

IFN Working Paper No. 1536, 2025

Sentence Length and Recidivism: Court Rulings based on BAC

Joakim Jansson, Per Pettersson-Lidbom, Mikael
Priks and Björn Tyrefors

Sentence Length and Recidivism: Court Rulings based on BAC*

Joakim Jansson,^{*} Per Pettersson-Lidbom,[♥] Mikael Priks[♠], Björn Tyrefors^{*}

September 14, 2025

Abstract

We study the effect of prison sentences on recidivism using a unique feature of sentencing for drunk driving in the Swedish court system. Below the blood alcohol concentration (BAC) of 1.0‰, individuals are never sentenced to prison and above 1.0‰, the average number of days sentenced to prison is essentially linearly increasing with the BAC level. We find that being sentenced to prison for one month reduces reoffending in the next five years by approximately 80 percent.

Keywords: Sentence length, Recidivism, DUI

JEL codes: K42, R41, C26

* We thank Randi Hjalmarsson and Matthew Lindquist for useful comments. Björn Tyrefors is grateful for financial support from Jan Wallanders och Tom Hedelius stiftelse (grant P23-0186).

^{*} Department of Economics and Statistics, Linnaeus University and the Research Institute of Industrial Economics (IFN), email: joakim.jansson@ifn.se.

[♥] Department of Economics, Stockholm University and the Research Institute of Industrial Economics (IFN), email: pp@ne.su.se.

[♠] Department of Economics, Stockholm University, email: mikael.priks@su.se.

^{*} Department of Economics, Gothenburg University and the Research Institute of Industrial Economics (IFN), email: bjorn.tyrefors@economics.gu.se.

1. Introduction

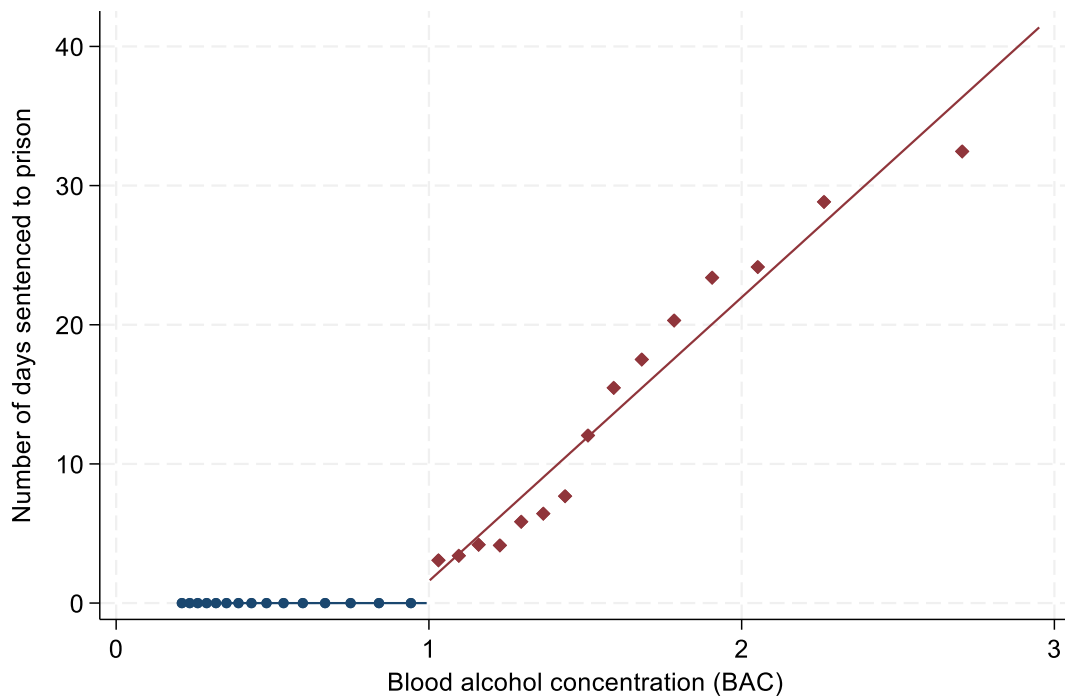
An oft-revisited question in criminology, as well as in economics and in the public debate, is whether the severity of punishment should discourage individuals from committing crimes. Despite a large amount of empirical research, there are surprisingly few studies that find a substantial deterrence effect of the severity of punishment (e.g., Durlauf and Nagin 2011; Chalfin and McCrary, 2017; Loeffler and Nagin, 2022; Hjalmarsson, 2024 and Alsan et al. 2025).¹ One potential reason for the lack of evidence is that most previous research has studied individuals who are not deterred by harsher sentences since they have typically committed several crimes before, as discussed by Durlauf and Nagin (2011), for example. In this paper, we will take a novel approach by studying individuals from the general public who may be considerably more responsive to the severity of punishment than hardened criminals.

We also make use of a novel identification strategy, which allows us to analyze the deterrence effect of the severity of punishment, i.e., very short prison sentences for driving under the influence (DUI) on recidivism. To analyze the dose-response relationship between the length of imprisonment and reoffending, we introduce a new instrumental variable approach. The instrument is constructed based on a *piecewise* linear feature of the Swedish legal system for DUI. Drivers who are caught with a blood alcohol content (BAC) lower than 1.0‰ are never imprisoned, while the average number of days sentenced to prison (the treatment dose) is essentially increasing *linearly* for drivers with a BAC value larger than 1.0‰.

Figure 1 shows a binned scatter plot of this bivariate relationship between sentence length and BAC used by Swedish courts. Such a scatterplot maps out the *nonparametric* relationship between these two variables, which allows researchers to directly assess the shape of the relationship, i.e., whether it is linear, as discussed by Starr and Goldfarb (2020) and Cattaneo et al. (2024). The plot in Figure 1 is based on the universe of 79,569 drunk drivers, 11,774 of whom are females, during the years 2008 to 2022. It reveals a (broken) hockey-stick pattern between BAC and average length in prison. The slope coefficient is zero up to the BAC-level 1.0‰, above which there is an abrupt change followed by an approximately linearly increasing function in sentence length for BAC above 1.0‰.

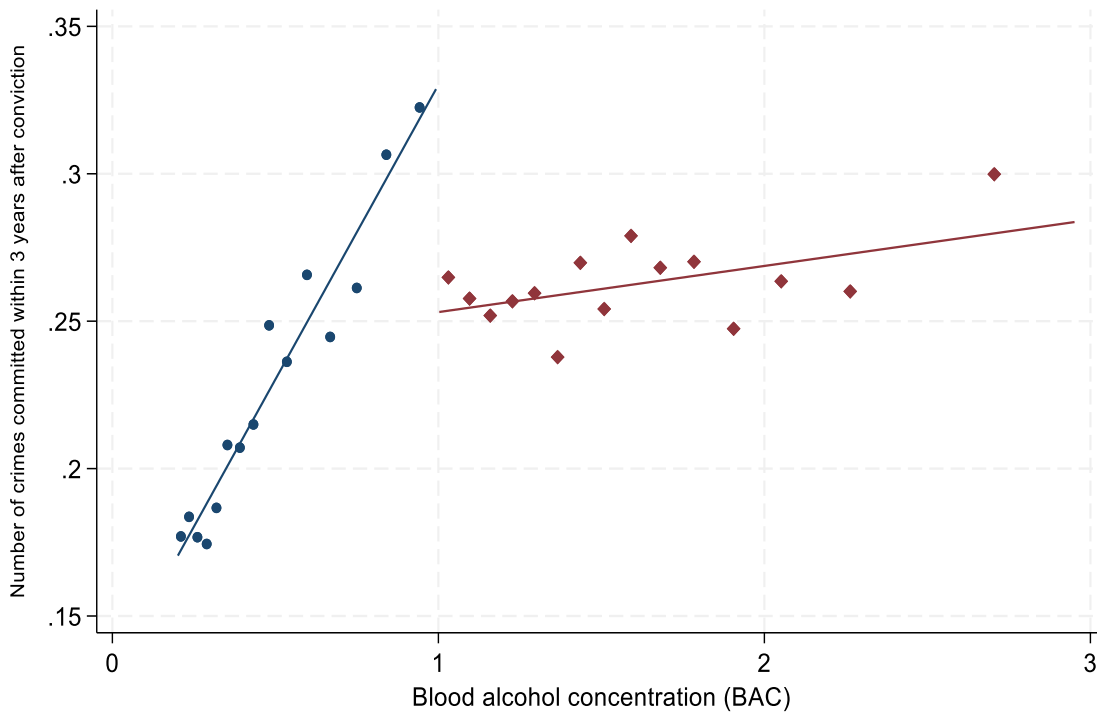
¹ In criminology, research about the severity of punishment is often classified into specific deterrence (i.e., own experience) or general deterrence (i.e., vicarious experience). However, since an individual might learn from both own experiences as well as from others experiences, specific and general deterrence cannot be separated from each as argued by Tomlinson (2016), Stafford and Warr (1993) and Paternoster and Piquero (1995), for example. For this reason, we refrain from interpreting our results as specific or general deterrence.

Figure 1. Binned scatterplot of the average number of days in prison and BAC levels



Notes: The binscatter plot was created by the Stata command binsreg (Cattaneo et al. 2025) with 15 quantile-spaced bins above and below BAC=1%. It also adds linear prediction lines to the plot. The slope coefficient is 0 below BAC=1%, while it is approximately 21 days in prison per 1.0 BAC level above.

Figure 2. Binned scatterplot of the average number of future crimes and BAC levels



Notes: The binscatter plot was created by the Stata command binsreg (Cattaneo et al. 2025) with 15 quantile-spaced bins above and below BAC=1%. It also adds linear prediction lines to the plot. The slope coefficient is approximately 0.19 crimes/BAC below BAC=1%, while it is 0.02 crimes/BAC above.

Intriguingly, the binned scatter plot relationship between BAC and reoffending (i.e., number of crimes committed within 3 years after conviction) shown in Figure 2 also features a similar *piecewise* linear relationship with an abrupt change at 1.0‰. The figure actually reveals a strong upward-sloping linear relationship for BAC values below 1.0‰, suggesting a strong positive correlation between alcohol consumption and crime when individuals cannot be sentenced to prison.² However, this relationship is significantly weakened above 1.0‰, where the relationship shifts to a nearly flat linear relationship for BAC above 1.0‰. In other words, Figure 2 also features a (broken) hockey-stick pattern, but a reversed one as compared to Figure 1.

We will argue that the two relationships displayed in Figures 1 and 2 must be causally related due to the close correspondence between the two figures. Indeed, under the assumption that average treatment effects are constant, there must exist a proportionality restriction between the slope coefficient in the reduced-form relationship (RF) and the slope coefficient in the first-stage relationship (FS) according to the logic of the classical linear instrumental variable method ($RF = \beta^{IV} \times FS$), as discussed by Kolesár (2013), for example.³ We therefore construct our instrument based on the very sharp changes in the slope coefficients of both the first-stage relationship and the reduced-form relationship when the BAC level is larger than 1.0‰. In other words, the interaction between an indicator variable of having a BAC level above 1.0‰ and the BAC level itself constitutes our instrumental variable.

The IV estimate of the dose-response relation between the length of imprisonment and recidivism can also be readily estimated based on the information about the four slope coefficients in Figures 1 and 2. It is simply the ratio between the difference in the two slope coefficients for being above or below the threshold 1.0‰, i.e., $\frac{0.02 - 0.19}{21 - 0}$. Thus, the IV estimate is equal to -0.008 .

This estimate is an order of magnitude larger than found in most previous studies. It corresponds to a reduction of 80 % in reoffending if an individual is sentenced to prison for one month as compared to the control group that is sentenced to pay fees only (i.e., -0.008 times 30 days divided by the average number of historic crimes just below BAC equal to 1.0‰, which is approximately 0.33). The treatment effect is found for all types of crime, where traffic-

² A substantial body of research in economics, criminology and public health documents a positive association between alcohol consumption and crime (see, e.g., Carpenter and Dobkin, 2010)

³ In the textbook model of IV, the proportionality restriction is typically derived under the assumption of a constant treatment effect, i.e., no heterogeneity in the treatment effect. However, Kolesár (2013) also shows that it holds under treatment heterogeneity as long as that the average treatment effect is constant.

related crime amounts to 54 percent of the effect, drug-related crime to 27 percent, violent crime to 12 percent and property crime to 6 percent. These effects on reoffending are not temporary but persist for at least 5 years after conviction. Consequently, incapacitation cannot explain these results since prison sentences are too short. We also find similar deterrence effects for men and women, which lends support to that average treatment effects are constant.

The obvious question is whether our proposed instrument is valid. To start with, the instrument is extremely strong statistically with an F -value above 5,000. Second, our instrument fulfills the monotonicity assumption by construction since there is just going to be a one-sided non-compliance problem in our IV design. This is because individuals below 1.0‰ constitute the control group since they are never sentenced to prison, while individuals above 1.0‰ constitute various treatment groups since they receive different treatment doses, i.e., sentence length, based on their BAC level. Third, we argue that our instrumental variable is also locally random, conditional on the BAC level. In other words, there is essentially a randomized experiment for each separate BAC level, and since the instrument is continuous, we have a continuum of experiments. Local randomness of the instrument around some fixed value of BAC is also supported by the fact that drivers cannot precisely predict their BAC level, since it, apart from the exact amount of alcohol consumption, depends on when and how much food has been consumed, body weight, time between consumption and police control, etc. The police are also unable to thwart the randomness in the BAC measure due to automated instruments for measuring BAC. We argue that the exclusion restriction is justifiable since there are no other treatments or confounders that have the same kind of hockey-stick relationship in treatment dose as displayed in Figure 1.

It is important to stress that our empirical design is based on a discontinuous (unconstrained) threshold regression model that is distinct from both a regression discontinuity (RD) design and a regression kinked (RKD) design (see, e.g., Hansen, 2022, for a discussion of different threshold models). A regression discontinuity design exploits that there is a discontinuous change in a treatment at a particular threshold value of the assignment variable to identify a binary treatment effect of interest. In contrast, we use a multivalued treatment setup (i.e., a dose-response relationship) and leverage the fact that there are discontinuous changes in (linear) slope coefficients above and below the BAC threshold along the whole range of values taken by the assignment variable. The attractiveness of our approach, as compared to an RD design, is that we, in this setting, can circumvent the violation of the exclusion restriction due to the problem of having multiple treatments (imprisonment and probation) that both change discontinuously at the 1.0‰ BAC threshold. In other words, we impose the

proportionality restriction to solve the problem with the violation of the exclusion restriction at the RD threshold. As a byproduct, our IV design also improves on external validity compared to an RD design since it is valid for a much broader population than only individuals close to the threshold. An RKD design relies on continuity at the kink (e.g., Card et al., 2015 and Hansen, 2017) and it is not therefore applicable in our setting since there is not only a kink but also a discontinuous jump in our treatment variable (length of sentence) at the threshold of 1.0‰.

We perform several validity and robustness checks of our IV design. First, we perform a refutability test based on a subpopulation that is not affected by the treatment, as discussed by Angrist and Krueger (1999). In fact, according to Swedish law, some occasions do not lead to a prison sentence even if drivers are caught with a BAC level above 1.0‰. One example is when an individual re-parks the car at a parking lot at night. Since there is no first-stage relationship for this subsample, there should not be any hockey-stick relationship in the reduced form between recidivism and BAC, but only a linear one. This is indeed what we find.

To test whether the instrument is likely to be as good as randomly assigned, conditional on the BAC level, we control both parametrically and non-parametrically for two of the most important predictors of crime: age and prior criminal history. The estimates are only slightly affected by the inclusion of these control variables, which suggests that the instrument is as good as randomly assigned. Relatedly, we also control for higher order of polynomials for the BAC level to assess the sensitivity to omitted variable bias caused by functional form misspecifications. We again find similar effects.

Finally, we perform an overidentification test to simultaneously assess both violation of the exclusion restriction and whether the assumption of a constant average treatment effect is valid. We perform this test by dividing our linear continuous instrument into 15 binary and mutually exclusive instruments and estimate this specification using a two-stage least squares (TSLS) estimator. Reassuringly, the TSLS estimator produces a similar result and we cannot reject the overidentifying restriction test.

Nonetheless, we also tried other instrumental variable estimators that have better properties than TSLS when it comes to many and weak instruments or certain types of violations of the exclusion restrictions, as discussed by Kolesár (2013) and Kolesár et al. (2015). Specifically, we used four other types of IV estimators: a limited-information-maximum-likelihood (LIML) estimator, a jackknife instrumental variables estimator (JIVE), an unbiased jackknife IV estimator (UJIVE), as well as a modified-bias-corrected two-stage

least squares (MBTSL) estimator. We find almost identical results across all IV estimators, which strongly suggests that we have estimated a causal treatment effect.

Why do we find such large effects when most of the previous literature does not? One explanation could be that European prison sentences involve more rehabilitation than U.S. prisons, and rehabilitation has been shown to reduce future crime, as discussed by Alsan et al. (2025) and Bhuller et al. (2020), for example. However, there is essentially no rehabilitation for individuals with short prison sentences during the time period we study.⁴

Instead, as mentioned above, we argue that a more plausible explanation is that we study a population consisting of individuals from the general public who, on average, have committed much fewer crimes rather than prolific offenders who may be less deterred by punishment. Moreover, in ordinary prisons, the potential deterrence effect is counteracted by criminogenic peer effects as well as negative labor-market effects, which tend to increase crime. This is not the case here since the inmates are placed in low-security prisons where the peers are not convicted for serious crimes and DUI is legally not a just cause for termination of employment. Thus, these mechanisms are unlikely to affect our results.

Another piece of evidence that lends support to our explanation is that we find similar magnitudes of the deterrence effect if the individual has served their sentence in prison or in the home environment being surveilled at all times with electronic monitoring (EM).⁵ Thus, the findings speak in favor of interpreting the effect as driven by sentence severity (set by the courts) rather than by facility assignment by the Swedish Prison and Probation Service. Hence, this result is in contrast to the literature on electronic monitoring, which typically has shown that recidivism is reduced more with EM than spending time in prison (e.g., Di Tella and Schargrotsky, 2013; Henneguette et al., 2016; Williams and Weatherburn, 2020 and Grenet et al. 2024). Again, this is likely due that negative peer- and labor-market effects are absent in the prisons we study.

We are not the first to study a population of drunk drivers. Hansen (2015) also analyzes the effect of BAC on sentencing and crime. He uses data from the state of Washington and exploits two BAC thresholds, above which the punishment, either in the form of imprisonment, probation or fees, or a combination, increases discontinuously. There are, however, some

⁴ The information about rehabilitation programs was provided by Martin Lardén, the Head of the Prison and Probation Service's treatment operations.

⁵ The Swedish Prison and Probation Service, may, under some criteria (e.g., acceptable accommodation, regular employment, drug and alcohol testing, household consent, and permission to home visits by probation officers), allow the individual to serve the sentence in the home environment being surveilled at all times with electronic monitoring (EM).

important differences between his and our study regarding both the study population, the empirical approach and the results. First, Hansen examines a population with an average BAC of 1.64‰, while our population has an average BAC of 0.9‰, i.e., only marginally above 0.8‰, which, for example, is the legal limit of DUI thresholds used by many U.S. States and England. This makes our population different from Hansen (2015). Second, Hansen (2015) uses an RD design to estimate a binary treatment effect, while we use an instrumental variable where we exploit a piecewise linear change in a policy rule to estimate multivalued treatment effects. Third, an RD design can only be used to analyze the reduced-form effect between harsher punishments and reoffending since there are multiple treatments that change discontinuously at the thresholds and thus violate the exclusion restriction of a fuzzy RD design.⁶ In contrast, our IV method offers the possibility to estimate the causal effect of sentence length on recidivism. Fourth, we find deterrence effects on all types of crime, while Hansen finds effects only on future drunk driving.

Our paper also adds to the literature that tries to use various identification strategies to estimate the causal relationships between harsher prison sentences and future crimes. One literature uses judge-fixed effects to get exogenous variation in the punishment (see, e.g., Green and Winik, 2010; Aizer and Doyle, 2015; Mueller-Smith, 2015 and Bhuller et al. 2020) Another literature has used various sanction reforms (see, e.g., Raphael and Ludwig, 2003; Vollaard, 2013 and Kuziemko, 2013). Some studies use discontinuities in punishment (see, e.g., Lee and McCrary, 2009; Hjalmarsson, 2009 and Rose and Shem-Tov, 2021). Most studies use binary treatments while only a few (e.g., Rose and Shem-Tov, 2021) estimate multi-valued treatments or dose-response functions using two-stage least squares, like we do.⁷ While the results vary across studies, most show limited deterrence effects.

In the following, we describe the background and data in Section 2. We then present the empirical design in Section 3, the results in Section 4, specification checks in Section 5, gender differences in Section 6, electronic monitoring vs. prison in Section 7 and concluding remarks in Section 8.

⁶ Interestingly, we also find a reduced form effect of approximately 0.8 at the threshold of 1.0 ‰ which is directly noticeable in Figure 2. This reduced form effect from the RD design is much larger than those found in Hansen (2015), which is approximately 0.2.

⁷ Multivalued treatments (e.g., the length of the prison sentence) are often binarized (i.e., an indicator for any prison sentence) despite the fact that it generally would lead to bias as discussed by Angrist and Imbens (1995), Marshall (2016) and Andresen and Huber (2021). Rose and Shem-Tov (2024) shows that recoding an ordered treatment into an indicator is no longer problematic under the assumption of “extensive margin compliers only”.

2. Background and Data

In this section, we describe (i) the Swedish laws against driving under the influence (DUI), (ii) the workings of BAC tests, (iii) the enforcement of a prison sentence, (iv) and the data we use in our empirical analysis.

The Swedish laws against driving under the influence

The provisions on DUI and aggravated DUI are found in the Act (1951:649) on punishment for certain traffic offences (the Traffic Offences Act).⁸ The Act states that if the BAC level is between 0.2‰ and 1.0‰, it is considered a “normal” DUI, while if it is equal to or above 1.0‰, it is considered an “aggravated” DUI.⁹

In the case of a normal DUI, a driver is fined and the penalty is usually imposed by way of summary imposition of a fine (“strafföreläggande”) without a formal court hearing. However, if the DUI is aggravated, then a court decides the penalty and the driver may be sentenced to one of two distinct treatments: either prison or probation.¹⁰ In the case of a prison sentence, the length of imprisonment ranges from 14 days up to two years, but most individuals are sentenced to one or a couple of months of imprisonment. In the case of a probation sentence, the trial period is typically set to two years. The court also takes into account whether the BAC level was below or above 1.5‰. If it was above, the probability of being sentenced to imprisonment is higher.


To fully describe the underlying sentencing behavior of the Swedish court when it comes to DUI, we will decompose the results in Figure 1 into two separate figures. Figure 3 plots the relationship between the BAC level and the probability of being sentenced to imprisonment and Figure 4 plots the relationship between the BAC level and sentence length, conditional on being sentenced to prison. Thus, by multiplying the probability of being sentenced to imprisonment displayed in Figure 3 with the length of sentence in Figure 4, we get the results in Figure 1. For example, the average of the first bin after the BAC level 1.0‰ in Figure 1 is approximately equal to $0.1 \times 31 = 3.1$.

⁸ Section 4 of the Act contains provisions on drunk driving of the so-called “normal” degree and Section 4a contains the rules on aggravated drunk driving.

⁹ These are relatively strict limits compared to most other industrialized countries that have set their legal limit at 0.5‰ or 0.8 ‰. Historically, Sweden have also been an early proponent of introducing tight legal limits on DUI. For example, a lower limit of 0.8‰ and an upper limited of 1.5‰ was implemented as early as 1941. In 1957, the lower limit was decreased to 0.5‰. In 1990, the lower limit was further decreased to 0.2‰ while the upper limit was decreased to 1.0‰ in 1994.


¹⁰ We use probation for both “villkorlig dom” and “skyddstillsyn.”

Figure 3. Binned scatterplot of the probability of sentenced to in prison and BAC levels



Notes: The binscatter plot was created by the Stata command binsreg (Cattaneo et al. 2025) with 15 quantile-spaced bins above and below BAC=1%.

Figure 4. Binned scatterplot of length of imprisonment and BAC levels



Notes: Notes: The binscatter plot was created by the Stata command binsreg (Cattaneo et al. 2025) with 15 quantile-spaced bins above BAC=1%. The number of days in prison is zero below BAC=1%.

Taken together, Figures 3 and 4 therefore reveal that the linear increase in average sentenced length as shown in Figure 1 is a combination of an increased probability of being sentenced to imprisonment (Figure 3) and longer prison sentences (Figure 4).

Importantly, expressing our treatment variable in terms of days sentenced to prison permits identification of a weighted average of the effect of one additional day in prison as discussed by Angrist and Imbens (1995) and Bhuller et al. (2020). Thus, this treatment effect captures a convex combination of the extensive margin effect of going to prison and the intensive margin effects of longer sentencing. However, in the case of a linear dose-response function, the treatment effect is equal to the effect of the intensive margin only.

The workings of BAC tests

Sweden has not only a relatively low DUI limit but also closely monitors its compliance. In fact, the Swedish police conducted approximately 1.5 million screening tests for DUI in 2024. A driver has to perform a screening breathalyzer test when caught by a police control officer. The test signals positive if the BAC level is above 0.2‰. If the test result turns out to be positive, the individual is required to take two additional Breathalyzer tests. The proof material in court is the average of these two new tests, less a 0.15‰ reduction (to account for, e.g., consumption of food or medicine with small amounts of alcohol). If the breathalyzer test is refused, the individual must go to the police station for a blood test instead. A blood test is also taken if there is reason to suspect an intake of illegal drugs. Importantly, there is no reason to believe precise sorting around certain BAC values. This is because individuals cannot determine exactly how a certain number of units transforms into different BAC levels. This depends on many factors, such as when and how much alcohol as well as food was consumed, body weight, etc. (NHTSA, 2016). It is, of course, also not possible to know when and where the police control takes place. The police, on their side, cannot manipulate the results from the equipment, which is automated. Interestingly, most DUIs happen between 10 and 12 in the morning, suggesting that the alcohol has been consumed the previous evening or during the night.

If a driver commits a DUI, then the driver's license may be revoked. This is handled by the Swedish Transport Agency (“Transportstyrelsen”). The limit for a driver's license to be confiscated in the event of drunk driving is 0.32‰. The suspension period is typically between 12 and 24 months, depending on the circumstances.

The enforcement of a prison sentence

The execution of prison sentences is handled by the Swedish Prison and Probation Service (Kriminalvården), which is a different governmental agency from the courts. A prison

sentence must be enforced as soon as possible. This typically means within three months after the sentence has become finalized and can no longer be appealed.

According to government legislation, if sentenced to a maximum of six months in prison, the individual can, under certain circumstances, serve the sentence under intensive supervision with electronic monitoring. The individual is then monitored 24 hours a day with the aid of a transmitter attached to the ankle. Individuals can apply for intensive supervision if they meet certain requirements, i.e., having a suitable accommodation and stable employment or a commitment to seek and obtain employment. They must also accept alcohol- and drug testing, and home visits from the Prison and Probation Service. The sentenced then has curfew at all times apart from certain pre-scheduled activities, such as work and studies. In fact, approximately 70 percent of those sentenced to imprisonment due to aggravated DUI actually serve it at home with electronic monitoring. However, we do not condition our empirical analysis based on whether the individual served their time in prison or not, since this is partly an individual choice variable and therefore potentially endogenous. Consequently, we define our treatment variable based on the formal sentences decided by Swedish courts rather than implementation decisions made by the Prison and Probation Service. Reassuringly, however, we provide suggestive evidence that both prison and electronic monitoring have similar effects on recidivism in section 6.

There are two major prison establishments in Sweden, “Östragården” and “Sörbyn”, where those sentenced for DUI are typically placed. They have a capacity of holding about 150 inmates altogether. These establishments have the security class 3, on a 1 to 3 scale, where 3 is the most open establishment and where no concrete measures are taken to stop escapes. In contrast to higher security prisons, inmates share rooms, they can spend much time outside and the facilities are only locked at night. This type of prison is similar to minimum security prisons in the U.S, which hold 14 percent of the prisoners (Federal Bureau of Prisons, 2025).

Data

Our data is collected from three different sources: the Swedish National Council for Crime Prevention (“BRÅ”), the National Forensics Centre (“NFC”), and The Swedish Prison and Probation Service (“Kriminalvården”).

The Swedish National Council for Crime Prevention provided data on crimes. This data includes individual information on age at crime, type of crime, gender, etc. NFC provided data on breathalyzer tests and data on blood tests. We have collected this data for the period between 2008 and 2022. Importantly, NFC only keeps the data for the 5 most recent years, so this means that we have had to collect this data at multiple points in time (2013, 2018 and 2023). The

Swedish Prison and Probation Service provided data on whether the individual was serving the sentence with electronic monitoring or in prison.

Our original sample is based on breathalyzer and blood tests from NFC. This data is merged to the data on crimes from the Swedish National Council for Crime Prevention and the Swedish Prison and Probation Service. We make several sample restrictions so that we can confidently link the individual's BAC level to the legislation on DUI. As a result, we only include individuals who have consumed alcohol and not used other drugs. We only analyze individuals who have committed only DUI and no other crime at the time of the BAC test. We analyze individuals who are 18 years or older since 18 is the legal limit for driving in Sweden. Consequently, we end up with 79,569 individuals (67,950 men and 11,774 women) for the period 2008 - 2022.

Figure 5 shows the frequency of the different BAC levels in our data set. Most individuals who are caught for drunk driving had BAC levels between 0.2‰ and 0.5‰, but as many as 30,052 individuals had a BAC equal to or larger than 1.0‰.

Table 1 presents summary statistics for the full sample but also for subsamples above and below the BAC threshold for aggravated DUI of 1.0‰, since our empirical design will exploit the difference in sentencing policy for those individuals with a BAC level above and below 1.0‰. There are 30,044 individuals with a BAC equal to or larger than 1.0‰ and 49,525 below. Table 1 shows summary statistics for the pretreatment variables (age at crime, gender, prior criminal history) as well as for the three dependent variables we use: the number of crimes committed one year after DUI, three years after DUI and five years after DUI. The dependent variables measuring the degree of recidivism for a certain time interval are prespecified by the Swedish National Council for Crime Prevention.¹¹ Thus, our hands are tied when it comes to the choice of how to measure reoffending. Table 1 reveals that the two subsamples are broadly similar regarding age, gender, and number of historic and future crimes, which lends support to that they are comparable.

¹¹See link

<https://bra.se/download/18.5e0f78b192bd39b23281c3/1730464445222/Variabelbeskrivning%20Lagf%C3%B6ringsregistret.pdf>

Figure 5. The number of observations based on levels of Blood Alcohol concentration (BAC).

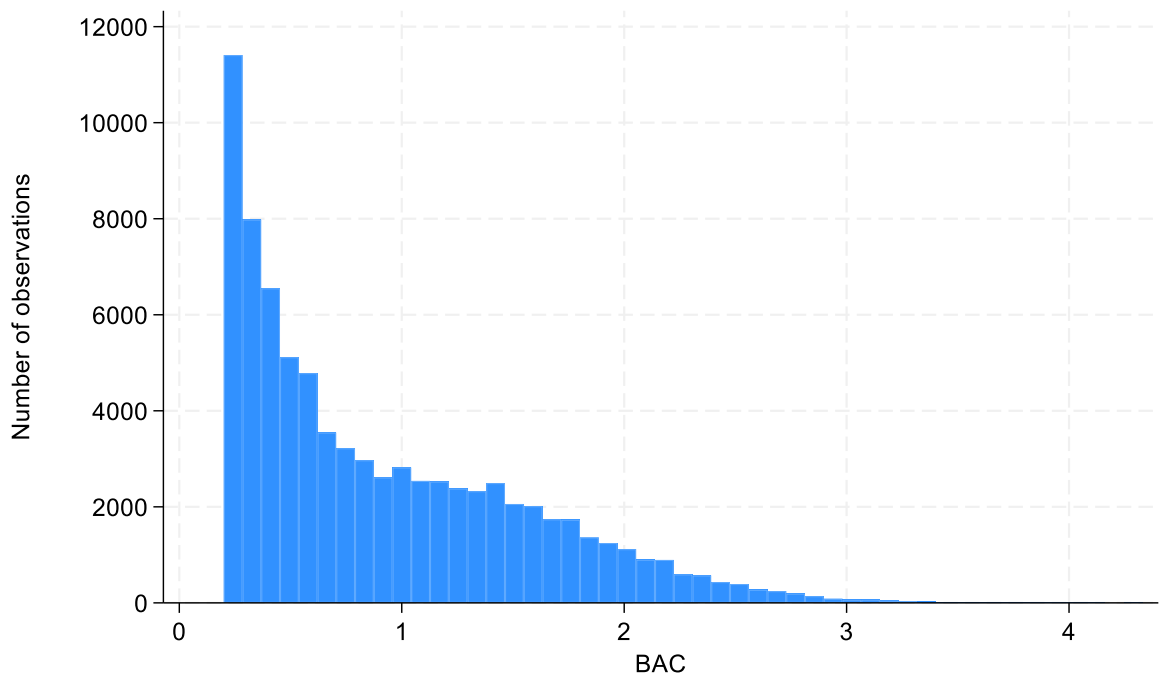


Table 1. Summary statistics

	Full sample	Below BAC 1.0‰.	Above BAC 1.0‰.
Age at crime	47	47	45
Percentage of females	15	13	17
Number of crimes committed within 1 year after conviction	0.08	0.08	0.08
Number of crimes committed within 3 years after conviction	0.22	0.21	0.24
Number of crimes committed within 5 years after conviction	0.32	0.31	0.35
Total number of historic crimes	1.70	1.50	2.02
Number of observations	79,569	49,525	30,052

Empirical Design

In this section, we describe our empirical approach. The population regression of interest is

$$(1) \quad Y_i = \alpha + \beta X_i + u_i,$$

where Y_i is a measure of reoffending (e.g., number of crimes committed during a certain time period after conviction) for individual i and X_i is the number of days in prison for individual i . The parameter of interest is β , which measures the population average treatment effect, i.e., a

linear dose-response relationship of the length of imprisonment on future crime activity. Thus, β measures how much an extra day in prison affects recidivism. If $\beta < 0$, then longer sentences reduce criminal activity, while if $\beta > 0$, then longer sentences increase criminal activity.

The problem of estimating (1) is that X_i is likely to be endogenous since individuals who are sentenced to different lengths in prison have different unobserved characteristics and are therefore not comparable. In fact, when we estimate equation (1) in our sample of drunk drivers, we find a positive and statistically significant effect (0.00072 with a standard error of 0.00032), which is not surprising given the well-established correlation of alcohol consumption and crime found in the literature as mentioned above. This finding strongly suggests that OLS is biased since those with longer sentences commit more crimes due to unmeasured individual confounders related to alcohol consumption.

To solve this endogeneity problem, we need to find a valid instrumental variable. We will use a policy rule used by Swedish courts in the case of DUI as an instrumental variable, as discussed previously. This policy rule determines the number of days in prison depending on the level of the blood alcohol concentration (BAC). If the BAC is below 1.0‰, the individual is never sentenced to jail, while if the BAC is above 1.0‰, the average number of days sentenced to prison increases almost linearly with the BAC level, as displayed in Figure 1. The figure also shows that the slope coefficient increases by approximately 24 days per 1 BAC level and that there is a discontinuous jump in the average number of days in prison at 1.0 of approximately 3 days.

We can now express our instrumental variable approach by two equations: the first-stage equation

$$(2) \quad X_i = \pi_0 + \pi_1 BAC_i + \pi_2 D_i + \pi_3 BAC_i \times D_i + v_i,$$

and the reduced form equation

$$(3) \quad Y_i = \delta_0 + \delta_1 BAC_i + \delta_2 D_i + \delta_3 BAC_i \times D_i + e_i,$$

where D_i is an indicator variable taking the value 1 if $BAC_i \geq 1.0\text{‰}$ and zero otherwise. The first-stage equation (2) is therefore simply a parametrization of the policy rule used by Swedish courts since it is a spline function defined by two piecewise linear segments, ($\widehat{\pi}_1 = 0$ and $\widehat{\pi}_3 \approx 21$) with one threshold jump at 1.0 ($\widehat{\pi}_2 \approx 3$). In other words, equation (1) is a discontinuous (unconstrained) threshold regression model, as discussed by Hansen (2022), for example.

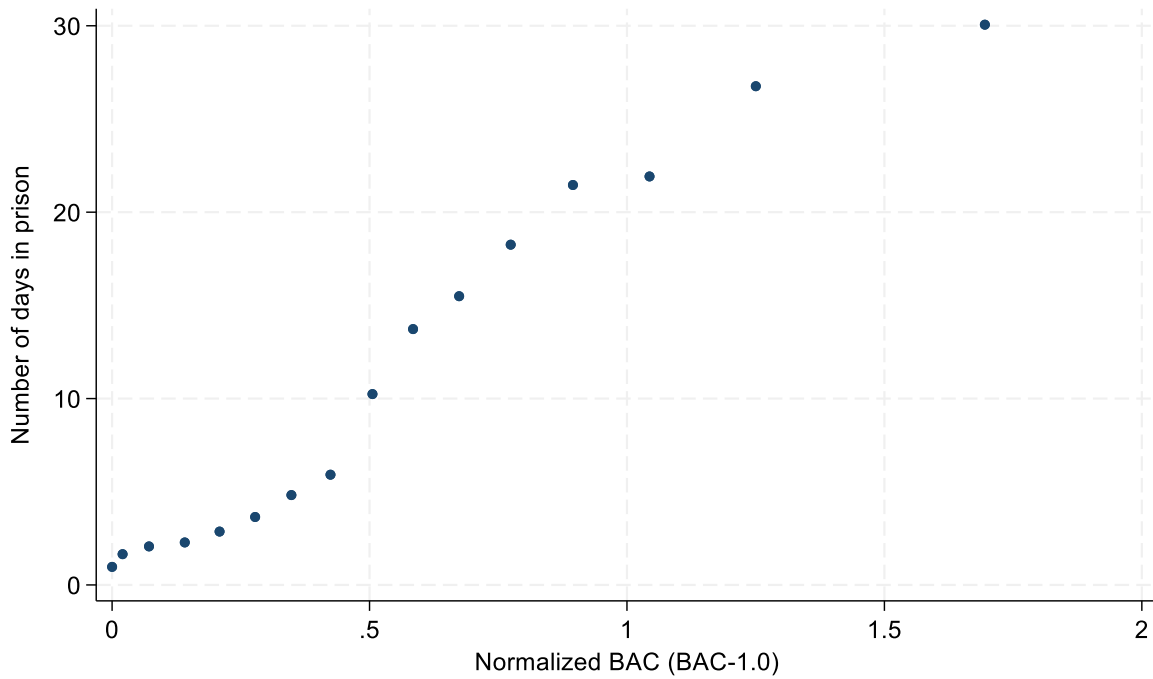
We will use $BAC_i \times D_i$ as our instrumental variable while simultaneously controlling for both BAC_i and D_i . Thus, we argue that our instrument is as good as randomly assigned and fulfills the exclusion restriction conditional on both BAC_i and D_i , i.e., $E[u | BAC \times D, BAC, D] = E[u | BAC, D]$. The reason why we need to control for not just BAC_i but also D_i is that there are two treatments (prison and probation) that discontinuously change at the BAC-level of 1.0‰. In other words, by including D_i , we are effectively holding the effect of probation on reoffending constant, i.e., the exclusion restriction is likely to be violated if not controlling for D_i . Consequently, the inclusion of D_i also implies that we are not exploiting the discontinuous jump in the treatment at the BAC level 1.0‰ as an exogenous source of variation. Our IV design is therefore distinct from an RD design. Moreover, since there is a discontinuous jump in treatment at 1.0‰, a regression-kink design is not applicable since it requires continuity at the threshold (e.g., Card et al. 2015 and Hansen 2017).

Our proposed IV estimator can now be expressed as $\beta = \delta_3 / \pi_3$. This means that the reduced form is proportional to the first-stage with a proportionality factor β . The proportionality restriction has been derived by assuming a constant and linear treatment effect, but, as Kolesár (2013) shows, it also holds as long as the *average* treatment effect is constant. Thus, the assumption of a constant average treatment effect restricts the treatment effects in two ways, as discussed by Kolesár. It requires that the source of heterogeneity in the individual treatment effects is unrelated to unobservables and it restricts the treatment effect to be linear. As a result, our identifying approach relies on linearity, i.e., that both the first-stage and the reduced-form relationships are approximately linear in parameters.

Nonetheless, if the constant average treatment is violated, we are still going to estimate a weighted average of causal responses to a unit change in treatment for those whose treatment status is affected by the instrument, as discussed by Angrist and Imbens (1995). They show that 2SLS applied to an IV model with variable treatment intensity, such as equation (1), has a LATE interpretation under treatment heterogeneity as long as monotonicity holds.


We will test whether linearity holds using the nonparametric binscatter approach developed by Cattaneo et al. (2024). Importantly, our instrumental variable is continuous in BAC so we do not need to rely on extrapolations for testing linearity. We will show binscatter plots for both the first stage and the reduced form, where we have conditioned on both BAC and D . Figure 6 displays the first-stage relationship, while Figure 7 shows the reduced form where we have plotted the normalized BAC ($BAC - 1.0\text{‰}$) on the x-axis. The evidence presented in both these graphs strongly supports that linearity approximately holds both for the first stage and the reduced form.

Figure 6. Binned scatter plot of the first-stage relationship conditional on BAC and D



Notes: The binscatter plot was created by the stata command binsreg (Cattaneo et al. 2025) with 16 quantile-spaced bins above BAC=1%. It displays a covariate-adjusted (D and BAC) binscatter.

Figure 7. Binned scatter plot of the reduced form relationship conditional on BAC and D



Notes: The binscatter plot was created by the stata command binsreg (Cattaneo et al. 2025) with 16 quantile-spaced bins above BAC=1%. It displays a covariate-adjusted (D and BAC) binscatter.

Another way of testing whether linearity and the exclusion restriction hold is to discretize our instrumental variable and to use an overidentifying restriction test. We create 15 binary and mutually exclusive instruments for each 0.1‰ BAC level as further discussed in section 6. We will also use other types of IV estimators that have better properties with many (and potentially weak) instruments than 2SLS, such as LIML, MBTSLS, JIVE, and UJIVE as noted in the introduction.

To test whether our instrumental variable is likely to be as good as randomly assigned conditional on BAC and D, we perform a number of other specification tests. We control for two of the most important factors related to recidivism: age at crime and the previous criminal history. We do this both parametrically and nonparametrically, i.e., one indicator variable for each age. To assess the sensitivity to omitted variable bias caused by functional form misspecifications of BAC, we also control for higher order polynomials of BAC. Finally, we perform a number of refutability or falsification exercises along the lines suggested by Angrist and Krueger (1999). These specification tests will be further discussed in the following section.

4. Results

In this section, we present the results from our IV approach. Table 2 shows the results for the first-stage effect (Column 1), the reduced-form effect (Column 2), as well as the IV estimate (Column 3) for individuals who are followed for three years after conviction.

The estimated first-stage slope coefficient is 21, which means that the average number of days sentenced to prison increases by approximately 2.1 days for each additional 0.1‰ level of BAC. This is the linear estimate of the slope coefficient of π_3 in the first-stage equation (2) and this estimate is therefore the best linear approximation to the binned scatter plot shown in Figure 6. The first-stage relationship is also very precisely measured with a t -statistics of nearly 73. The single instrumental variable is therefore extremely statistically strong since the F -statistics is 5,286.

Table 2. First-stage, reduced-form and instrumental-variable estimates

Dependent variables	Number of days in prison	Number of crimes committed within 3 years after conviction	Number of crimes committed within 3 years after conviction
	First-stage (1)	Reduced form (2)	IV (3)
<i>Independent variables</i>			
Instrument: $D \times BAC$	21 (0.29)	-0.17 (0.02)	
Number of days in prison			-0.0081 (0.0009)
Number of observations	79,569	79,569	79,569

Notes: This table reports the first-stage estimate (Column 1), the reduced-form estimate (Column 2) and the corresponding IV estimate (Column 3). The first-stage cluster robust F-statistics is 5,286. The standard errors reported within parenthesis are clustered at the individual level.

The estimated reduced form slope coefficient is -0.17, which means that the average total number of future crimes within three years is reduced by approximately 0.017 crimes for each additional 0.1% level of BAC. This estimate of -0.17 can be compared to the mean of 0.3 for the control group, i.e., those individuals who are never sentenced to prison with a BAC slightly below 1.0%. This implies that the reduced form effect is approximately 60%. Again, the estimate, -0.17, of the slope coefficient of δ_3 is the best linear approximation to the binned scatter plot shown in Figure 7. This slope coefficient is precisely estimated with a tight 95% confidence interval between -0.21 and -0.13.

Importantly, the binned scatter plots as displayed in Figures 6 and 7 suggest that the first-stage and the reduced-form relationships are both approximately linear. As a result, the proportionality restriction of the IV estimator is likely to hold, i.e., $\beta = \delta_3/\pi_3$. This means that the IV estimate is -0.0081 ($= -0.17/22.22$), which is displayed in Column 3 of Table 1. In other words, one month in prison decreases the total number of crimes by 0.24 (-0.0081×30). This also means that reoffending decreases by 80% when compared to the average number of 0.30 crimes committed by the control group after three years since conviction. The IV estimate is also highly statistically significant with a t-statistics larger than 8.

We also find similar results when we analyze both shorter (one year) and longer (five years) time periods since conviction. These results are displayed in Table 3 and show that recidivism decreases after only one year of the sentence. The effect of being sentenced to prison for one month reduces the total number of crimes by approximately 90% as compared to the

average number of crimes committed by the control group $((-0.0037 \times 30) / 0.12)$. The effect continues to hold for five years after convictions with an effect of approximately 80% for a one-month sentence $((-0.0112 \times 30) / 0.44)$. Thus, these effects on reoffending are highly persistent over time and incapacitation can therefore not explain our results.

Table 3. The effect of prison length on recidivism for the number of crimes committed within 1 year and 5 years after conviction

Dependent variables	Number of crimes committed 1 year after conviction	Number of crimes committed 5 years after conviction
Number of days in prison	-0.0037 (0.0004)	-0.0112 (0.0012)
Number of observations	79,569	79,569

Notes: The standard errors reported within parenthesis are clustered at the individual level.

Next, we analyze the effect of the length of imprisonment on different types of crimes. Table 4 displays the results from four types of crimes: traffic crimes in Column 1, drug-related crimes in Column 2, violent crimes in Column 3, and property crimes in Column 4. Panel A shows the results for the number of crimes committed within 1 year after conviction, while panels B and C show the corresponding results for the crimes committed within 3 years and 5 years after conviction, respectively. Starting with the results for the number of crimes committed within 1 year after conviction, Panel A reveals that all types of crimes are reduced, but to different magnitudes. Traffic crimes are reduced the most with an effect of -0.0023 and they therefore constitute 62% of the reduction in the total number of crimes of -0.0037 that we reported in Table 3. Similarly, drug-related crimes are reduced by 19%, violent crimes are reduced by 8% and thefts are reduced by 11% in relation to the total number of crimes. Again, Panels B and C show that these effects on the four different crime categories are highly persistent across time. All the estimates except one are also statistically significant at conventional levels or lower.

Table 4. The effect of prison length on recidivism for different types of crimes

Dependent variables	Traffic (1)	Drugs (2)	Violence (3)	Property (4)
Panel A: Number of crimes committed within 1 years after conviction				
Number of days in prison	-0.0023 (0.0003)	-0.0007 (0.0001)	-0.0003 (0.0001)	-0.0004 (0.0001)
Panel B: Number of crimes committed within 3 years after conviction				
Number of days in prison	-0.0044 (0.0006)	-0.0022 (0.0004)	-0.0010 (0.0002)	-0.0005 (0.0003)
Panel C: Number of crimes committed within 5 years after conviction				
Number of days in prison	-0.0057 (0.0007)	-0.0036 (0.0004)	-0.0014 (0.0002)	-0.0011 (0.0004)
Number of observations	79,569	79,569	79,569	79,569

Notes: This table reports the IV estimate for four different categories of crime: traffic in Column 1, drugs in Column 2, violence in Column 3 and property in Column 4. The standard errors reported within parenthesis are clustered at the individual level.

5. Specification checks

In this section, we perform several specification checks of our instrumental variable approach. Specifically, we assumed that the reduced form slope coefficient is equal to the first-stage coefficient multiplied by the average treatment effect, i.e., $\delta_3 = \pi_3 \times \beta$. This proportionality restriction can be violated for three different reasons: (i) the exclusion restriction fails, (ii) the random assignment fails, (iii) the average treatment is heterogeneous and non-linear.

One way of jointly testing for violation of any of these three reasons is to make an overidentifying restriction test. We perform such a test by discretizing our continuous instrument into a larger number of binary and mutually exclusive instruments. Specifically, we create 16 separate indicator variables for every 0.1% BAC level over 1.0 except for BAC values above 2.5%, i.e., $D_1=1$ if $BAC \geq 1.0$ & $BAC < 1.1$, and zero otherwise, $D_2=1$ if $BAC \geq 1.1$ & $BAC < 1.2$, and zero otherwise, ..., $D_{16}=1$ if $BAC \geq 2.5$.

Table 5 shows the first-stage (Column 1) and reduced-form results (Column 2) from using our 15 binary and mutually exclusive instruments. As a result, we need to exclude one of them in order to avoid perfect multicollinearity and we have chosen to exclude D_1 . The interpretation of the first-stage and reduced-form estimates is therefore relative to the omitted indicator D_1 . For example, the first stage estimate of D_2 is 0.96, which means that those with a

BAC value between 1.1‰ and 1.2‰ are sentenced to 1 more day in prison relative to those with a BAC value between 1.0‰ and 1.1‰. The reduced-form effect is similarly interpreted, namely those with a BAC level between 1.1‰ and 1.2‰ commit 0.029 fewer crimes than those with a BAC level between 1.0‰ and 1.1‰. We can now scale the reduced form effect with the first-stage effect to get the single IV or Wald estimate, which is being reported in Column 3. The corresponding Wald estimate is therefore -0.0305. Thus, Table 5 reports 15 distinct Wald estimates. Reassuringly, most of these single Wald estimates are of roughly similar magnitudes and also broadly similar to the IV estimate of -0.0081 reported for the single instrument in Table 2.

In addition, Table 5 shows that both the first stage and the reduced form are monotonically increasing in the BAC level. The results presented in Table 5 therefore lend strong support for that the proportionality restriction is likely to hold, at least approximately. We can also conduct a formal overidentifying restriction test by using two-stage least squares (2SLS). The 2SLS estimate is -0.0064 with a standard error of 0.0008. The 2SLS is therefore smaller than the estimate of -0.0081 from the single instrument in Table 2. However, and as discussed by Angrist and Imbens (1995), the 2SLS estimate provides one way of combining a set of mutually orthogonal binary instruments into a single new weighted average, where 2SLS gives more weight to individual Wald estimates that are closer to the center of the distribution of the instruments. Consequently, the 2SLS estimator gives more weight to the dummy instrument D_7 in Table 5 since this Wald estimate is -0.0064. Turning to a test of whether this set of instruments is weak, an F -test yields a value of 549, which suggests that the instruments are strong and that we have enough statistical power for the overidentifying restriction test to be useful. Reassuringly, we cannot reject the overidentifying restriction test since the J -statistic is 18.7 with a corresponding p-value of 0.18. We have also tried other types of more robust instrumental variable estimators in the case of many and weak instruments, such as LIML, MBTSLS, JIVE, and UJIVE. They all give nearly identical results as the 2SLS, which suggests that the 2SLS is reliable.

Table 5. First-stage, reduced form and instrumental variable estimates with many instruments

	First-stage (1)	Reduced form (2)	”Wald estimate” (3)
D ₂	0.96 (0.26)	-0.029 (0.019)	-0.0305
D ₃	1.54 (0.27)	-0.038 (0.020)	-0.0244
D ₄	3.00 (0.33)	-0.084 (0.019)	-0.0279
D ₅	4.85 (0.37)	-0.068 (0.022)	-0.0139
D ₆	11.58 (0.40)	-0.097 (0.021)	-0.0084
D ₇	13.96 (0.42)	-0.089 (0.026)	-0.0064
D ₈	16.72 (0.45)	-0.138 (0.023)	-0.0083
D ₉	18.98 (0.53)	-0.155 (0.025)	-0.0082
D ₁₀	19.82 (0.63)	-0.164 (0.027)	-0.0083
D ₁₁	21.44 (0.65)	-0.187 (0.028)	-0.0087
D ₁₂	24.06 (0.80)	-0.191 (0.038)	-0.0079
D ₁₃	25.37 (0.88)	-0.227 (0.035)	-0.0089
D ₁₄	27.76 (1.11)	-0.197 (0.041)	-0.0071
D ₁₅	30.29 (1.22)	-0.205 (0.043)	-0.0068
D ₁₆	30.17 (0.75)	-0.267 (0.040)	-0.0088
Number of observations	79,569	79,569	


Notes: This table reports the first-stage estimates (Column 1), the reduced-form estimates (Column 2) and the corresponding Wald estimates (Column 3). These estimates correspond to equations (2), (3) and (1) respectively. The first-stage cluster robust F-statistics is 549. The standard errors reported within parenthesis are clustered at the individual level.

Another specification test that we perform is a refutability or a falsification test, as discussed by Angrist and Krueger (1999), for example. The idea is based on finding a “sub-population in which the “treatment effect” should not be observed, either because the sub-population is thought to be immune to the treatment or did not receive the treatment” (Angrist and Krueger 1999, p. 1326). It turns out that, according to Swedish DUI law, there are some well-defined circumstances in the case of driving under the influence that are not considered

as aggravated even though the BAC level is larger than 1.0‰. In those cases, drunk drivers are only fined and never sentenced to prison. For example, if a drunk driver reparks the car in a parking lot, then this is not considered aggravated. Consequently, these individuals did not receive the treatment (length of imprisonment), and therefore, there should not be any treatment effect for this subpopulation. To test this, we conduct a refutability test based on these individuals. The prediction is that there should not be a break in the reduced-form relationship at BAC level 1.0‰ since there is no first-stage relationship for this subpopulation.

In our data, there are 2,458 cases where the individuals have been fined despite that they have had a BAC level above 1.0‰. Figure 8 shows the reduced-form relationship between the number of future crimes and the BAC level using the binned scatter approach by Cattaneo et al. (2024). We find no break in the relationship at 1.0‰. Indeed, there is a linear relationship across all BAC values.

Figure 8. Refutability test based on exceptions to the law

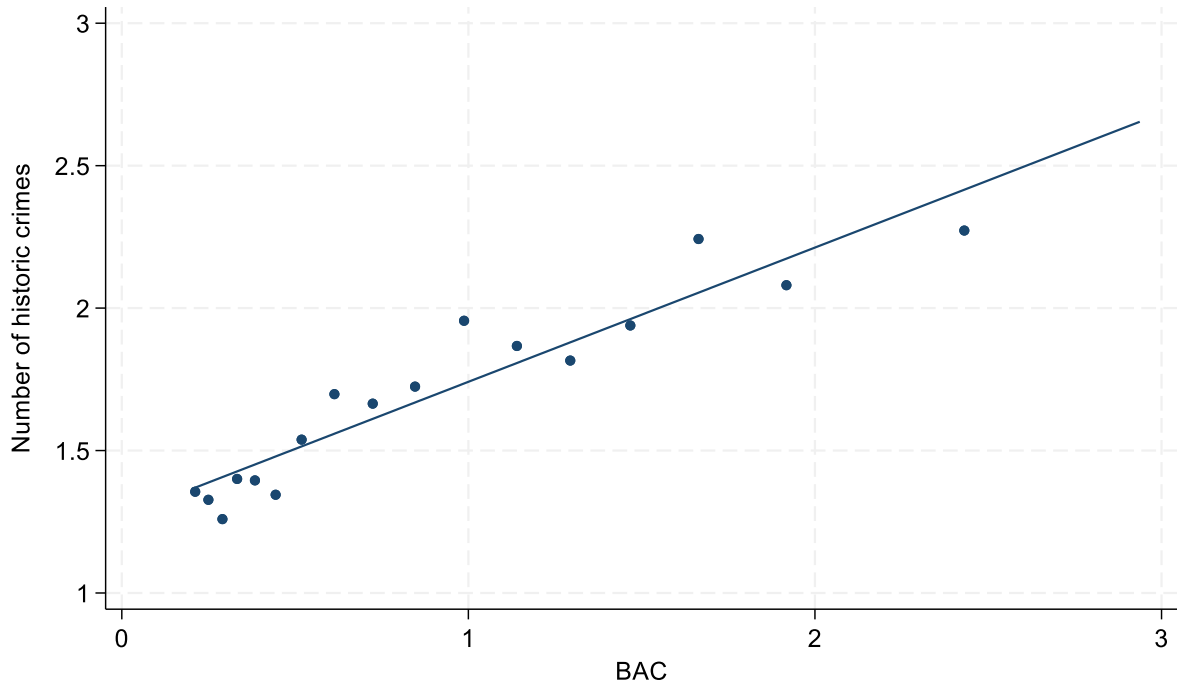


Notes: The binscatter plot was created by the stata command `binsreg` (Cattaneo et al. 2025) with 10 quantile-spaced bins. The number of observations above BAC level 1.0‰ is 2,458. It also adds a linear prediction line to the plot.

We can also perform a refutability test based on historical crimes since very few of the drunk drivers in our data have been previously sent to prison. This means that there should not be a break in the reduced-form relationship at BAC level 1.0‰ since there is essentially no

first-stage relationship for this population. Figure 9 shows the reduced-form relationship between the number of future crimes and the BAC level. Again, there is no break but instead a linear relationship between the number of historic crimes and BAC.

Figure 9. Refutability test based on historical crimes



Notes: The binscatter plot was created by the stata command binsreg (Cattaneo et al. 2025) with 17 quantile-spaced bins. It also adds a linear prediction line to the plot.

We can also use the data on historical crimes to test whether our instrument is as good as randomly assigned, i.e., whether there is any omitted variable bias (OVB). It is noteworthy that historical crime is a powerful predictor of future crimes, which is of importance when assessing OVB. Table 6 shows the IV results when controlling for the number of historical crimes, both parametrically (linear) and non-parametrically (full set of binary indicators). These two estimates should be compared to -0.0081 in Table 2. The conclusion is that controlling for the number of historic crimes has limited effects on the IV estimate, irrespective of whether one controls for the crime history, linearly or fully non-parametrically.

Another key predictor of crime is the individual's age at crime. However, it is important to keep in mind that age at crime may not be a valid control variable since it occurs simultaneously with the crime itself and therefore might be endogenous to crime. Nonetheless, we still test whether controlling for the age at crime has an impact on the IV estimate. Table 7 displays these results. The estimate is -0.0060 for both the linear and non-parametric

specification. Controlling for age thus has a larger impact than historic crimes on the original IV of -0.0081 in Table 2, but these two estimates are still not statistically significantly different from each other.

Table 6. Controlling for criminal history

	Linear (1)	Non-parametrically (2)
Number of days in prison	-0.0071 (0.0009)	-0.0068 (0.0008)
Number of observations	79,569	79,569

Notes: This table reports IV estimate when controlling for criminal history, i.e., the cumulative number of crimes committed before the sentence, linearly (Column 1) and nonparametrically (Column 2). The standard errors reported within parenthesis are clustered at the individual level.

Table 7. Controlling for age at crime

	Linear (1)	Non-parametrically (2)
Number of days in prison	-0.0060 (0.0009)	-0.0060 (0.0009)
Number of observations	79,569	79,569

Notes: This table reports IV estimates when controlling for age at crime linearly (Column 1) and nonparametrically (Column 2). The standard errors reported within parenthesis are clustered at the individual level.

Another issue that can lead to OVB bias is that we have misspecified the functional form of BAC, i.e., that we need to control for a higher-order polynomial in BAC for the instrument to be uncorrelated with the error term. Table 8 shows the results when we control for a second and third-order polynomial. The IV estimate is smaller than 0.0081 when controlling for a second-order polynomial in BAC, but also larger when controlling for a third-order polynomial. This implies that our results are robust to misspecification of the functional form of BAC.

Table 8. Controlling for different polynomials in BAC

	Second order	Third order
	(1)	(2)
Number of days in prison	-0.0057 (0.0012)	-0.0106 (0.0037)
Number of observations	79,569	79,569

Notes: This table reports IV estimate when controlling for a second-order polynomial in BAC (Column 1) and a third-order polynomial (Column 2). The standard errors reported within parentheses are clustered at the individual level.

Finally, we have also estimated a specification where we include all these control variables simultaneously, i.e., a third-order polynomial in BAC and controls for criminal history and age. This IV estimate is -0.0081 with a standard error of 0.0034, i.e., identical to the original IV estimate in Table 2.

6. Gender differences

In this section, we present the results on whether there are any gender differences since women only make up 15% of the total sample. Tables 9 and 10 show the first-stage, reduced-form and instrumental-variable estimates for men and women, respectively. The size of the first stage 21 is the same as for the total sample in both tables, while the reduced form effect is smaller for women, -0.18 versus -0.10. Thus, the IV estimate -0.0049 is also smaller for the sample of females than the IV estimate -0.0083 for the sample of men. However, the effect in percentage terms is of similar magnitudes both for men and women since females commit fewer crimes than men. In other words, men in the control group commit on average 0.33 crimes while the corresponding number for women is 0.17. Hence, the effects for both men and women correspond to a reduction of 80 % in recidivism if an individual is sentenced to prison for one month as compared to the control group that is sentenced to pay fees only. The similar effects for men and women bolster our assumption that the average treatment effect is approximately constant.

Table 9. First-stage, reduced form and instrumental variable estimates: males

Dependent variables	Number of days in prison	Number of crimes committed within 3 years after conviction	Number of crimes committed within 3 years after conviction
	First-stage (1)	Reduced form (2)	IV (3)
<i>Independent variables</i>			
Instrument: $D \times BAC$	21 (0.7)	-0.18 (0.02)	
Number of days in prison			-0.0083 (0.0016)
Number of observations	67,795	67,795	67,795

Notes: This table reports the first-stage estimate (Column 1), the reduced-form estimate (Column 2) and the corresponding IV estimate (Column 3). The first-stage cluster robust F-statistics is 4,447. The standard errors reported within parenthesis are clustered at the individual level.

Table 10. First-stage, reduced form and instrumental variable estimates: females

Dependent variables	Number of days in prison	Number of crimes committed within 3 years after conviction	Number of crimes committed within 3 years after conviction
	First-stage (1)	Reduced form (2)	IV (3)
<i>Independent variables</i>			
Instrument: $D \times BAC$	21 (0.7)	-0.10 (0.03)	
Number of days in prison			-0.0049 (0.0016)
Number of observations	11,774	11,774	11,774

Notes: This table reports the first-stage estimate (Column 1), the reduced-form estimate (Column 2) and the corresponding IV estimate (Column 3). The first-stage cluster robust F-statistics is 926. The standard errors reported within parentheses are clustered at the individual level.

7. Prison vs. electronic monitoring

In this section, we analyze whether the magnitudes of the deterrence effects differ if an individual has served the sentence in prison or in the home environment being surveilled at all times with electronic monitoring (EM).

This comparison is made possible since the severity of punishment is determined by the Swedish courts, as discussed above, but the Swedish Prison and Probation Service, may, under

some criteria (e.g., acceptable accommodation, regular employment, drug and alcohol testing, household consent, and permission to home visits by probation officers), allow the individual to serve the sentence in the home environment rather than in prison.

We have collected data from the Swedish Prison and Probation Service and matched it with the data from the Swedish National Council for Crime Prevention. We find that out of 11,552 observations where an individual was sentenced to prison, 2,845 (26%) served it in prison, while 7,965 (74%) were monitored in their homes electronically.¹²

Table 11 shows the first-stage, reduced-form and instrumental-variable estimates for three years after conviction for the subsample where they served their sentence in prison, while Table 12 displays the corresponding results for those with electronic monitoring. Although the first-stage and reduced-form estimates differ between the two subsamples due to the differences in the probability of placement, the magnitude of the IV estimates is quite similar, i.e., -0.0113 vs. -0.0097. Thus, the deterrence effect of being sentenced to prison is not affected by its execution, i.e., prison vs electronic monitoring.

A caveat is in place regarding this conclusion since the choice of treatment, i.e., prison vs. EM, is partly endogenous. Nonetheless, “one expects that those least likely to recidivate are assigned to EM as opposed to prison,” as argued by Hjalmarsson (2024, p. 47). Thus, this reasoning suggests that the deterrence effect of EM on reoffending that we find may be biased upwards, while the deterrence effect of prison on recidivism can be biased downwards. As a result, the true deterrence effect of prison cannot be smaller than the true effect of EM on recidivism. Nonetheless, finding similar results suggests that it is the severity of punishment decided by the Swedish courts rather than the specific facility where the sentence is served that deters individuals from committing crimes.

¹² We fail to merge 742 observations.

Table 11. First-stage, reduced form and instrumental variable estimates for the prison sample

Dependent variables	Number of days in prison	Number of crimes committed within 3 years after conviction	Number of crimes committed within 3 years after conviction
	First-stage (1)	Reduced form (2)	IV (3)
<i>Independent variables</i>			
Instrument: $D \times BAC$	11 (0.33)	-0.13 (0.02)	
Number of days in prison			-0.0113 (0.0020)
Number of observations	70,857	70,857	70,857

Notes: This table reports the first-stage estimate (Column 1), the reduced-form estimate (Column 2) and the corresponding IV estimate (Column 3). The standard errors reported within parenthesis are clustered at the individual level.

Table 12. First-stage, reduced form and instrumental variable estimates for the EM sample

Dependent variables	Number of days in prison	Number of crimes committed within 3 years after conviction	Number of crimes committed within 3 years after conviction
	First-stage (1)	Reduced form (2)	IV (3)
<i>Independent variables</i>			
Instrument: $D \times BAC$	19 (0.32)	-0.18 (0.02)	
Number of days in prison			-0.0097 (0.0010)
Number of observations	75,966	75,966	75,966

Notes: This table reports the first-stage estimate (Column 1), the reduced-form estimate (Column 2) and the corresponding IV estimate (Column 3). The standard errors reported within parentheses are clustered at the individual level.

8. Concluding remarks

Traditional economics theory suggests that longer prison sentences should deter recidivism. However, the previous empirical literature on the length of imprisonment and crime does not give clear support to this theory. We study a prison population that, on average, commits about the same number of crimes as the population at large, the drunk drivers. We exploit a unique feature in the sentencing for drunk driving applied in the Swedish court system. Individuals are

not sentenced to prison if caught driving with a BAC level below 1.0. Above 1.0, the number of expected days in prison is increasing almost linearly with the BAC level. We find that the piecewise linear relationship appears in the reduced form between BAC level and the number of reoffenses, and the evidence is consistent only with the interpretation that prison sentences drive this effect. In other words, in line with the deterrence theory that not only the certainty of punishment but also its severity should deter crime, our results show very large and robust crime-reducing effects of short prison sentences.

We believe an overall takeaway is that, in order to understand the behavioral mechanisms to what extent the severity of punishment deters crimes, it is important to focus not only on prolific offenders but also on populations that are more similar to the population at large.

References

Aizer, Anna and Doyle, John (2015), “Juvenile Incarceration, Human Capital and Future Crime: Evidence from Randomly Assigned Judges”, *Quarterly Journal of Economics*, 130: 759-853.

Alsan, Marcella, Barnett, Arkley, Hull, Peter and Yang, Crystal S. (2025), ““Something Works” in U.S. Jails: Misconduct and Recidivism Effects of the IGNITE Program”, *The Quarterly Journal of Economics*, 140: 1367-1415.

Andresen, Martin E. and Huber, Martin (2021) “Instrument-based estimation with binarised treatments: issues and tests for the exclusion restriction”, *The Econometrics Journal*, 24: 536–558.

Angrist, Joshua D., and Imbens, Guido W. (1995), “Two-stage least squares estimation of average causal effects in models with variable treatment intensity”, *Journal of the American statistical Association*, 90, 431-442.

Angrist, Joshua D. and Krueger, Alan B. (1999), “Empirical strategies in labor economics”, in *Handbook of labor economics*, 3: 1277-1366, Elsevier.

Bhuller, Manudeep, Dahl, Gordon B., Löken, Katrine V. and Mogstad, Magne (2020), “Incarceration, Recidivism, and Employment”, *Journal of Political Economy*, 128: 1269-1324.

Card, David, Lee, David S., Pei, Zhuan, and Weber, Andrea (2015), “Inference on causal effects in a generalized regression kink design”, *Econometrica*, 83: 2453-2483.

Carpenter, Cristopher and Dobkin, Carlos (2010), “Alcohol Regulation and Crime”, in *Controlling crime: Strategies and tradeoffs* (pp. 291-329). University of Chicago Press.

Cattaneo, Matias D., Crump, Richard K., Farrell, Max H. and Feng, Yingjie (2024), “On Binscatter”, *American Economic Review*, 114: 1488 – 1514.

Cattaneo, Matias D., Crump, Richard K., Farrell, Max H and Feng, Yingjie (2025), “Binscatter Regressions”, *Stata Journal*, 25: 3-50.

Chalfin, Aaron and McCrary, Justin (2017), “Criminal Deterrence: A Review of the Literature”, *Journal of Economics Literature*, 55: 5-48.

Di Tella, Rafael and Schargrodsky, Ernesto (2013), “Criminal Recidivism after Prison and Electronic Monitoring,” *Journal of Political Economy*, 121: 28–73.

Durlauf, Steven N. and Nagin, Daniel S. (2011), “Crime and Imprisonment”, *Criminology & Public Policy*, 10: 13-54.

Federal Bureau of Prisons (2025),
https://www.bop.gov/about/statistics/statistics_inmate_sec_levels.jsp, accessed September 2025.

Green, Donald P. and Winik, Daniel (2010), ”Using Random Judge Assignments to Estimate the Effects of Incarceration and Probation on Recidivism Among Drug Offenders”, *Criminology*, 48: 357-387.

Grenet, Julien, Grönqvist, Hans and Niknami, Susan (2024), “The Effects of Electronic Monitoring on Offenders and their Families”, *Journal of Public Economics* 230, Article 105051.

Hansen, Benjamin (2015), “Punishment and Deterrence: Evidence from Drunk Driving”, *American Economic Review*, 105: 1581-1617.

Hansen, Bruce E. (2017), “Regression kink with an unknown threshold”, *Journal of Business & Economic Statistics*, 35, 228-240.

Hansen, Bruce E. (2022). *Econometrics*, Princeton University Press.

Henneguelle, Anais, Benjamin, Monnery and Annie, Kensey (2016), “Better at Home than in Prison? The Effects of Electronic Monitoring on Recidivism in France,” *Journal of Law and Economics*, 59: 629–667.

Hjalmarsson, Randi (2009), ”Crime and Expected Punishment: Changes in Perceptions at the Age of Criminal Majority”, *American Law and Economics Review*, 11: 209-248.

Hjalmarsson, Randi (2024) “Determinants of Recidivism: How Criminal Justice Interactions and the Post-Release Environment Affect Repeat Offending” SNS Research Report.

Kolesár, Michal (2013), “Estimation in an Instrumental Variables Model with Treatment Effect Heterogeneity”, Mimeo, Princeton University.

Kolesár, Michal, Chetty, Raj, Friedman, John N., Glaeser, Edward, and Imbens, Guido W. (2015), “Identification and Inference with Many Invalid Instruments”, *Journal of Business and Economic Statistics*, 33: 474-484.

- Kuziemko, Ilyana (2013), "How Should Inmates be Released from Prison? An Assessment of Parole Versus Fixed Sentence Regimes", *The Quarterly Journal of Economics*, 178: 371-424.
- Lee, David S. and McCrary, Justin (2009), "The Deterrent Effect of Prison: Dynamic Theory and Evidence", Working paper, No. 550, Industrial Relations Section, Princeton University.
- Loeffler, Charles E. and Nagin, Daniel S. (2022), "The Impact of Incarceration on Recidivism", *Annual Review of Criminology*, 5: 133-152.
- Marshall, John (2016), "Coarsening Bias: How Coarse Treatment Measurement Upwardly Biases Instrumental Variable Estimates", *Political Analysis*, 24: 157-17.
- Mueller-Smith, Michael (2015), "The Criminal and Labor Market Impacts of Incarceration", Mimeo, Columbia University.
- National Highway Traffic Safety Administration (NHTSA) (2016), <https://www.nhtsa.gov/sites/nhtsa.gov/files/809844-theabcsofbac.pdf>, accessed June 2025.
- Paternoster, Raymond and Piquero, Alex R. (1995) "Reconceptualizing Deterrence: An Empirical Test of Personal and Vicarious Experiences", *Journal of Research in Crime and Delinquency*, 32: 251-286.
- Raphael, Steven and Ludwig, Jens (2003), "Prison Sentence Enhancements: The Case of Project Exile." in Jens Ludwig and Philip J. Cook. (eds.), *Evaluating Gun Policy: Effects on Crime and Violence*, Washington, DC: Brookings Institution Press.
- Rose, Evan K. and Shem-Tov, Yotam (2021), "How Does Incarceration Affect Reoffending? Estimating the Dose-Response Function", *Journal of Political Economy*, 129: 3302-3356
- Rose, Evan K. and Shem-Tov, Yotam (2024), "On Recoding Ordered Treatments as Binary Indicators", forthcoming in *Review of Economic and Statistics*.
- Stafford, Mark and Warr, Mark (1993), "A Reconceptualization of General and Specific Deterrence", *Journal of Research in Crime and Delinquency*, 30: 123-135.
- Starr, Evan and Goldfarb, Brent D. (2020), "Binned Scatterplots: A Simple Tool to Make Research Easier and Better", *Strategic Management Journal*, 41: 2261-2274.
- Tomlinson, Kelli D. (2016), "An Examination of Deterrence Theory: Where Do We Stand?" *Federal Probation*, 80: 33-38.
- Vollaard, Ben (2013), "Preventing Crime through Selective Incapacitation", *Economic Journal*, 123: 262-284.
- Williams, Jenny and Weatherburn, Don (2020), "Can Electronic Monitoring Reduce Reoffending?," *Review of Economics and Statistics*, 104: 232-245.