Table 1 shows current supply-side estimate for the UK cannabis market size. Supply-side estimates tend to be substantially higher than demand-side estimates, and of course vary with the assumed seizure rate. For example, Groom *et al.* [1998] present a supply-side estimate of 1800t assuming a 10% seizure rate. As true seizure rates are unknown, a useful thought experiment is to ask what seizure rate would reconcile the demand-side estimates shown in Figure 2A with the total quantity of seized cannabis. Figure 2B shows the implied seizure rate that reconciles the KP estimates, other demand-side estimates, and our forensic estimates of market size. Two points are of note: (i) reconciling with demand-sized estimates implies seizure rates have declined over time, from 25% in the late 1990s to around 10% by 2010; (ii) given our market size estimate is higher than demand-side estimates, the forensic approach suggests seizure rates are even lower, at around 5% for the two years in which we implement the method.

Taken together, this evidence highlights the importance of establishing accurate methods for estimating the market size for cannabis, and to understand the key sources of discrepancy between demand-side and forensic approaches. This is the focus of the remainder of our analysis.

3 A Forensic Approach

The approach we develop uses the tools of forensic economics in that it exploits data on *licit* markets to provide insights for the measurement of *illicit* activities [Zitzewitz 2012]. The intuition is that while the consumption of cannabis is hard to measure, its use entails the use of highly *complementary* inputs, such as rolling papers and roll-your-own (RYO) tobacco. Crucially, these complementary inputs are *legal* and so reliable sales data exist for them. To formalize this intuition,

[7] Estimating production directly may be possible for some drugs, for example, using aerial or satellite photography. However this method is impractical for cannabis where production is dispersed geographically and much of which occurs indoors [Kilmer *et al.* 2011].

we first assume legal inputs can be put to two potential uses: (i) rolling papers and RYO tobacco can be combined to produce legal hand-rolled cigarettes: we refer to this as Sector 1; (ii) the two inputs can be combined with cannabis to produce illegal joints: we refer to this as Sector 2. To be clear, cannabis is certainly also consumed in other forms (beyond joints). Hence our approach provides a *lower* bound estimate of the total market size. We later discuss modifications to this baseline set-up to account for additional uses of both inputs.[8]

Consider an individual i who participates in the legal Sector 1. This consumer can combine RYO tobacco (denoted as ryo_{1ij} for RYO tobacco type j, measured in units of weight) and rolling papers (denoted as pap_{1ik} for paper type k, measured in units of regular sized rolling papers) to produce legal hand-rolled cigarettes. Denote R_{1i} as user i's input ratio of RYO tobacco to rolling papers, in other words, their tobacco content per hand-rolled cigarette:

$$R_{1i} = \frac{\sum_j ryo_{1ij}}{\sum_k pap_{1ik}}. \tag{1}$$

In Sector 2, a user i combines cannabis, denoted can_{2i}, with RYO tobacco of type j (ryo_{2ij}) and rolling papers of type k (pap_{2ik}) to produce joints. The ratio R_{2i} is the combined quantity of both tobacco and cannabis per joint for individual i. Users may use more than one rolling paper to construct a joint, and so θ_i denotes the number of papers per joint user i uses. Hence,

$$R_{2i} = \frac{\sum_j ryo_{2ij} + can_{2i}}{\sum_k pap_{2ik}/\theta_i}. \tag{2}$$

We denote user i's quantity of cannabis per joint by β_i, resulting in the following input relationship,

$$\beta_i = \frac{can_{2i}}{\sum_k pap_{2ik}/\theta_i}. \tag{3}$$

Within each sector, we assume users behave in a homogeneous way with respect to these input ratios. That is, for users in Sector 1 we assume $R_{1i} = R_1$, and for users in Sector 2, $R_{2i}/\theta_i = R_2/\theta$ and $\beta_i/\theta_i = \beta/\theta$.

We now aggregate the consumed quantities within each sector across users and product types, and define $ryo_1 = \sum_j \sum_i ryo_{1ij}$ and $ryo_2 = \sum_j \sum_i ryo_{2ij}$ as the total weight of RYO tobacco used in each sector, and $pap_1 = \sum_k \sum_i pap_{1ik}$ and $pap_2 = \sum_k \sum_i pap_{2ik}$ as the total number of rolling

[8] Very few reliable data sources exist providing information on the form in which cannabis is consumed. One source is survey data for Ireland from 2002/3: this suggests that among respondents reporting having used cannabis (herbal or resin) in the past month, 98% said that smoking joints was one of the ways they had consumed the substance, with the second most popular method being pipes (7%) [UNODC 2006]. For the US, Schauer *et al.* [2016] find that among those who have ever used marijuana, 89% report having used joints. Among last-month users, around 50% of respondents report last month use of joints (and prevalence of combusted modes of use account for more than 90% of last-month users).

papers used in each sector. Most importantly, $can_2 = \sum_i can_{2i}$, representing the aggregate market size for cannabis.

As detailed below, our data contains detailed information about total sales of RYO tobacco and rolling papers, product-by-product: we denote these total sales of RYO tobacco product j and rolling paper k as ryo_{Tj} and pap_{Tk} respectively. We then have the measurable total sales of rolling paper as $pap_T = \sum_k pap_{Tk}$, summing across product types. As this input can be used in one of the two sectors described above, we have that:

$$pap_1 + pap_2 = pap_T. \tag{4}$$

For the other legal input, RYO tobacco, we denote measured total sales as $ryo_T = \sum_j ryo_{Tj}$, again summing across product types. It is well recognized that there is both a sizeable market in illegal import of RYO tobacco, as well as a significant amount of cross-border shopping of this product into the UK [Cullum and Pissarides 2004, HMRC 2011]. We denote the total quantities of RYO tobacco from both illicit sources as ryo_T^* so that given the uses of RYO tobacco in both sectors:

$$ryo_1 + ryo_2 = ryo_T + ryo_T^\star. \tag{5}$$

To summarize, the forensic approach is encapsulated in the set of input parameters, $\mathbf{\Pi} = (R_1, R_2, \theta, \beta)$, that apply at the level of outputs produced in each sector, and market level adding up restrictions for each legal input that requires knowledge of market level parameters (pap_T, ryo_T, $ryo_T^\star$).

Three further points are of note. First, the cannabis content of joints, β, likely varies with type/strength of cannabis. Indeed, evidence suggests the type of cannabis being consumed in the UK (and US) has slowly changed over the past 20 years, with relatively more consumption now of higher-potency sinsemilla which might reduce the cannabis content per joint [Kilmer *at al.* 2011, Bryan *et al.* 2013]. Hence any attempt to construct a long time-series of the cannabis market size using the forensic approach would need to adjust for these underlying changes in cannabis types being consumed.

Second, this concern is however *common* to both forensic and demand-side approaches: as described earlier, the demand-side approach also requires an assumption on the cannabis content per joint. The constructed time-series in Figure 2A takes the value suggested by Kilmer and Pacula [2009] of $\beta = .4g$, based on their evidence review. This is applied across years and user types (occasional and regular), without accounting for potential changes in the types of cannabis consumed over time. This overlap in input parameters between demand-side and forensic approaches is useful to anchor the comparison between the market size estimates derived from each.

We can then be precise about the difference in market size estimates that arise from differences in methodology rather than parameter inputs, as well as making clear what adjustments to the other inputs in the demand-side approach are needed to reconcile the two sets of estimates.

In short, the contrast between the forensic approach and existing approaches is that: (i) we entirely avoid having to make assumptions on survey non-response and misreporting, that might arise for survey-based estimates; (ii) we avoid making assumptions on seizure rates; (iii) we replace those assumptions with required knowledge of input parameters (R_1, R_2, θ); (iv) the forensic and demand-side approaches overlap in their derived use of parameter β.

The third important point to emphasize is that when describing the input ratios (R_1, R_2) in the two sectors we are *not* claiming a Leontief technology is used in both sectors, that forces inputs to be used in some fixed shares. If that were the case then it would imply these input ratios are insensitive to the relative prices of RYO tobacco, rolling papers and cannabis, that is unlikely to be true. Rather, we emphasize that underlying the forensic approach is some underlying individual utility maximization problem where individuals: (i) make extensive and intensive margin choices over the consumption of hand-rolled cigarettes and joints; (ii) the input ratios to use in each sector depending on relative prices. In general, *all* input parameters depend on a vector of relative prices $\mathbf{p}$ so that for the set of input parameters,

$$\mathbf{\Pi} = \mathbf{\Pi}(\mathbf{p}). \tag{6}$$

Hence as the price of inputs (RYO tobacco, rolling paper, cannabis) change, then input ratios and demands on the intensive and extensive margin will be impacted. Moreover, as the price of substitute/complement goods to hand-rolled cigarettes and joints vary, such as pre-rolled cigarettes or alcohol for example, these technology parameters and underlying demands will also be impacted [Rothwell *et al.* 2015]. However, the focus of this paper is to provide a proof-of-concept that the forensic approach can be implemented: we do so by calibrating the model for the two years of data (2008/9) that we have related to the legal inputs of RYO tobacco and rolling paper products. Over such a short window we treat the price vector $\mathbf{p}$ as fixed, and put to one side issues related to the longer term variation in these input parameters caused by price variation. Of course, these sources of variation open up a rich set of possibilities for future research, and we return to these issues in our conclusions.

Combining the input intensities in (1)-(3) with the market level adding up restrictions in (4) and (5) produces a system of five equations in five unknowns (pap_1, pap_2, ryo_1, ryo_2, can_2). This can be solved to obtain a unique solution for the aggregate market size for cannabis:

$$can_2 = \frac{\beta}{\beta + \theta R_1 - R_2} \left(R_1 pap_T - ryo_T - ryo_T^{\star} \right). \tag{7}$$

Figure 3 illustrates the forensic approach: we represent the available inputs in the form of an Edgeworth box, where the vertical axis represents rolling paper and the horizontal axis represents RYO tobacco. Each point in the box represents a potential allocation of inputs between Sectors 1 and 2, and the origins of the two sectors are located at opposite ends of the Edgeworth box. The rays (with slopes R_1 and $(R_2 - \beta)/\theta$, respectively) represent the input restrictions in each of the sectors (equations (1) and (2)). The intersection of the rays locates the allocation of inputs across sectors that is consistent with the input restrictions. This pins down precisely how rolling papers and RYO tobacco are split between the sectors and thus reveals how these legal inputs are divided between licit and illicit uses. The market size for cannabis can then be directly estimated given the papers allocated to Sector 2, using input restriction (3).

Figure 3 makes clear that a necessary condition to ensure a *unique* market size estimate for can_2 is for the slopes of the rays to differ: $R_1 \neq (R_2 - \beta)/\theta$. Intuitively, this requires the RYO intensities (net of any cannabis input in a joint) in the two sectors must differ. Although in principle either of these intensities could be larger or smaller, our baseline parameters (discussed in the next Section) indicate that RYO intensity is higher in Sector 1, so $R_1 > (R_2 - \beta)/\theta$. Hence we focus our exposition on this case. Figure 3 also makes precise that in order for the solution to be *feasible*, namely to obtain non-negative values for the five unknown quantities, the total market ratio of RYO and paper, $(ryo_T + ryo_T^{\star})/pap_T$, needs to be in the interval $[(R_2 - \beta)/\theta; R_1]$, and Sector 2 intensity R_2 needs to be (weakly) larger than β, that naturally implies the weight per joint is at least as great as the cannabis content per joint.

Figure 3 can be used to see how changes in parameters translate into changes in the market size estimate. For example, an increase in the tobacco content of legal hand-rolled cigarettes, R_1, requires that everything else equal, more rolling papers must then be allocated to the illegal Sector 2: this results in an increased estimate for the cannabis market size, can_2. We summarize these comparative static results as follows,

$$can_2 = \mathbf{C}(\underset{(+)}{pap_T}, \underset{(-)}{ryo_T}, \underset{(-)}{ryo_T^{\star}}, \underset{(+)}{R_1}, \underset{(+)}{R_2}, \underset{(-)}{\theta}, \beta). \tag{8}$$

Only the effect of an increased cannabis content per joint, β, is ambiguous: this is because of two opposing effects: an increased β reduces the overall market size of the illicit Sector 2 (as we hold constant pap_T and ryo_T, but increases the amount of cannabis for a given market size of Sector 2.

In the sensitivity analysis below, we study how reasonable variation in the input parameters (R_1, R_2, β, θ) affects our estimate of aggregate market size for cannabis.

A special case is the scenario in which Sector 2 does not consume any RYO tobacco (so $ryo_2 = 0$, hence $R_2 = \beta$ and all joints are constructed purely from cannabis). As can_2 is increasing in R_2, this represents a lower bound on the size of the cannabis market, all else equal. Intuitively, this is a lower bound estimate because when $ryo_2 = 0$, then holding constant ryo_T, it must be the case that the size of the legal sector is maximized. This bound is given by:

$$\underline{can_2} = \frac{\beta}{\theta R_1}\left(R_1 pap_T - ryo_T - ryo_T^{\star}\right). \tag{9}$$

This case is of further interest because it can be computed even when no information on R_2 is available, relating to input intensity in the illegal sector. In Figure 3, this situation is characterized by all RYO tobacco being used in Sector 1, and the size of Sector 2 being determined from the remaining papers.

4 Data

To calibrate the forensic approach and estimate the aggregate market size for cannabis requires: (i) market level data on two legal products: rolling papers and RYO tobacco (pap_T, ryo_T) and on total illegal imports of RYO tobacco ($ryo_T^{\star}$); (ii) estimates of the input parameters $\mathbf{\Pi} = (R_1, R_2, \theta, \beta)$. We discuss each in turn.

4.1 Market Level Parameters

We obtain information on the legal market level factors based on a *Nielsen* panel data set on point-of-sales information covering tobacco and related products in the UK. This data covers the near universe of sales outlets, accounting for 94% of all UK sales of such products. Hence, in sharp contrast to demand-side survey estimates, these sales data achieve a high level of coverage and are not subject to concerns over non-response or misreported sales. This represents a major advantage of the forensic approach over the demand-side approach.[9]

The *Nielsen* point-of-sales data are recorded by product based on unique product barcodes. The data cover 72 unique loose-leaf or RYO tobacco products (j) and 86 unique rolling paper products (k), for each month of sales from January 2008 to December 2009. Each month's data is

[9] The *Nielsen* sales figures exclude small tobacco kiosks, motorway service stations, prisons, and army outlets. To account for the 6% of sales that originate from such outlets, we scale up sales figures by 1/0.94 for both rolling paper and RYO tobacco.

broken down into 11 UK regions, corresponding to 'TV region equivalents' which are regularly used for *Nielsen*'s marketing analyses. We aggregate across products, regions and months to construct total annual UK sales for the legal inputs of rolling paper and RYO tobacco (pap_T, ryo_T). For rolling paper products, the data distinguishes between different sizes/types of rolling paper k (e.g. regular, King Size, King Slim Size). We normalize paper sizes so that our unit of measurement is regular length papers. For RYO tobacco, the market quantity is measured in metric tonnes.[10]

The quantity of illegal imports and cross-border shipments of RYO tobacco into the UK ($ryo^{\star}_T$) is derived from official statistics from the UK tax authority [HMRC 2011]. HMRC estimate illegal imports of hand-rolling tobacco from constructing the tax gap for this product, so the difference between the tax revenue expected from total consumption minus the tax revenue actually obtained. Total consumption of RYO tobacco ($ryo_T + ryo^{\star}_T$) is based on survey-responses from the General Lifestyle Survey (GLS) on the quantity of hand-rolled cigarettes consumed (but no information on the consumption of joints is available), while HMRC directly calculate the total tax revenues from legal sales of RYO tobacco (ryo_T), taking into account duty-free and EU duty paid, as well as cross-border shopping [HMRC 2010, Table 4.8]. Given the GLS does not ask about the use of RYO tobacco in joints (ryo_2), this procedure likely underestimates $ryo^{\star}_T$. As (8) shows, this leads us to potentially overestimate the market size for cannabis. To address this concern, we later also construct market size estimates using HMRC's published upper bound value for $ryo^{\star}_T$.

4.2 Input Parameters

To derive the input parameters (R_1, R_2, β, θ) we draw on two types of evidence base. We first examine academic research papers in the fields of epidemiology, health policy, medicine, chemistry, psychiatry and pharmacology. We complement this with parameter estimates from factual expert accounts/case studies in judicial studies, and the reports of health and drug-related organizations. The full set of sources used are in Appendix Tables A1-A4, where each corresponds to direct estimates of the input parameters (R_1, R_2, β, θ).[11]

For the legal Sector 1, Appendix Table A1 lists studies providing estimates of the RYO intensity for hand rolled cigarettes, R_1. The estimates are spread over the range $.44g - .92g$, with a mean (median) of $.65g$ ($.67g$). For the illegal Sector 2, there are fewer studies that estimate R_2, the combined weight of cannabis and RYO tobacco per joint. However, the available evidence,

[10] We normalize King Size and King Slim Size to regular sized papers. To do so we use area measurements based on the leading UK brand of rolling papers.

[11] We make no attempt to weight studies according to their authority or sample size for example. To do so lies beyond the scope of our analysis. We also do not claim that our list of sources is comprehensive: the sensitivity analysis allows us to examine how our market size estimate changes over plausible ranges of input values that we have been able to find.

summarized in Table A2, provides a very consistent picture: all five studies imply R_2 is close to $1g$ per joint (the estimates range between $.89g$ and $1.09g$).[12]

There is a substantial body of evidence measuring cannabis content per joint (β). Table A3 reports these studies, grouped by cannabis type (resin, herbal, other). For studies that report ranges, we take the midpoint. Where we list several values for the same study, we include the average over these values. From the 24 sources listed, the mean (median) value is $.36g$ ($.37g$).

Our baseline estimates for (R_1, R_2, β) take the mean values derived from this evidence base.

As discussed above, β is also a key parameter for demand-side estimates: in their evidence review, Kilmer and Pacula [2009] take $.4g$ as preferred value, and also consider the $.3g - .5g$ range. Remarkably, the preferred estimate taken in demand-based approaches lines up very closely with our summary of the evidence ($.36g$ versus $.4g$), despite being based on different sources of underlying evidence. We take this overlap in values are reassuring and potentially strengthening the credibility of some of the other input parameter estimates required to calibrate the forensic approach. As discussed earlier, Table A3 shows the values for β do vary by cannabis type (being generally smaller for cannabis in resin form relative to being in herbal form). Without further information on the composition of cannabis being consumed over time, neither the demand-side nor our approach can utilize these differences.[13]

The final technology parameter is the number of papers per joint, θ. The academic literature has focused far less on this parameter, partly because the number of papers is not directly linked to the health effects of smoking. The two studies we have located on this parameter are a survey of users and one expert statement, and they are not entirely clear on the size of rolling paper they refer to. We thus take take an intermediate value of 1.98 regular sized rolling papers per joint (corresponding to one King Size paper). For each market level factor (pap_T, ryo_T, $ryo_T^{\star}$) and input parameter (R_1, R_2, β, θ) we show the sensitivity of our market size estimate to unilaterally changing its value holding all other inputs constant.[14]

Two further points are of note. First, as discussed earlier, we expect the input parameters $\mathbf{\Pi}(.)$ to generally vary over time because of: (i) changes in the types of cannabis consumed over time; (ii) changes in the relative price of legal and illegal inputs used in Sectors 1 and 2, as well as changes in prices for products that might be substitutes/complements for hand-rolled cigarettes and joints. Both factors might indeed be able to explain some of the variation in the parameter

[12] Two of the studies do not report total weight, and so we use our β estimate to infer the quantity of interest.

[13] Trautmann *et al.* [2013] report the BCS 2009/2010 did include questions on the use of cannabis types: marijuana, skunk, hash and hash oil. This showed that 71% of cannabis users had used herbal cannabis, 38% took hash, 6% hash oil and 6% did not know which type they had used. 29% of users had consumed multiple types.

[14] Tables A1-A4 use the evidence base reporting the input parameters described in the forensic approach. Table A5 provides other estimates from the literature, for which the relevant input parameters Π can be inferred.

values reported from the evidence review in Tables A1 to A4 (that cover a wide time period). To assess whether input parameters might be changing over time, we present baseline estimates for the cannabis market size using the average input value derived from using all the studies reported, and also just focusing on those studies using data from 2000 or later.

Finally, we note that our baseline parameter estimates imply the conditions for uniqueness and feasibility of (8) are satisfied. Our baseline estimates imply the two sectors use different quantities of smoking material per paper $(R_1, \frac{R_2}{\theta})$: Sector 1 uses $.65g$ per paper while Sector 2 uses $1/1.98 = .51g$ per paper. This confirms we are in the case where $R_1 > (R_2 - \beta)/\theta$ as discussed in Section 3, and suggests that for a given number of rolling papers, cannabis does not crowd-out RYO one-to-one. The feasibility conditions are also satisfied as $R_2 = 1 > \beta = .36$. On the uniqueness constraint $R_1 > (R_2 - \beta)/\theta$, given our baseline values for (R_2, β, θ) this will hold whenever $R_1 > (1 - .36)/1.98 = .323$. We note than none of the estimates of R1 reported in Table A1 violate this restriction, even those based exclusively on UK data sources. However, because $(ryo_T + ryo^{\star}_T)/pap_T = .574$, when only using the three UK-based studies for R_1, this violates the need to be in the interval $[(R_2 - \beta)/\theta; R_1]$ even though the interval itself exists. This violation occurs because the UK-based estimates for R_1 are all at the low end of the range reported in Table A1. This might arise because two of the three UK studies use small laboratory samples. Clearly, more work needs to be done to establish a UK-specific evidence base.

5 Proof-of-Concept Results

5.1 Baseline Measures

Table 2 provides our baseline proof-of-concept forensic measure of the market size for cannabis. For the input parameters (R_1, R_2, β), we take the mean values derived from our evidence base, and for θ we use a value of 1.98 regular sized rolling papers per joint. These input values are shown in Columns 1-4. The market level parameters used are shown in Columns 5-7. Inserting all these into (7) provides our first estimate for the size of the UK cannabis market (averaged over 2008/9). This is reported in the first row of Table 2: the market size estimate is 734.8 metric tonnes for 2008/9. As Figure 2A shows, this is nearly double estimates based on the demand-side approach for the same period. Given the average price of cannabis in 2009 is around £4 per gram [Bryan *et al.* 2013], our estimate suggests the annual revenues generated in the market for cannabis to be close to £2.94 billion. The sheer scale of economic activity taking place in this market, irrespective of any additional private and social costs associated with the market, highlights the importance for policy of accurately measuring the market size.[15]

As discussed above, a concern is that the input parameters (R_1, R_2, β, θ) might be changing over time. To assess the issue, the second row shows our derived market size estimate using only evidence published since 2000 or later (again taking the mean values of (R_1, R_2, β) in this subset of studies and retaining $\theta = 1.98$). This results in small upward adjustments of R_1 (the tobacco content of legal hand-rolled cigarettes) and β (the cannabis content of joints). The derived value of R_2 remains unchanged at 1 (the total content of joints) because as noted earlier, the evidence base for this input is more consistent than for other input parameters. The resulting cannabis market size estimate is 873.2 metric tonnes for 2008/9 (so 19% larger than the earlier estimate). We take this as our preferred value for the remainder of the discussion.[16]

Two further points are then of note. First, given a closed form solution for the market size measure in (8), we can derive the *elasticity* of the market size with respect to each input parameter (R_1, R_2, β, θ) as well as the market level parameters (pap_T, ryo_T, $ryo^{\star}_T$). Figure 4 shows the implied elasticities: our market size measure is particularly responsive to changes in R_1, the input ratio in the legal Section 1. The elasticity is actually larger with respect to this input ratio than with respect to the parameters directly relating to the illicit Sector 2 (R_2, β, and θ). This is because of the quantitative dominance of Sector 1: our baseline estimate implies that the illicit Sector 2 accounts for 19% of the market for RYO tobacco, and 35% of the (size-adjusted) market for rolling papers. Hence in our method, understanding behavior in Sector 1 is crucial for making inferences about the illegal Sector 2. This counts as a practical advantage in methodology relative to demand- and supply-side approaches, as behavior in the legal Sector 1 behavior is relatively easy to observe and measure.

Second, a novel insight that arises from the forensic approach is that taxing legal inputs that are complementary to cannabis use can be a policy instrument through which to regulate the overall market size for cannabis. As we discuss in more detail in the final Section, the *Nielsen* sales data actually includes product prices, but with only two years of data and no regional variation in UK tax rates, it is not possible to use this to precisely estimate the cross price elasticity of cannabis with respect to legal inputs, although this would be possible with a longer time-series. However, there are a number of caveats to using this policy instrument: (i) it directly impacts the size of the legal Sector 1; (ii) it might cause cannabis users to switch to non-joint forms of consumption;

[15] Bryan *et al.* [2013, p12] report time series evidence on cannabis prices in England and Wales, based on data from Drugscope and the Independent Drug Monitoring Unit. Prices vary across types and quality of cannabis, with the £4 price per gram being an average across these types.

[16] Focusing on an even narrower time period generates very similar estimates. For example, basing our estimate on studies published around a ±3 year window from 2008/9 implies $R_1 = .66$ (based on 10 studies), $\beta = .40$ (based on 9 studies) and the derived value of R_2 remains unchanged at 1. The resulting market size estimate is $875t$, almost identical to the estimate based on all post-2000 studies.

(iii) the total price per joint is largely made up of the price of cannabis rather than the price of legal inputs, so this cross price elasticity might be relatively low.[17]

Our baseline estimate is more than double the corresponding demand-side estimate derived from the methods in Pacula and Kilmer [2009] for the same years ($380.5t$). Hence despite both approaches using having similar values for the one parameter in common β ($.36g$ versus $.4g$), they deliver very different measures of the market size for cannabis. It is fundamental to understand what drives the measurement gap arising between the methods. We tackle the issue in two ways: (i) we first ask what would need to be the adjustments made in demand-side approaches in order to reconcile with the forensic measure; (ii) what adjustments to how the forensic approach is implemented could be made to derive estimates closer to those from established demand-side methods. The next two subsections discuss each in turn.

5.2 Reconciling Demand-side Measures with the Forensic Measure

The demand-side approach uses as input values, separately for regular and occasional users: (i) prevalence rates (denoted δ_1); (ii) days cannabis is consumed per year (δ_2); (iii) joints consumed per day (δ_3); (iv) cannabis content per joint (denoted β as in the forensic method); (v) a re-scaling factor to offset non-response and misreporting (ρ). The approach also uses information on population shares of those aged 15-64, but this derives from census data and we assume those figures are not the source of discrepancy in methods. We therefore focus on the demand-side parameters ($\delta_1, \delta_2, \delta_3$) as elicited from survey data, and the re-scaling factor ρ.

Column 1 of Table 3 shows the parameter values ($\delta_1, \delta_2, \delta_3, \beta, \rho$) used in the KP application of the demand-side approach, split by regular and occasional users. Column 2 then shows the change needed if *each and every* one of these parameters were to be increased in the same proportion, to match the baseline forensic measure of the market size of $873.2t$. We see that if each parameter were increased by 21%, the demand-side approach generates a market size estimate matching the forensic estimate.[18]

In concrete terms, the reconciliation shown in Column 2 implies that if prevalence rates rose from 4.6 to 5.6% (7.9 to 9.6%) for regular (occasional) users, the number of days joints were consumed for regular (occasional) users rose from 150 to 183 days (30 to 36 days) etc. then the two estimates would be reconciled. The under-reporting rate would need to rise from 21% to

[17] Deriving sales-weighted prices for rolling papers and RYO tobacco in our data and combining this with cannabis price data from Bryan *et al.* [2013], we find that in the illegal Sector 2, 91.7% of the value of a joint derives from the cannabis content, 7.6% from the RYO tobacco and .7% from the rolling paper.

[18] Note that a 21% increase in each parameter results in a more than 21% increase in the demand-side estimate because these inputs are multiplied together when calculating the market size.

24% in this thought experiment. Are such increases plausible? Studies validating the extent of misreporting in demand-side estimates find rates of under-reporting on the extensive margin of cannabis to vary between 20 and 40% based on the comparison between self-reported usage and biological markers of usage [Fendrich *et al.* 2004, Harrison *et al.* 2007, Kilmer and Pacula 2009]. Hence the implied increase in ρ from 21% to 24% is entirely plausible.

A related thought experiment to reconcile demand-side and forensic measures is conducted in Column 3: this repeats the analysis allowing only $(\delta_1, \delta_2, \delta_3, \rho)$ to vary, holding β fixed because demand-side and forensic approaches use similar imputed values for β. The result shows that if the discrepancy arises from this subset of parameters, then we need each to increase by 30% in order to reconcile the measures. In this scenario, the implied increase in ρ from 20% to 26% remains entirely plausible given the evidence from validation studies.

The final set of thought experiments are in Columns 4-8 where we consider changing one parameter in $(\delta_1, \delta_2, \delta_3, \rho)$ at a time, and then derive how large that univariate change would need to be to reconcile the market size measures. Hence those values framed in a box in Columns 4-8 are those that change in each thought experiment, all other inputs take the same value as the baseline demand-side estimate in Column 1. This exercise shows that if the extent of misreporting and under-reporting were the *sole* source of the discrepancy, the re-scaling factor ρ would have to more than treble, rising to 65% to reconcile the two sets of estimates. This seems implausible given the validation studies above do not suggest this degree of bias in self-reports. If the sole source of the discrepancy was either due to only prevalence rates (δ_1) (non-response on the extensive margin of cannabis use), or due only the number of days per year on which joints are consumed (δ_2) or the number of joints consumed per day (δ_3) (relating to misreporting on the intensive margin), then each parameter would need to rise by a factor of 2.3 to reconcile the estimates.[19]

5.3 Reconciling Forensic Measures with the Demand-side Measure

Table 4 summarizes how the forensic approach can be modified so the implied measure of the cannabis market size converges to that derived from the demand-side approach. Given the framework underlying the forensic approach, we can derive alternative market size estimates by changing underlying parameter values (R_1, R_2, β, θ), or by modifying assumptions about the modelling framework itself.

[19] For completeness we note that if the cannabis content per joint β were for some reason the source of discrepancy, this would have to rise by a factor of 2.3 from $.4g$ to $.92g$ per joint. Given the independent evidence reviews conducted by Kilmer and Pacula [2009] and our review as summarized in Table A3 (in which the highest value of β is $.62g$), this is implausible (even more so because the evidence suggests β has likely fallen over time as cannabis content has changed over time).

To begin with, Rows 1 and 2 show our earlier market size estimates based on the forensic approach, using the entire evidence base (Row 1) or restricting parameter values to be derived from studies from 2000 onwards. Columns 1-8 are as described earlier for Table 2. Columns 9 and 10 make precise the link between the forensic, demand- and supply-side estimates: (i) Column 9 shows the implied increase needed in each parameter of the demand-side method $(\delta_1, \delta_2, \delta_3, \beta, \rho)$ to reconcile with the forensic estimate; (ii) Column 10 shows the implied seizure rate that would reconcile the forensic and supply-side estimates given official data on aggregate drug seizures. We see that our baseline forensic estimate in Row 2 would match demand-side estimates for the same period if each parameter in $(\delta_1, \delta_2, \delta_3, \beta, \rho)$ is scaled up by 21%. Alternatively, our baseline forensic estimate would match supply-side estimates assuming a seizure rate of 5%.

The first variation we consider is to use median rather than mean values from the evidence base for the parameters (R_1, R_2, β), while still setting $\theta = 1.98$. Taking median values increases R_1 and β slightly (the tobacco content of legal hand-rolled cigarettes and the cannabis content per joint respectively), while R_2 remains with the same value. Given the comparative statics shown in (8), Rows 3 and 4 show that as expected this causes the forensic estimates to become *larger* (where in Row 4 we again restrict to studies based on 2000 or later data), and thus can only be reconciled with even greater proportionate changes in the demand-side parameters $(\delta_1, \delta_2, \delta_3, \beta, \rho)$.

To get a sense of how much the input parameters would need to deviate from their mean/median values to match the demand-side estimate, Figure 5 uses a Tornado diagram to show the cannabis market size estimate resulting from a $\pm 10\%$ univariate change in an input parameter (R_1, R_2, β, θ) or the quantity of illegal imports and cross-border shipments of RYO tobacco into the UK ($ryo^{\star}_T$). The lower axis shows the resultant market size estimate, and compares this to the standard KP measure as well as the other demand-side measures listed in Table 1. The upper axis shows the implied seizure rate that allows the market size measure to reconcile with supply-side measures.

Two results follow. First, the market size estimate is relatively insensitive to $\pm 10\%$ univariate changes in R_2, β, θ, and $ryo^{\star}_T$: this was previously highlighted in the elasticity estimates and Figure 5 confirms that univariate changes in these parameters are unlikely to close the gap between forensic- and demand-side market size estimates. Second, as with the earlier elasticity analysis, the forensic measure is most sensitive to changes in R_1. Indeed, we now see that a 10% reduction from its mean in R_1 would reduce the market size estimate to *below* that suggested by demand-side studies. Such a 10% reduction in R_1 lies within the bounds of estimates in the evidence base: Table A1 shows estimates of R_1, the tobacco content of hand-rolled cigarettes, ranging from $.44g$ to $.92g$ (where the mean value was $.661g$). As we discuss in the final Section, obtaining more precise estimates of R_1 is an important next step for the development and further application of the forensic approach.

The next set of adjustments take alternative estimates of the quantity of illegal imports of RYO tobacco into the UK ($ryo_T^\star$): as discussed in Section 3, this market level parameter is based on official statistics from the UK tax authority (HMRC) but is likely to be downwards biased. Row 5 in Table 4 shows how the forensic measure of the cannabis market size changes if we take the *upper* bound value for $ryo_T^\star$ from HMRC. Doing so causes the forensic estimate to fall to $449.8t$, and so be more closely aligned to the standard demand-side estimate of $380.5t$, and indeed to overlap with some modifications of the KP demand-side approach as reported in Table 1. More precisely, this estimate can be reconciled with the standard demand-side estimate if each and every one of the demand-side parameters $(\delta_1, \delta_2, \delta_3, \beta, \rho)$ is scaled up by 4% (well within plausible values), and implies a seizure rate of 11% in order to be consistent with supply-side estimates.[20]

The next modification involves a departure from the modelling framework in Section 3: the basic model assumed rolling papers and RYO tobacco are used in either Sector 1 or 2. However, it is possible that some of these legal inputs are put to other uses. For example, some fraction of rolling papers and RYO tobacco might simply be wasted or lost. Alternatively, joints might also be produced combining rolling papers with tobacco from pre-rolled cigarettes. We can straightforwardly re-construct market size estimates in such scenarios when there is wastage from both Sectors 1 and 2, by introducing a generic third sector (Sector 3), to which we allocate a particular share of overall sales volume for each legal input.

Rows 6 and 7 in Table 4 show how the forensic market size measure changes as we assign different percentages of these legal inputs to Sector 3. Row 6 shows that if we assume 5% of pap_T and 1% of ryo_T are used in Sector 3, the forensic market size estimate falls to just under $600t$, so still 58% larger than the demand-side estimate. This estimate can be reconciled with the standard demand-side estimate if demand-side parameters are scaled up by 11%, and implies a seizure rate of 8% to match with supply-side estimates. As Row 7 shows, assuming an alternative ratio so that a greater fraction of ryo_T is used in Sector 3 relative to the fraction of pap_T used, the forensic market size estimate increases from its baseline value to $1094.7t$.[21]

[20] For completeness we note that using the lower bound estimate from HMRC on $ryo_T^\star$ causes the forensic market size estimate to increase to $1310.7t$, so 50% larger than the preferred baseline forensic measure reported in Row 2. We have based our input values for RYO tobacco on the *Nielsen* sales data. HMRC provide an alternative figure from the duties collected on hand-rolled tobacco. Using the HMRC figures for the 2008/9 and 2009/10 tax years, increases the average RYO quantity over the relevant period by about 8% (3% when taking into account illegal imports), and the resulting quantitative estimate of the market size for cannabis falls by 21% to $686t$. This estimate remains more than 80% higher than the baseline demand-side based estimate, and is reconcilable with supply-side estimates assuming a seizure rate of around 7%.

[21] In scenarios where both legal inputs have the *same* fraction allocated to Sector 3 have relatively minor impacts on the quantitative estimate of the market size for cannabis. This is because the overall market ratio of RYO tobacco to rolling papers stays unaffected in these scenarios.

We next consider the bound for the market size given in (9), when we assume Sector 2 uses no RYO tobacco and so all joints are constructed purely from cannabis. Row 8 shows that making this assumption lowers the forensic measure to $474t$, that is still 26% larger than demand-side estimates. The estimate matches the demand-side measure if the demand-side parameters are scaled up by 5%, and implies a seizure rate of 10% to be consistent with supply-side estimates.

Our baseline approach assumes all forms of rolling paper are utilized in Sectors 1 and 2. We can relax this assumption and exploit the fact that rolling papers vary in their size and thickness (regular, king size, king slim size). We thus examine how estimates of the market size vary under alternative assumptions about the use of specific type of rolling paper k in the illicit Sector 2. We first assume rolling papers of type 'king size' are assigned exclusively to Sector 2, and *all* other rolling papers are used in Sector 1. Then equation (3), together with our baseline values for β and θ, provides an estimate for the market size for cannabis. Doing so results in a reduced market size of $713t$, corresponding to 82% of the baseline forensic measure reported in Row 2. If we go one step further and assume rolling papers of type 'slim king size' are assigned exclusively to Sector 2, and *all* other rolling papers k are used in Sector 1 then we have a resulting forensic estimate of $384.5t$, that actually matches estimates from existing demand-side studies. In other words, the two approaches can be reconciled if we make this alternative assumption on precisely which rolling papers products are exclusively assigned to the illicit sector.

All the various adjustments gone through in Table 4 are summarized in Figure 6. This reiterates that in most modifications to the forensic method, be they in terms of changing underlying parameter values for (R_1, R_2, β, θ), or the market level input for $ryo^{\star}_{T}$, or by modifying the modelling framework itself, the resulting estimates are larger than corresponding demand-side estimates for the same period. However, there are some scenarios in which the two approaches generate near identical market sizes. This emphasizes that one way to move forward is to optimally combine both approaches, in particular, to add questions to the individual surveys from which demand-side estimates are derived. The most vital additional questions that are needed are those relating to the tobacco content of hand-rolled cigarettes (R_1) that forensic estimates are most sensitive to, and those that help pick out a scenario from those shown in Figure 6, especially related to the use of *types* of rolling paper (size and thickness) when constructing cannabis joints.

5.4 Validation

Our final set of results are designed to underpin the plausibility of the forensic estimates we derive for the cannabis market size in 2008/9. We do so using two strategies: (i) decomposing market size estimates by month and region; (ii) using administrative records on hospital admissions to examine the correlation between the aggregate number of cannabis-related hospital admissions, and our regional estimates of the cannabis market size.

An important property of the forensic method is that because it is based on detailed point-of-sales data for legal inputs, it is straightforward to derive estimates of the cannabis market size decomposed by month and region. The monthly estimates allow us to conduct a useful validation exercise because it is well established that many forms of crime, including property theft and violent crime, exhibit robust *seasonal* patterns across countries and time periods, and that there is a link between the size of illicit drug markets and non-drug crime [Grogger and Willis 2000, Corman and Mocan 2000, Adda *et al.* 2014]. Survey-based data are often unsuited to address the issue because interviewing periods are concentrated over particular months of the year. Figure 7 shows the forensic estimates of the UK cannabis market size by month, averaged over 2008/9. This reveals variation in market size over the year: the first quarter shows particularly high values, as do the summer months of June to August. These variations should be investigated further as this evidence provides potentially novel policy insights for decisions over the *intertemporal* allocation of resources to regulate illicit drug markets or to treat drug users. It also highlights the importance for demand-side estimates to correct for potential seasonality bias.

We next use the forensic method to derive estimates of the cannabis market size by UK region. Again, such decompositions are difficult to investigate using either demand- or supply-side approaches. For demand-side methods this reflects the prohibitive costs of collecting large enough samples that are *representative* of regions. For supply-side studies the constraint preventing regional estimates being constructed is that seizures are concentrated at border crossings rather than in the region of consumption or supply. However, when decomposing forensic estimates by region, one caveat is that the market parameter related to illegal RYO imports ($ryo^{\star}_{T}$) is only available at a national level: we therefore assume it is distributed across regions in proportion to observed RYO sales in the *Nielsen* data. Doing so we find a regional cross-section of market sizes according with intuition: (i) the London region is the largest single market, accounting for 32% (35%) of the nationwide market in 2008 (2009); (ii) market size estimates are reassuringly stable within region over the two years of data (the R-squared of a regression of market size against region fixed effects is above .99).

Building on these regional estimates, we can further validate our market size estimates using an alternative data source. More precisely, we use administrative records from the National Health Service to measure cannabis-related hospital admissions in England, aggregate these to the same regions as in the *Nielsen* data, and then compare these hospitalization numbers to our estimates of the aggregate market size for cannabis. Data on hospital admissions are drawn from the *Inpatient*

Hospital Episode Statistics. These provide an administrative record of every inpatient health episode, defined as a single period of care under one consultant in an English National Health Service hospital. As such, they constitute the most comprehensive data source on health service usage for England, and have been used in earlier work on illicit drug markets in the UK [Kelly and Rasul 2014]. Inpatients include all those admitted to hospital with the intention of an overnight stay, plus day case procedures when the patient is formally admitted to a hospital bed. As such, these records cover the most serious health events. For each patient-episode event, the data record the date of admission and ICD-10 diagnoses codes in order of importance (ICD codes being the international standard diagnostic classification for epidemiological and clinical use). Based on ICD-10 classifications, we calculate the total number of hospital admissions where cannabis usage is listed as either a primary or secondary cause of admission.[22]

Figure 8 shows a scatter plot of our regional estimates of the *total* market size for cannabis against the *total* number of cannabis-related hospital admissions in the region, where we use the average over 2008/9 for both. The two series are closely related: their correlation coefficient is .69. Overall, this validation exercise compares favorably to similar exercises conducted for demand-side estimates [Fendrich *et al.* 2004, Harrison *et al.* 2007, Kilmer and Pacula 2009], and suggests our method for measuring the aggregate demand for cannabis matches other policy-relevant markers of the size of this illicit drug market.[23]

6 Discussion

Measurement has always been central to economics. Measurement issues remain critical in the context of illicit drug markets, that by their nature do not lend themselves to being quantified using conventional tools. We have developed a novel approach to measuring the aggregate demand for cannabis, the most widely consumed illicit drug on Earth. Quantifying the market size for cannabis is a vital input for any cost-benefit analysis of policy interventions in this illicit market, that remain the subject of enormous amounts of policy discussion [Miron 2010, Bryan *et al.* 2013]. This statistic is also of intrinsic value because it: (i) underpins knowledge of the monetary value

[22] Hospital admissions record the Local Authority District (LAD) of admission. To link hospital admissions locations to the corresponding TV region, we assign LADs to TV regions using data and maps from ONS [2014]. This assignment is proximate because TV region borders do not follow exactly along LAD boundaries. However this measurement error is unlikely to impact our conclusions because large cities are typically contained within a TV region. Indeed, our results are robust to using an alternative algorithm to match LADs to TV regions that allows for LADs to cross TV region boundaries.

[23] Figure 8 plots *total* numbers on each axis. An alternative would be to scale each series by the number of cannabis users. This would obviously entail the use of richer data that is subject to the same set of issues demand-side market estimates are subject to.

of the market, and hence the revenues accruing to criminal organizations supplying cannabis; (ii) provides an indication of the tax revenues that might be generated from the legalization of cannabis.

Our approach takes ideas from forensic economics, exploiting data on *licit* markets to measure the amount of *illicit* activity being undertaken [Zitzewitz 2012]. Using the insight that the consumption of cannabis is often combined with the use of highly complementary legal inputs in the form of roll-your-own tobacco and rolling papers, we show proof-of-concept for a forensic approach that specifies how legal and illegal inputs are combined in legal markets (the production of hand-rolled cigarettes) and illegal markets (the production of cannabis joints). These input relationships, along with market adding up conditions, can then be combined to infer the aggregate size of the market for cannabis. We show our approach can be readily calibrated by combining available point-of-sales data on the legal inputs of roll-your-own tobacco and rolling papers, with evidence on the relevant input parameters. We also show our approach yields a range of market size estimates based on alternative scenarios, and future research could be directed to help rule-in or rule-out some of these scenarios.

Our analysis sets out a clear agenda for future research: we see the best way forward as one that *combines* demand-side and forensic approaches. For example, established surveys from which demand-side estimates are derived should be extended to elicit information on key parameters required for the forensic method. There are three classes of information that would have especially high returns to being collected. First, the relative quantities of cannabis consumed in joint and non-joint form: this would show the extent to which the forensic approach underestimates the total market size once other forms of cannabis consumption are accounted for. Second, on the input parameter set required for the forensic approach, $(R_1, R_2, \beta, \theta)$, the elasticity estimates in Figure 4 highlight the relative returns to obtaining more precise estimates for each parameter. The highest returns relate to pinning down the tobacco content of hand-rolled cigarettes (R_1). The elasticity analysis shows the next highest returns would be to establishing more precise estimates of the number of rolling papers per joint (θ), and then on the total content of joints (R_2). Third, the analysis embodied in Figure 6 showed how the forensic measure of the cannabis market size varied under different scenarios. There are returns to collecting information in surveys to help pick out a scenario from those shown, especially related to the use of *types* of rolling paper used when constructing cannabis joints.

Once such data collection innovations are in place, it would be possible to construct a time-varying input parameter set, $\mathbf{\Pi}(t)$, that would allow for a complete forensic time-series market size for cannabis to be derived. With such a time-series (that could also be decomposed by region or month for example), it would be possible to move forward in two directions. First, to estimate own price elasticities for cannabis, and cross price elasticities with respect to the complementary legal inputs of RYO tobacco and rolling papers, as well as cross price elasticities with respect to complements/substitute goods such as alcohol, pre-rolled cigarettes and non-joint forms of cannabis consumption.[24] Ultimately, this allows researchers to estimate the underlying utility maximization problem that drives the forensic approach and is embodied in the input restrictions (1)-(3). This would neatly complement and extend existing work that structurally estimates individual demand for cannabis, such as Jacobi and Sovinksy [2016], accounting for non-random selection into consumption and using the demand system to back out price elasticities of cannabis demand. Second, having available such a time-series would enable researchers to conduct further validation exercises, say using the legalization of cannabis in some jurisdictions to then compare demand-side and forensic estimates as the market transitions from illicit to licit status, or to test the effectiveness of policies aiming to disrupt illicit drug markets.[25]

Ultimately, pushing forward this research agenda on methods and data collection can help bridge the gap economists currently face between research that identifies the causal impacts of policy interventions in illicit drug markets, cost-benefit analysis of those interventions, and informing important policy debates [Cook *et al.* 2013].

References

[1] ADDA.J, B.MCCONNELL AND I.RASUL (2014) "Crime and the Depenalization of Cannabis Possession: Evidence from a Policing Experiment," *Journal of Political Economy* 122: 1130-202.

[2] ATHA.M.J AND S.T.DAVIS (2011) Taxing the UK Cannabis Market, Independent Drug Monitoring Unit Report.

[3] BHATTACHARYA.J, C.GATHMANN AND G.MILLER (2013) "The Gorbachev Anti-Alcohol Campaign and Russia's Mortality Crisis," *American Economic Journal: Applied Economics* 5: 232-60.

[24] Descriptive evidence on such cross-price demand effects in tobacco-related products in the UK is provided in Rothwell *et al.* [2015].

[25] Although we have studied the market for cannabis, we hope our work encourages researchers to use similar forensic approaches to measure key quantities in other illicit markets. A potential example where such techniques could be applied relate to the production of illegal alcohol: this market has raised huge public health issues, especially in some transition economies, and its production requires the use of sugar, whose licit use could be estimated from dietary consumption data for example [Bhattacharya *et al.* 2013].

[4] BECKER.G.S (1968) "Crime and Punishment: An Economic Approach," *Journal of Political Economy* 76: 169-217.

[5] BRAMLEY-HARKER.E (2001) Sizing the UK Market for Illicit Drugs, RDS Occ. Paper 74.

[6] BRYAN.M.L, E.DEL BONO AND S.PUDNEY (2013) Licensing and Regulation of the Cannabis Market in England and Wales: Towards a Cost Benefit Analysis, mimeo, ISER Essex.

[7] CAULKINS.J.P, B.KILMER, P.H.REUTER AND G.MIDGETTE (2014) "Cocaine's Fall and Cannabis's Rise: Questions and Insights Based on New Estimates of Consumption and Expenditures in US Drug Markets," *Addiction* 110: 728-36.

[8] COOK.P.J, S.MACHIN, O.MARIE AND G.MASTROBUONI (2013) *Lessons from the Economics of Crime*, CESifo Seminar Series: MIT Press.

[9] CORMAN.H AND N.MOCAN (2000) "A Time-Series Analysis of Crime, Deterrence and Drug Abuse in New York City," *American Economic Review* 90: 584-604.

[10] CULLUM.P AND C.PISSARIDES (2004) The Demand for Tobacco Products in the UK, GES WP, London: UK.

[11] DOBKIN.C AND N.NICOSIA (2009) "The War on Drugs: Methamphetamine, Public Health, and Crime," *American Economic Review* 99: 324-49.

[12] DONOHUE.J.J, B.EWING AND D.PELOQUIN (2011) "Rethinking America's Illegal Drug Policy", in P.J.Cook, J.Ludwig and J.McCrary (eds.) *Controlling Crime: Strategies and Tradeoffs* NBER Conference Report, University of Chicago Press.

[13] FARRELLY.M, J.BRAY, G.ZARKIN, B.WENDLING AND R.PACULA (1999) The Effects of Prices and Policies on the Demand for Marijuana: Evidence from the National Household Surveys on Drug Abuse, NBER WP6940.

[14] FENDRICH.M, T.P.JOHNSON, J.S.WISLAR, A.HUBBELL AND V.SPIEHLER (2004) "The Utility of Drug Testing in Epidemiological Research: Results from a General Population Survey," *Addiction* 99: 197-208.

[15] FERGUSSON.D.M AND L.J.HORWOOD (1997) "Early Onset of Cannabis Use and Psycho-social Adjustment in Young Adults," *Addiction* 92: 279-96.

[16] FISMAN.R AND S-J.WEI (2004) "Tax Rates and Tax Evasion: Evidence from "Missing Imports" in China," *Journal of Political Economy* 112: 471-500.

[17] GALENIANOS.M, R.L.PACULA AND N.PERSICO (2012) "A Search-Theoretic Model of the Retail Market for Illicit Drugs", *Review of Economic Studies* 79: 1239-69.

[18] GROGGER.J AND M.WILLIS (2000) "The Emergence of Crack Cocaine and the Rise in Urban Crime Rates," *Review of Economics and Statistics* 82: 519-29.

[19] GROOM.C, T.DAVIES AND S.BALCHIN (1998) "Developing A Methodology For Measuring Illegal Activity For The UK National Accounts," *Economic Trends* 536: 33-71.

[20] HALL.W AND L.DEGENHARDT (2009) "Adverse Health Effects of Non-medical Cannabis Use," *Lancet* 374: 1383-91.

[21] HARRISON.L.D, S.S.MARTIN, T.ENEV AND D.HARRINGTON (2007) Comparing Drug Testing and Self-Report of Drug Use Among Youths and Young Adults in the General Population, DHHS Publication No. SMA 07-4249.

[22] HMRC (2011) *Measuring Tax Gaps 2011*, Supplementary Tables.

[23] HOME OFFICE (2015) *Drug Misuse: Findings from the 2014/15 Crime Survey for England and Wales*, D.Lader (ed.), Statistical Bulletin 03/15, Second edition, Supplementary Tables.

[24] JACOBI.L AND M.SOVINSKY (2016) "Marijuana on Main Street? Estimating Demand in Markets with Limited Access," *American Economic Review* 106: 2009-45.

[25] KELLY.E AND I.RASUL (2014) "Policing Cannabis and Drug Related Hospital Admissions: Evidence from Administrative Records," *Journal of Public Economics* 112: 89-114.

[26] KILMER.B, J.P.CAULKINS, R.L.PACULA AND P.H.REUTER (2011) "Bringing Perspective to Illicit Markets: Estimating the Size of the US Cannabis Market," *Drug and Alcohol Dependence* 119: 153-60.

[27] KILMER.B AND R.L.PACULA (2009) Estimating the Size of the Global Drug Market: A Demand-side Approach, RAND Europe.

[28] KUZIEMKO.I AND S.D.LEVITT (2004) "An Empirical Analysis of Imprisoning Drug Offenders," *Journal of Public Economics* 88: 2043-66.

[29] MARIE.O AND U.ZOLITZ (2017) "High Achievers? Cannabis Access and Academic Performance," *Review of Economic Studies*, forthcoming.

[30] MARSHALL.K.S, L.GOWING AND R.ALI (2011) "Pharmacotherapies for Cannabis Withdrawal," *Cochrane Database for Systematic Reviews* 1: CD008940.

[31] MIRON.J.A (2010) The Budgetary Implications of Drug Prohibition, mimeo, Harvard University.

[32] MOTT.J (1985) "Self-Reported Cannabis Use in Great Britain in 1981," *British Journal* 80: 37-43.

[33] ONDCP (2004) The Economic Costs of Drug Abuse in the United States 1992-2002, Executive Office of the President.

[34] ONS (2014) *United Kingdom: Local Authority Districts, Counties and Unitary Authorities, March 2009*, Office for National Statistics.

[35] PUDNEY.S (2007) Rarely Pure and Never Simple: Extracting the Truth from Self-reported Data on Substance Use, CeMMaP WP 11/07.

[36] PUDNEY.S, C.BADILLO, M.BRYAN, J.BURTON, G.CONTI AND M.IACOVOU (2006) "Estimating the Size of the UK Illicit Drug Market," in N.Singleton, R.Murray, and L.Tinsley (eds.) *Measuring Different Aspects of Problem Drug Use: Methodological Developments*, Home Office.

[37] ROTHWELL.L, J.BRITTON AND I.BOGDANOVICA (2015) "The Relation Between Cigarette Price and Hand-rolling Tobacco Consumption in the UK: An Ecological Study," *BMJ Open* 5:e007697.

[38] SCHAUER.G.L, B.A.KING, R.E.BUNNELL, G.PROMOFF AND T.A.MCAFEE (2016) "Toking, Vaping, and Eating for Health or Fun: Marijuana Use Patterns in Adults, US, 2014," *American Journal of Preventive Medicine* 50: 1-8.

[39] SUKHTANKAR.S (2012) "Sweetening the Deal? Political Connections and Sugar Mills in India," *American Economic Journal: Applied Economics* 4: 43-63.

[40] TRAUTMANN.F, B.KILMER AND R.TURNBULL (2013) *Further Insights into Aspects of the EU Illicit Drugs Market*, Luxembourg: Publications Office of the EU.

[41] UNODC (2006) *Bulletin on Narcotics: Review of the World Cannabis Situation* 51: No. 1, 2.

[42] UNODC (2011) *World Drug Report 2011*, United Nations Office on Drugs and Crime.

[43] VAN OURS.J (2003) "Is Cannabis a Stepping-stone for Cocaine?," *Journal of Health Economics* 22: 539-54.

[44] VAN OURS.J AND J.WILLIAMS (2009) "Why Parents Worry: Initiation Into Cannabis Use by Youth and Their Educational Attainment," *Journal of Health Economics* 28: 132-42.

[45] ZITZEWITZ.E (2012) "Forensic Economics," *Journal of Economic Literature* 50: 731-69.

A References for Appendix Tables

A.1 References for Table A1

CAMPAIGN FOR TOBACCO-FREE KIDS (2009) Fixing Federal Tax Rates on Roll-Your-Own (RYO) Tobacco, Washington, DC.

CDC (2012) Consumption of Cigarettes and Combustible Tobacco: United States 2000-2011, *Morbidity and Mortality Weekly Report (MMWR)* 61: 565-69, Centers for Disease Control and Prevention.

COCKEYED.COM (2011) How Much is Inside a Pack of Rolling Tobacco?, http://www.cockeyed.com/inside/tobacco/tobacco.html, accessed January 25, 2011.

DARRALL.K.G AND J.A.FIGGINS (1998) "Roll-your-own Smoke Yields: Theoretical and Practical Aspects," *Tobacco Control* 7: 168-75.

DYMOND.H (1996) "Making Habits of Roll-your-own Smokers in the Netherlands and Tar and Nicotine Yields from the Resultant Products," *Tobacco Science* 40: 87-91.

END SMOKING NZ (2011) Tax on the RYO Cigarette as Smoked is Still 30% Less; Many Poor Smokers Switch to RYOs, Fatally Delaying Quitting, http://www.endsmoking.org.nz/RYOhalfprice.htm, accessed 1 December 2011.

FOREY.B, J.HAMLING, A.THORNTON AND P.LEE (2012) *International Smoking Statistics: A Collection of Worldwide Historical Data*, United Kingdom, Web Edition.

PWC (2005) Research Report on the Illegal Tobacco Market, PriceWaterhouseCoopers.

PWC (2007) Illegal Tobacco Trade: Costing Australia Millions: Strategies to Curb the Supply and Use of Illegal Tobacco Products, PriceWaterhouseCoopers.

PWC (2010) Australia's Illegal Tobacco Market: Counting the Cost of Australia's Black Market, PriceWaterhouseCoopers.

SHAHAB.L, R.WEST AND A.MCNEILL (2008) "The Feasibility of Measuring Puffing Behaviour in Roll-your-own Cigarette Smokers," *Tobacco Control* 17: i17-23, Supplement I.

A.2 References for Table A2

ABT ASSOCIATES (2001) *What America's Users Spend on Illegal Drugs 1988-2000*, ONDCP.

ATHA.M AND S.BLANCHARD (1985) "Appendix: Cannabis Use in Britain," in A.Henman, R.Lewis and T.Malyon (Editors), *Big Deal: The Politics of the Illicit Drugs Business*, Pluto Press: London and Sydney.

BUCHANAN.B AND D.O'CONNELL (1998) "Survey on Cannabis Resin and Cannabis in Unsmoked Handrolled Cigarettes Seized in the Republic of Ireland," *Science & Justice* 38: 221-4.

FAIRBAIRN.J.W, I.HINDMARCH, S.SIMIC AND E.TYLDEN (1974) "Cannabinoid Content of some English Reefers," *Nature* 249: 276-8.

GETTMAN.J (2007) Lost Taxes and Other Costs of Cannabis Laws, The Bulletin of Cannabis Reform.

GHODSE.H (2010) *Ghodse's Drugs and Addictive Behaviour: A Guide to Treatment*, Cambridge University Press: 4th Edition.

HUMPHREYS.I.J AND J.R.JOYCE (1982) "A Survey of the Cannabis Content of Unsmoked Reefer Cigarettes," *Journal of the Forensic Science Society* 22: 291-2.

IVERSEN.L.L (2008) *The Science of Cannabis*, Oxford University Press.

KILMER.B, J.P.CAULKINS, B.M.BOND AND P.H.REUTER (2010) Reducing Drug Trafficking Revenues and Violence in Mexico. Would Legalizing Cannabis in California Help?, RAND Corporation.

LEGLEYE.S, C.B.LAKHDAR AND S.SPILKA (2008) "Two Ways of Estimating the Euro Value of the Illicit Market for Cannabis in France," *Drug and Alcohol Review* 27: 466-72.

MACCOUN.R.J AND P.H.REUTER (2001) *Drug War Heresies. Learning from Other Vices, Times, and Places*, RAND Studies in Policy Analysis, Cambridge University Press.

MCBRIDE.A.J (1995) "Cannabis Use in a Drug and Alcohol Clinic Population," *Drug and Alcohol Dependence* 39: 29-32.

RIGTER.H AND M.VAN LAAR (2002) "Epidemiological Aspects of Cannabis Use," *Cannabis 2002 Report*: 11-36.

ROBSON.P (2009) *Forbidden Drugs*, Oxford University Press.

SLACK.A, D.O'DEA, I.SHEERIN, D.NORMAN, J.WU AND G.NANA (2008) New Zealand Drug Harm Index, Technical report, BERL House, Report to the New Zealand Police.

UNODC (2006) Bulletin on Narcotics: Review of the World Cannabis Situation, Technical Report LVIII, Nos. 1 and 2, UNODC.

WILKINS.C, K.BHATTA AND S.CASSWELL (2002) "A 'Demand Side' Estimate of the Dollar Value of the Cannabis Black Market in New Zealand," *Drug and Alcohol Review* 21: 145-51.

WILKINS.C, J.L.REILLY, M.PLEDGER AND S.CASSWELL (2005) "Estimating the Dollar Value of the Illicit Market for Cannabis in New Zealand," *Drug and Alcohol Review* 24: 227-34.

A.3 References for Tables A3 and A4

CHAIT.L AND J.PIERRI (1989) "Some Physical Characteristics Of Nida Cannabis Cigarettes," *Addictive Behaviors* 14: 61-7.

FAIRBAIRN.J.W, I.HINDMARCH, S.SIMIC AND E.TYLDEN (1974) "Cannabinoid Content of some English Reefers," *Nature* 249: 276-8.

MCBRIDE.A.J (1995) "Cannabis Use in a Drug and Alcohol Clinic Population," *Drug and Alcohol Dependence* 39: 29-32.

TYLE.A (1995) *Street Drugs*, Hodder and Stoughton, 2nd edition.

UNODC (2006) Bulletin on Narcotics: Review of the World Cannabis Situation, Technical Report LVIII, Nos. 1 and 2, UNODC.

Figure 1A: Ranking of Cannabis Use by Prevalence, 2009

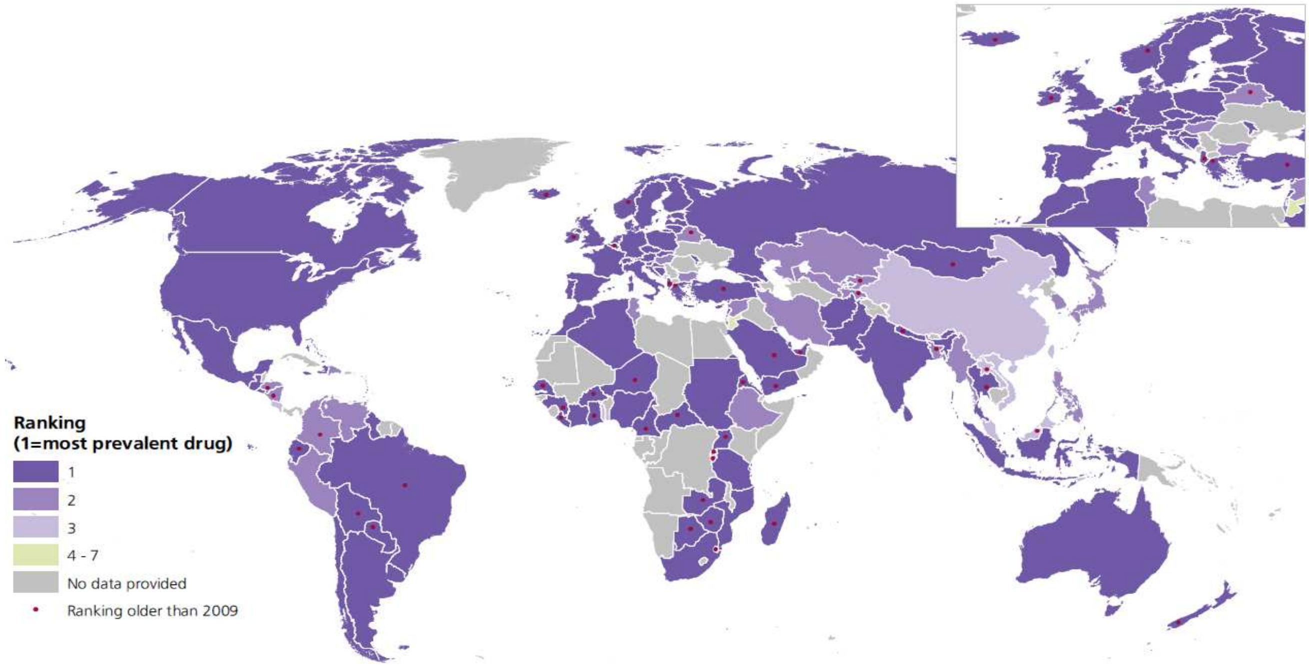

Figure 1B: Prevalence of Cannabis Use, 2009

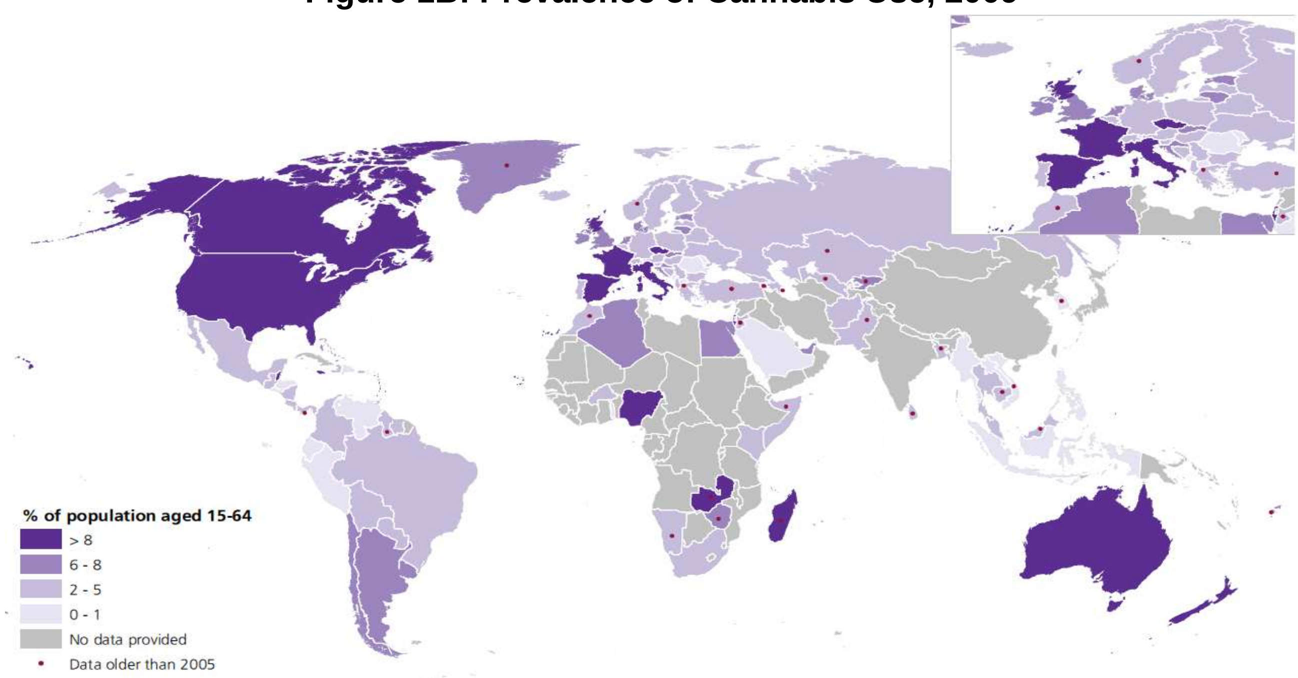

Source and Notes: Chapter 5 (Map 40 (p177) and Map 42 (p188), World Drug Report, 2011 [UNODC]. Figure 1A is based on figures from 2009 (or the latest year available). In both Figures, the boundaries and names shown and the designations used on the maps do not imply official endorsement or acceptance by the United Nations. The dotted line represents approximately the line of control in Jammu and Kashmir agreed upon by India and Pakistan. The final status of Jammu and Kashmir has not been agreed upon by the parties.

Figure 2A: Estimates for the Market Size for Cannabis

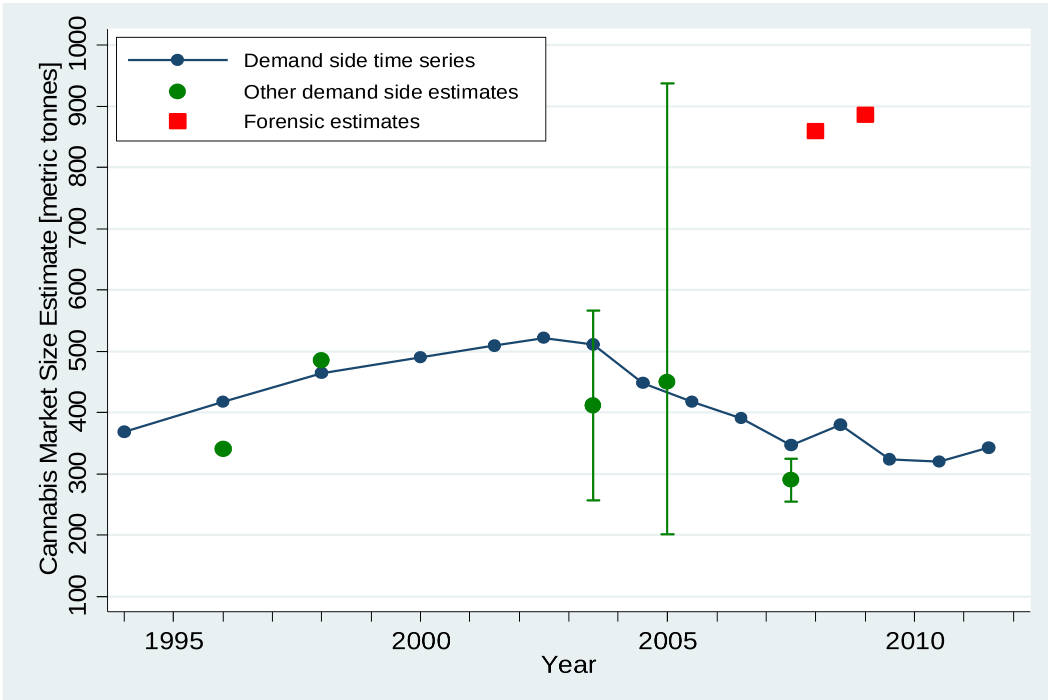

Figure 2B: Implied Seizure Rates

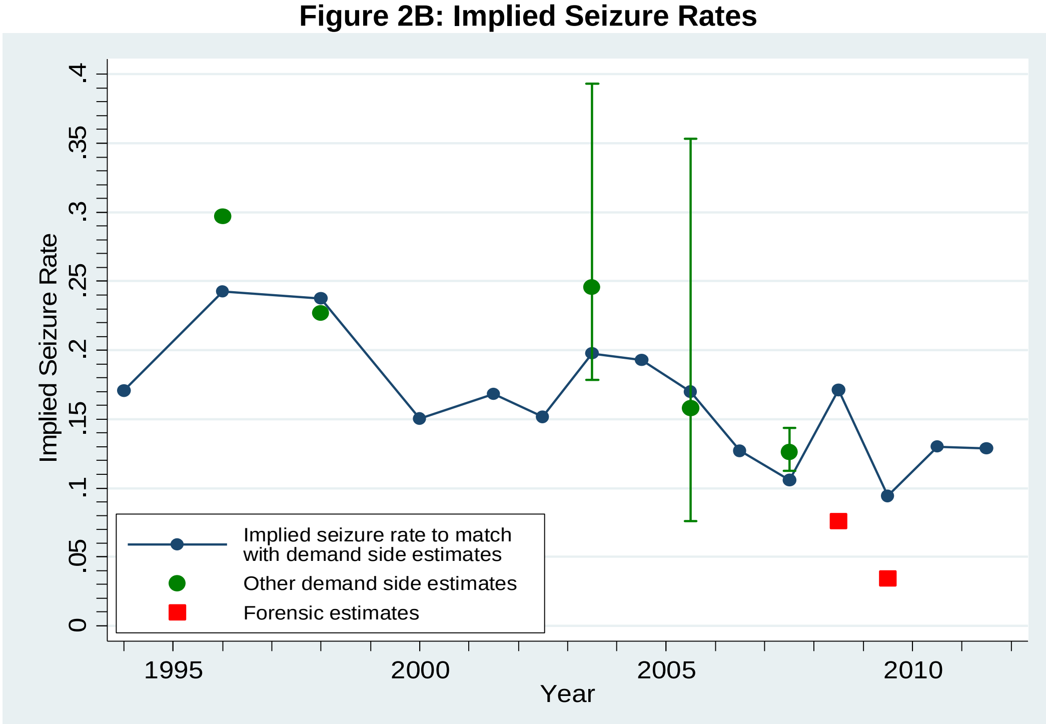

Notes: In Figure 2A, the 'Demand-side time series' applies the cannabis market-size estimation approach of Kilmer and Pacula [2009] to UK data for all years from 1994 to 2012. The method uses survey-based prevalence rates, estimated separately for 'regular users' and 'occasional users' (these types are distinguished by survey responses as to whether the individual has consumed cannabis in last month or the last year, respectively). Information on prevalence rates is obtained from EMCDDA [2015a]. The other data requirements, for each type of user, needed to implement this method are: (i) the number of days per year cannabis is typically consumed; (ii) the amount of cannabis consumed per use-day. Kilmer and Pacula [2009]'s review of the evidence suggest baseline values of 2.5 (1.25) joints per day for regular (occasional) users, and 0.4 gram of cannabis per joint (for both user types). These inputs are then used to construct annual consumption per user, and combined with population statistics for individuals 15-64 to arrive at an estimate of the aggregate market size for cannabis. Annual population estimates are taken from ONS data. In a final step, underreporting is taken into account by re-scaling the resulting market size estimates by 20%. The 'previous estimates' shown on Figure 2A are those reported in EMCDDA [2012, Table 5.1], supplemented with two earlier UK-based studies [Groom et al. 1998, Bramley-Harker 2001]. Where possible we show the range of estimates provided in these earlier studies. Figure 2A also shows the baseline estimates from the forensic approach for 2008 and 2009. Figure 2B graphs the ratio of annual seizures of cannabis in the UK relative to the demand-side time series estimate in Figure 2A. This shows the implied seizure rate that would reconcile the demand-side estimate with the level of annual seizures. The source for seizure data is EMCDDA [2015b,c,d].

Figure 3: The Forensic Approach

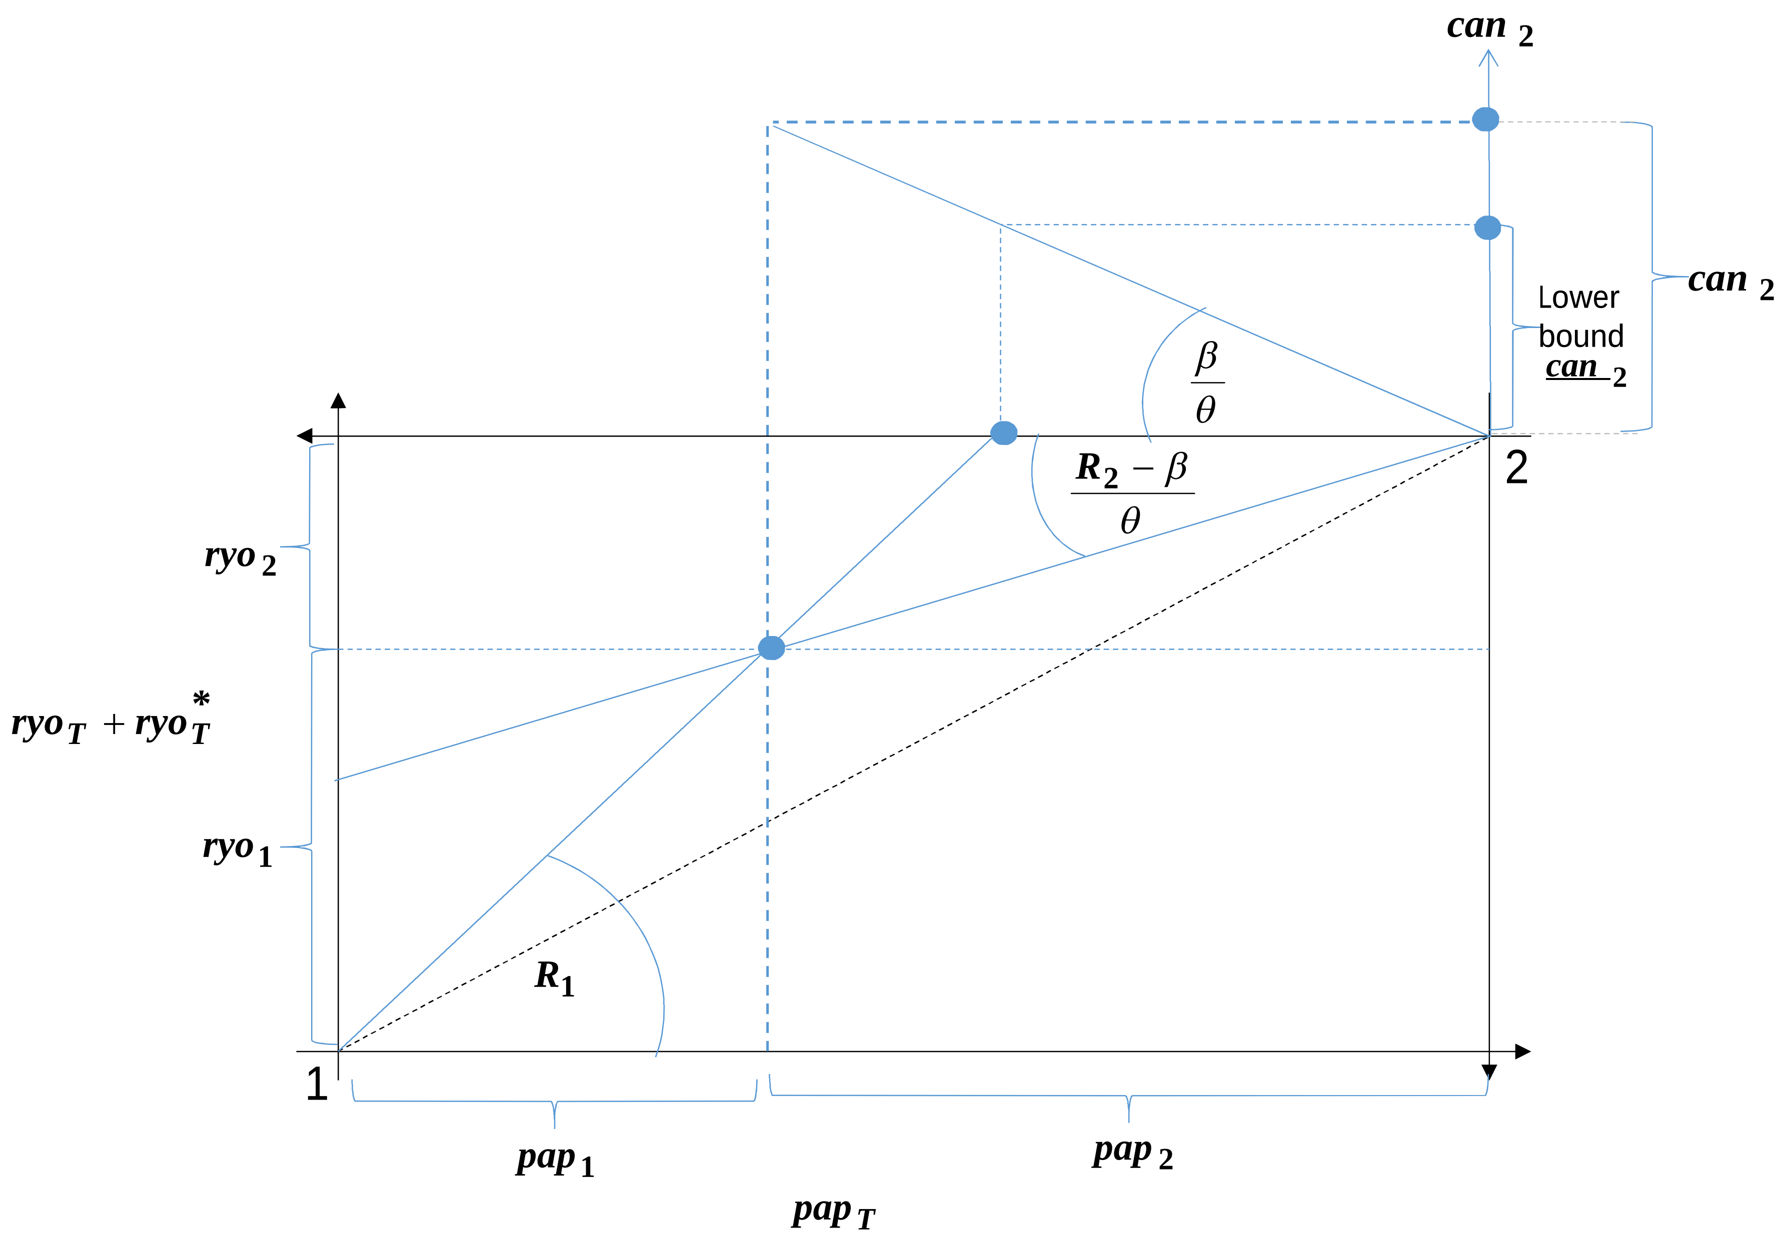

Figure 4: Elasticities of the Forensic Market Size Estimate

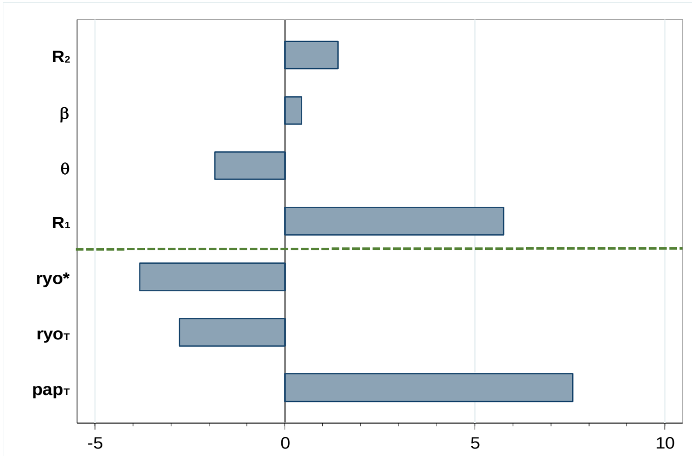

Figure 5: The Forensic Estimate and Input Parameters

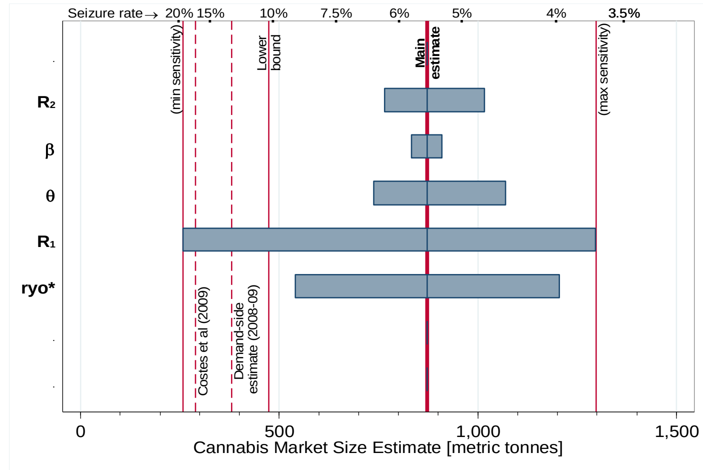

Figure 6: Reconciling Forensic Measures with Demand-Side Estimates

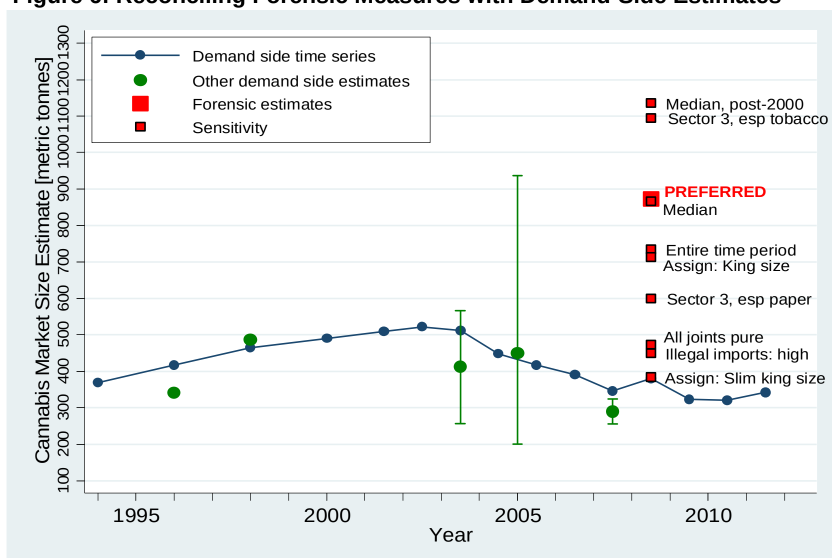

Notes: Figure 4 shows the implied elasticity of our baseline forensic estimate of the aggregate market size for cannabis in the UK (averaged over 2008 and 2009). Figure 5 is a Tornado figure that shows how the forensic market size estimate varies as we change each parameter value by plus or minus 10 per cent from its baseline value, holding all other parameters constant. Figure 5 also superimposes the contemporaneous estimates of the market size from the existing literature of demand-side studies: the studies shown are those reported in EMCDDA [2012, Table 5.1], supplemented with two earlier UK-based studies [Groom et al. 1998, Bramley-Harker 2001]. The upper axis in Figure 5 indicates the implied seizure rates for each market size, based on UK seizure data in EMCDDA [2015b,c,d]. In Figure 6,the 'Demand-side time series' applies the cannabis market-size estimation approach of Kilmer and Pacula [2009] to UK data for all years from 1994 to 2012. The method uses survey-based prevalence rates, estimated separately for 'regular users' and 'occasional users' (these types are distinguished by survey responses as to whether the individual has consumed cannabis in last month or the last year, respectively). Information on prevalence rates is obtained from EMCDDA [2015a]. The other data requirements, for each type of user, needed to implement this method are: (i) the number of days per year cannabis is typically consumed; (ii) the amount of cannabis consumed per use-day. Kilmer and Pacula [2009]'s review of the evidence suggest baseline values of 2.5 (1.25) joints per day for regular (occasional) users, and 0.4 gram of cannabis per joint (for both user types). These inputs are then used to construct annual consumption per user, and combined with population statistics for individuals 15-64 to arrive at an estimate of the aggregate market size for cannabis. Annual population estimates are taken from ONS data. In a final step, underreporting is taken into account by re-scaling the resulting market size estimates by 20%. The 'previous estimates' shown on Figure 2A are those reported in EMCDDA [2012, Table 5.1], supplemented with two earlier UK-based studies [Groom et al. 1998, Bramley-Harker 2001]. Where possible we show the range of estimates provided in these earlier studies. Figure 2A then also shows the baseline estimates from the forensic approach for 2008 and 2009, as well as various alternative forensic estimates that make different underlying assumptions or use alternative inputs. These include: (i) taking HMRC's upper bound estimates for the quantity of illegally imported roll-your-own tobacco; (ii) taking median estimates of technology parameters, or taking median estimates from studies from 2000 onwards; (iii) allowing for a generic third sector to exist in which some percentage of the total inputs of rolling papers and roll-your-own tobacco are utilized (or simply wasted); (iv) exclusively assigning some rolling papers (King Size, Slim King Size) to be used only in Sector 2 and that no other types of rolling paper are used in Sector 2.

Figure 7: Monthly Variation in Forensic Market Size Estimates

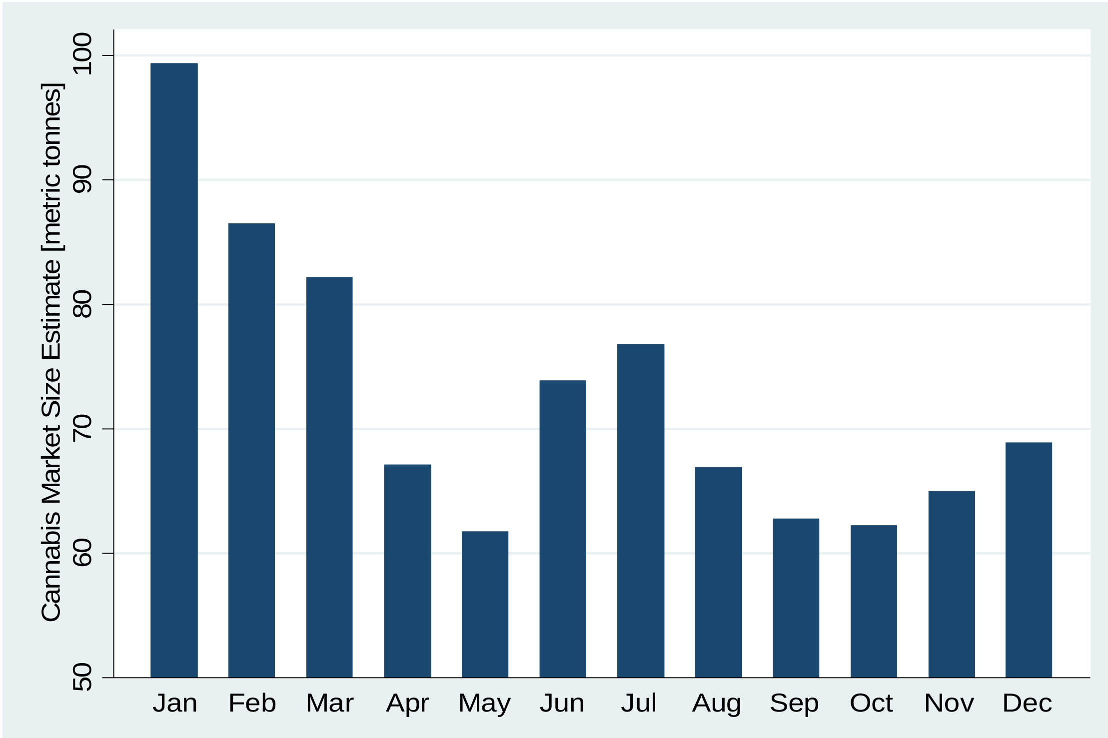

Figure 8: Validation with Administrative Records on Hospital Admissions

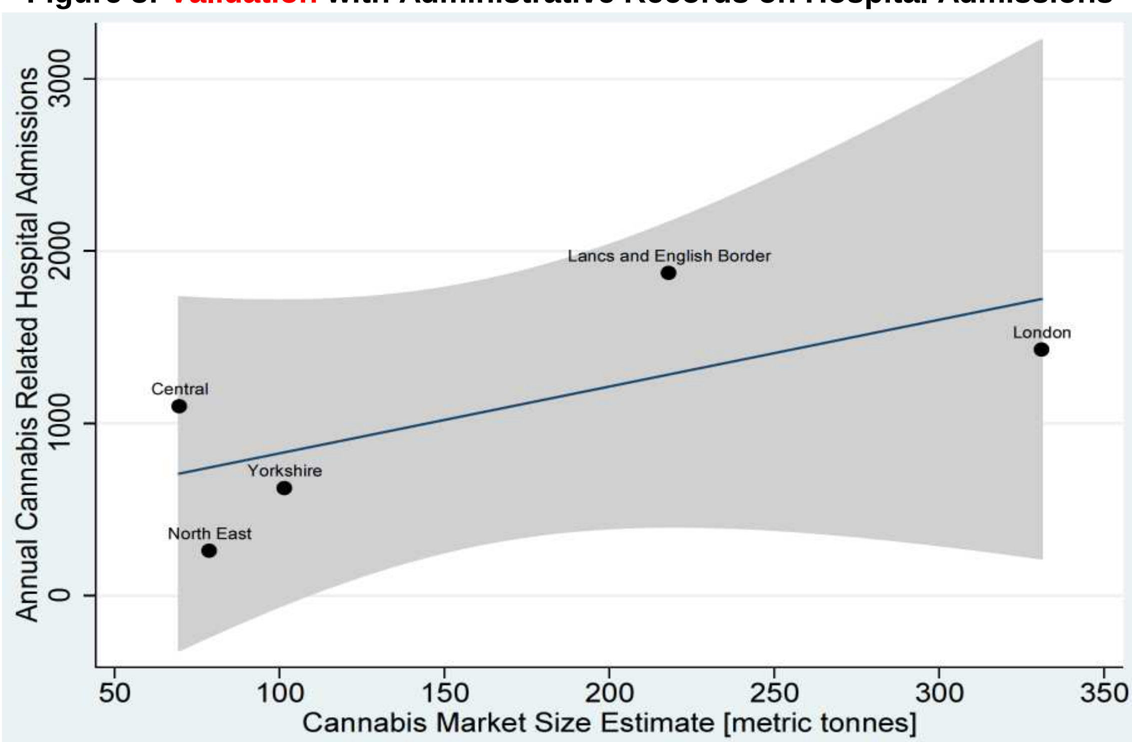

Notes: Figure 7 shows the cannabis market size estimate by month derived from the forensic approach, where monthly estimates are averaged over 2008 and 2009. Figure 8 shows a scatter plot of estimates of the market size for cannabis broken down by UK TV Region, against hospital admissions for cannabis related causes, based on National Health Service administrative records on hospital admissions, drawn from the Inpatient Hospital Episode Statistics. These provide an administrative record of every inpatient health episode, defined as a single period of care under one consultant in an English National Health Service hospital. For each patient-episode event, the data record the date of admission and ICD-10 diagnoses codes in order of importance. Based on ICD-10 classifications, we calculate the total number of hospital admissions where cannabis usage is listed as either a primary or secondary cause of admission. In the scatter plot, both series are averaged over 2008 and 2009.

Table 1: Current Estimates of the UK Market Size for Cannabis[a]

	Time Period	Preferred Market Size Estimate [tonnes]	Notes
(a) Demand-Side Estimates			
Groom et al. [1998][b]	1996	341	Table 6
Bramley-Harker [2001]	1998	486	Table 4.2
Pudney [2006]	2003/4	412 (257 - 567)	Table S4.1
Kilmer and Pacula [2009]	2005	450 (201 - 937)	Table 7
Costes et al [2009]	2007/8	290 (255 - 325)	as reported in EMCDDA 2012, Table 5.1
(b) Supply-side Estimates			
Groom et al. [1998][c]			
Seizure rate 5%	1996	3759	
Seizure rate 10%	1996	1836	
Seizure rate 15%	1996	1194	
Seizure rate 20%	1996	874	

Notes:

[a] The studies shown are those reported in EMCDDA [2012, Table 5.1], supplemented with two earlier UK-based studies [Groom et al. 1998, Bramley-Harker 2001].

[b] These are initial estimates of 'consumption' before the balancing procedure of Groom et al. [1998] is implemented.

[c] This is our own calculation based on Table 6 in Groom et al. [1998]. The initial estimates ('import' and 'domestic production') are taken before balancing procedure is implemented.

Table 2: Baseline Forensic Estimates of the Market Size for Cannabis

	R_1	**β**	**R_2**	**θ**	**pap_T**	**ryo_T**	**ryo_T***	**can_2**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
Variable	Sector 1 intensity	cannabis / joint	Sector 2 intensity	paper/joint	Papers total	RYO legal total	RYO illegal imports	**Cannabis Market Size Estimate**
Units	[g/paper]	[g/joint]	[g/joint]	[paper/joint]	[million paper]	[metric tonnes]	[metric tonnes]	[metric tonnes]
(1) All studies	0.648	0.363	1	1.98	17748	4293	5900	734.8
(2) Studies based on data from 2000 or later	0.661	0.402	1	1.98	17748	4293	5900	873.2

Notes: Each row represents an alternative estimate of the aggregate market size for cannabis in the UK using the forensic approach. Columns 1 to 4 show the value of the technology parameters assumed in the calculation, and Columns 5 to 7 show the aggregate input values related to the legal sector used in each calculation. Column 8 shows the resultant market size estimate. Row (1) shows our baseline estimate using technology parameter estimates using all study sources shown in Appendix Tables A1 to A5. Row (2) shows our market size estimate based only on those studies that provide technology estimates based on data from and including 2000.

Table 3: Reconciling Demand-Side Market Size Estimates with the Baseline Forensic Estimate

		Baseline Demand-side estimate	Adjustment to all parameters	Adjustment to all parameters except beta	Element-by-element adjustments				
					Under-reporting	Prevalence	Consumption days/year	Joints per day	g/joint
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Required Adjustment:		-	21% increase	30% increase					
Rescaling Factor to Correct for Nonresponse and Under-reporting		20%	24%	26%	65%	20%	20%	20%	20%
Regular users	Prevalence	0.046	0.056	0.059	0.046	0.106	0.046	0.046	0.046
	Days	150.3	182.5	193.4	150.3	150.3	344.9	150.3	150.3
	Joints	2.50	3.04	3.22	2.50	2.50	2.50	5.74	2.50
	g/joint	0.40	0.49	0.40	0.40	0.40	0.40	0.40	0.92
Occasional users	Prevalence	0.079	0.096	0.102	0.079	0.181	0.079	0.079	0.079
	Days	29.9	36.2	38.4	29.9	29.9	68.5	29.9	29.9
	Joints	1.25	1.52	1.61	1.25	1.25	1.25	2.87	1.25
	g/joint	0.40	0.49	0.40	0.40	0.40	0.40	0.40	0.92
Implied market size [metric tonnes]		380.5	873.2	873.2	873.2	873.2	873.2	873.2	873.2

Notes: The Table shows how the baseline estimate from our forensic approach can be reconciled with adjustments to the standard demand-side approach of Pacula and Kilmer [2009], so that they both generate an implied market size for cannabis of 873.2 metric tonnes. Column 1 shows the standard parameter values used for the Pacula and Kilmer [2009] demand-side approach. The method uses survey-based prevalence rates, estimated separately for 'regular users' and 'occasional users' (these types are distinguished by survey responses as to whether the individual has consumed cannabis in last month or the last year, respectively). Information on prevalence rates is obtained from EMCDDA [2015a]. The other data requirements, for each type of user, needed to implement this method are: (i) the number of days per year cannabis is typically consumed; (ii) the amount of cannabis consumed per use-day. These inputs are then used to construct annual consumption per user, and combined with population statistics for individuals 15-64 to arrive at an estimate of the aggregate market size for cannabis. Annual population estimates are taken from ONS data. In a final step, underreporting is taken into account by re-scaling the resulting market size estimates by 20%. Column 2 shows the adjustment required if all parameters were to be adjusted in the same proportion, so that the demand-side estimate would reconcile with the forensic estimate. Column 3 repeats this but holds constant one technology parameter: beta. Columns 4 to 8 then vary one parameter at a time, showing the extent to which that single parameter needs to be adjusted to match the forensic baseline estimate.

Table 4: Reconciling Forensic Market Size Estimates with the Demand-Side Estimate

	R_1 (1)	β (2)	R_2 (3)	θ (4)	pap_T (5)	ryo_T (6)	ryo_T* (7)	can_2 (8)	(9)	(10)
Variable	Sector 1 intensity	cannabis / joint	Sector 2 intensity	paper/joint	Papers total	RYO legal total	RYO illegal imports	**Cannabis Market Size Estimate**	**Implied Scale-up to Demand-side Estimate**	**Implied Seizure Rate**
Units	[g/paper]	[g/joint]	[g/joint]	[paper/joint]	[million paper]	[metric tonnes]	[metric tonnes]	[metric tonnes]		
Baseline Estimates										
(1) All studies	0.648	0.363	1	1.98	17748	4293	5900	734.8	17%	7%
(2) Studies based on data from 2000 or later	0.661	0.402	1	1.98	17748	4293	5900	873.2	21%	5%
Estimate (R_1, β) Using Median Rather than Mean Estimate										
(3) All studies	0.665	0.370	1	1.98	17748	4293	5900	867.2	21%	6%
(4) Studies based on data from 2000 or later	0.700	0.400	1	1.98	17748	4293	5900	1135.2	29%	4%
Bounds on ryo_T*										
(5) Upper bound	0.661	0.402	1	1.98	17748	4293	6650	449.8	4%	11%
Generic Third Sector										
(6) Paper: 5%, Tobacco: 1%	0.661	0.402	1	1.98	16861	4250	5841	599.4	11%	8%
(7) Paper: 1%, Tobacco: 5%	0.661	0.402	1	1.98	17571	4078	5605	1094.7	28%	4%
(8) All joints pure (no RYO content)	0.661	0.402		1.98	17748	4293	5900	474.2	5%	10%
Assign Paper Types to Illegal Sector 2										
(9) King Size papers assigned exclusively to Sector 2		0.402		1.98				713.1	16%	7%
(10) Slim King Size papers assigned exclusively to Sector 2		0.402		1.98				384.5	0%	12%

Notes: Each row represents an alternative estimate of the aggregate market size for cannabis in the UK using the forensic approach. Columns 1 to 4 show the value of the technology parameters assumed in the calculation, and Columns 5 to 7 show the aggregate input values related to the legal sector used in each calculation. Column 8 shows the resultant market size estimate. Column 9 shows how the demand-side estimates following the methodology of Pacula and Kilmer [2009] would need to be scaled-up in each of their inputs to reconcile with the market size estimate shown in Column 8. Column 10 shows the implied seizure rate that would reconcile the forensic market size estimate in each row with the level of annual seizures. The source for seizure data is EMCDDA [2015b,c,d]. Row (1) shows our baseline estimate using technology parameter estimates using all study sources shown in Appendix Tables A1 to A5. Row (2) shows our market size estimate based only on those studies that provide technology estimates based on data from and including 2000. Rows (3) and (4) show estimates based on median (rather than mean) values of technology parameters derived from these studies. Row (5) uses the upper bound estimate, provided by HMRC with regards to illegal imports of roll-your-own tobacco. Rows (6) and (7) allow for a generic third sector to exist in which some percentage of the total inputs of rolling papers and roll-your-own tobacco are utilized (or simply wasted). Row (6) considers the scenario in which 5% of rolling papers and 1% of all roll-your-own tobacco is used in Sector 3. Row (7) considers the scenario in which 1% of rolling papers and 5% of all roll-your-own tobacco is used in Sector 3. Row (8) considers the scenario in which all joints only contain cannabis (and no tobacco). Rows (9) and (10) make further assumptions on the exclusive use of some kinds of rolling papers in the illegal Sector 2. Row (9) assumes that King Size rolling paper products are used only in Sector 2 and that no other types of rolling paper are used in Sector 2. Row (10) assumes that Slim King Size rolling paper products are used only in Sector 2 and that no other types of rolling paper are used in Sector 2.

Table A1: Estimates of R_1 from Legal RYO Sector

Parameter	Estimate	Source	Country	Population	Source Field	Notes
Tobacco content of RYO cigarettes	0.44g	Rosenberry et al (2013)	US	Lab study	Epidemiology / Health policy	0.43-0.45g, depending on home versus lab.
	0.455g	Laugesen et al. (2009)	New Zealand	Lab study, male volunteer smokers	Epidemiology / Health policy	
	0.48g	Gallus et al. (2014)	UK	survey of current smokers	Medical	Median weight for sample from England. Median weight across all included European countries 0.75g.
	0.49g	Dymond (1996, Tobacco Science)			Chemistry	Cited in Darrall & Figgins (1998)
	0.49g, 0.54g	End Smoking NZ (2011)	New Zealand	Survey evidence	Health policy	Citing Tobacco Use Survey 2006 (Ministry of Health) and AC Nielsen smoking survey reports to the Ministry of Health, 1983-2005.
	0.505g	Darrall & Figgins (1998, Tobacco Control)	UK	Lab study with volunteer smokers	Chemistry	Range 0.226-0.919. 26 participants with 20 RYO cigarettes each
	0.511g	Shahab et al. (2008)	UK	Lab study with volunteer smokers	Epidemiology	(95% CI 0.476g to 0.548g). Length 70mm (pre-determined from experiment setting).
	0.5g	Forey et al. (2012)	US		Health policy	"averages calculated from Wald and Nicolaides-Bouman (1991)"
	0.63g	Cornelsen et al (2013)	Ireland	Expert statement	Health policy	Range in the literature indicated as 0.4-0.8g.
	0.70g	Laugesen (2012)	New Zealand	Expert statement	Medical	'standard RYO cigarette'
	0.70g, 0.77g	PWC (2005, 2007, 2010)	Australia	Factual Statement	Health policy	"average based on rolling habits"
	0.717g	www.cockeyed.com (2011)	US	Self-observation	Self-observation	Self-observation: preparing RYO from a given pack. Original weight 0.78g. 76.2mm length, no filter. Adjusted weight to 70mm
	0.75g	Darrall & Figgins (1998, Tobacco Control)		Factual statement	Chemistry	Required amount of tobacco for automatic cigarette maker, designed for use with hand-rolling tobacco
	0.765g	Campaign for Tobacco-Free Kids (2009)	US	Market research	Health policy	Evidence from RYO brands supplying both tobacco and papers (own calculation)
	0.80g	Fu et al (2014)	Spain	Expert statement	Medical	Middle case. Scenarios considered are 0.5g, 0.8g, 1.0g
	0.874g	Campaign for Tobacco-Free Kids (2009)	US	Market research	Health policy	Evidence from RYO starter kits (own calculation)
	0.88g	Gallus et al. (2013)	Italy	RYO users	Medical	Mean among regular and occasional users (Median=0.63g).
	0.92g	CDC (2012)	US	Factual statement	Medical	Based on conversion formula in Master Settlement Agreement

Table A2: Estimates of R_2 in Illicit Sector

Parameter	Estimate	Source	Country	Population	Source Field	Notes
Weight per joint in grams	0.891g	Chait et al (1989), own calculation	US	Joints for research and medical use	Psychiatry	Study reports total weight of both cannabis and tobacco.
	0.927g	Fairbairn et al. (1974), own calculation	UK (London, Leeds)	Lab study using volunteer samples	Pharmacology	Study reports total weight of both cannabis and tobacco. Reported figure is sample average, including both herbal and resin
	1.0g	UNODC (2006)	Netherlands	Dutch coffee shop cigarettes	Medicine	Study reports joints as containing 0.9g tobacco, 0.1g cannabis.
	1.00g	McBride (1995), own calculation	UK (Wales)	Community drug clinic attendees	Psychiatry	Study reports tobacco content of joint as 0.79 of King size cigarette. Estimate based on assumption that King size cigarette contains 0.8g tobacco, and parameter $\beta=0.363$.
	1.09g	Tyle (1995), own calculation	UK	Factual statement	Drugs policy	Study reports ratio of herbal:tobacco as 1:2. Estimate based on assumption that $\beta=0.363$.

Table A3: Estimates of β (Content of Joints)

Parameter	Estimate	Source	Country	Population	Source Field	Notes
Content of one joint: RESIN	0.137g	Humphreys & Joyce (1982)	UK	Lab study of seizures	Forensics	N=188
	0.1g	Buchanan & O'Connell (1998)	Republic of Ireland	Lab study of seizures, 1980-1996	Justice	N=2025
	0.212g	own calculation based on Fairbairn et al. (1974)	UK (London, Leeds)	Lab study using volunteer samples	Pharmacology	N=7
	0.35g	McBride (1995, p.30)	UK (Wales)	Community drug clinic attendees	Psychiatry	SD=0.14
Content of one joint: HERBAL	0.197g	Humphreys & Joyce (1982)	UK	Lab study of seizure	Forensics	N=54. 16.7% contain more than 0.35g. 11.1% in range 0.2-0.35g. 44.4% less than 0.1g.
	0.26g	Buchanan & O'Connell (1998)	Republic of Ireland	Lab study of seizures, 1980-1996	Justice	N=179
	0.294g	Own calculation based on Fairbairn et al. (1974)	UK (London, Leeds)	Lab study using volunteer samples	Pharmacology	N=7
	0.39g	Abt Associates (2001)	US	System to Retrieve Drug Evidence (STRIDE), data from 1993	Drug policy	
	0.3g	Robson (2009, p.70)	UK (?)	Factual statement	Drugs policy	"An average marijuana roll-up or `spliff' contains around 300mg of herbal material, though `fatties' or `blunts' may pack up to 1g." (p.70)
	0.3g-0.5g	Ghodse (2010, p.96)	UK	Factual statement	Pharmacology	"An average marijuana cigarette in the UK usually contains 300-500mg of herbal material with a THC content of perhaps 1-2%" (p.96)
	0.46g	Kilmer et al. (2010)	US	Survey in Arrestee Drug Abuse Monitoring Program	Drug policy	Estimate is recovered from responses on $ paid per joint, and $ paid per gram. 95% CI: 0.43-0.50g.
	0.4g	MacCoun & Reuter (2001)	US	Factual statement	Drug policy	"If the average joint has 0.4 grams, ..." (p.343)
	0.5g	Iversen (2008, p.14)	UK (?)	Factual statement	Pharmacology	"with or without added tobacco - which assists the otherwise often erratic burning of the marijuana" (p.14)
	0.5g	NNICC Report 1997, as cited in Abt Associates (2001)	US		Drug policy	
	0.5g	California Commission on Peace Office Standards and Training (1992), as cited in Gettman (2007)	US	Factual statement for "typical homemade"	Justice	Slender "matchstick"=0.34g (3/16th inches diameter), a typical homemade=0.50g (5/16th inches diameter), tobacco cigarette-refilled with marijuana 0.9g
	0.62g	McBride (1995, p.30)	UK (Wales)	Survey of 100 community drug clinic attendees	Psychiatry	SD=0.3
Content of one joint: OTHER	0.1g; 0.25g	UNODC (2006, p.43)	NL	Coffee shop cigarette vs street sales cigarette. Factual statement/expert witness	Drug policy	
	0.2-0.5g	Rigter & van Laar (2002)	NL, D, US	Factual statement	Epidemiology	"When smoked with tobacco, for instance, one gram may be processed into two to five joints." (p.21). Refers to survey conducted in three cities (Amsterdam, Bremen, San Francisco)
	0.33g	Atha & Blanchard (1985), as cited in McBride (1995, p.31)	UK	LCC survey	Drug policy	No distinction Resin vs Herbal
	0.33g	Legleye et al. (2008)	France	Factual statement	Drug policy	"six to nine joints per barrette (smallest quantity sold in France) if this weighs between 2 and 3 g." (p. 468)
	0.33g; 0.5g	Slack et al (2008)	New Zealand	Factual statement	Drug policy	Appendix Table 6 (p.64); p.3
	0.5g	Iversen (2008, p.12)	UK (unclear)	Factual statement		"... as the average joint only contains a total of around 500mg of herbal cannabis or resin..." (p.12)
	0.5g	ESR (1998), as cited in Wilkins et al (2002, p.147)	New Zealand		Drug policy	
	0.5g	Wilkens et al (2005, pp.228-9)	New Zealand	Factual statement	Drug policy	

Table A4: Estimates of e (Papers per Joint)

Parameter	Estimate	Source	Country	Population	Source Field	Notes
Number of papers per joint	Single rolling paper	UNODC (2006, p.43)	UK	Factual statement	Drugs policy	"cannabis cigarettes smoked in the United Kingdom and in Ireland are typically mixed with tobacco and a single rolling paper is used"
	3 cigarette papers	McBride (1995, p.30)	UK (Wales)	Community drug clinic attendees	Psychiatry	Reflects modal response. Most frequent responses: 3 cigarette papers (75.3% of responses), 5 cigarette papers (20.6%). Mean: 3.36 cigarette papers.

Table A5: Other Estimates

Tobacco content per joint	Joints sold in coffee shop sales: 0.1 cannabis, 0.9g tobacco	UNODC (2006, p.43)	NL	Coffee shop sales. factual statement	Drugs policy	
	1-to-2 = herbal-to-tobacco	Tyle (1995, p.145)	UK	Factual statement	Drugs policy	"The ratio of grass to tobacco is usually two to one in favour of the tobacco." (p.145)
	0.79 of a king size cigarette (mean)	McBride (1995, p.30)	UK (Wales)	Community drug clinic attendees	Psychiatry	Most frequent responses: 3/4 of a king size cigarette (42.3%), 1 king size cigarette (34.0%)
Total weight of joint (including cannabis and tobacco)	0.867g; 0.915g	Chait et al. (1989, p.64)	US	NIDA marijuana cigarette analysis	Psychiatry	Joints produced for research and medical use. Length 85mm. Numbers correspond to joints of different potency
	0.927	Own calculation based on Fairbairn et al. (1974)	UK (London, Leeds)	Lab study using volunteer samples	Pharmacology	N=23. Sample average including both herbal and resin.