—Preliminary Technical Report—

Quantifying the Effectiveness of the Winnebago County Misdemeanor Drug Diversion Program

Chad D. Cotti*, PhD, MPA
M. Ryan Haley†, PhD

Department of Economics
University of Wisconsin Oshkosh
800 Algoma Blvd
Oshkosh, WI 54901

19 October 2011

Abstract

We use a Cox proportional hazard model and propensity-score matching to estimate how much the Winnebago County Misdemeanor Drug Diversion Program (MDDP) impacts the likelihood of criminal re-offense. We obtain this estimate by analyzing data on misdemeanor drug offenders arrested in Winnebago County during 2010 and 2011. Within this time frame, we find that the hazard rate of re-offense is 75% lower (per day) than those not treated with the MDDP program. After adjusting for potential selectivity, we find that completing the MDDP correlates with an 18% reduction in the likelihood of criminal re-offense over the two years studied. Overall, this suggests that completing the MDDP meaningfully reduces recidivism.

*cottic@uwosh.edu
†haley@uwosh.edu

1 Introduction and Motivation

The Winnebago County District Attorney's (DA) office handles hundreds of misdemeanor drug offenders each year. Under ideal circumstances the DA would have the resources to prosecute and incarcerate these offenders in a manner befitting their crime. However, prosecution and incarceration are costly, jails/prisons are overcrowded, the DA is understaffed, and funding is currently tight. Together these issues create an environment where a) offenders who might otherwise do jail time are instead fined and released; b) some offenses that would normally be aggressively prosecuted are prosecuted less aggressively or perhaps not at all; and c) potential offenders get the impression that the DA is "soft on crime." Obviously, these are not optimal from a jurisprudential or societal-safety perspective.

To circumvent these shortcomings, Winnebago County DA Christian Gossett and the Winnebago County Safe Streets Initiative have installed a series of diversion programs designed to educate first-time or minor offenders. Among these programs is the Misdemeanor Drug Diversion Program (MDDP), which offers guidance, education, drug testing, and deferred adjudication instead of traditional prosecution. The program has two main goals. The first is to obtain a recidivism rate lower than what would have obtained had the offender been prosecuted, and the second is to save money by obviating prosecution and incarceration.

The primary advantage of the MDDP (for offenders) is the possibility of deferred adjudication. Specifically, if the offender successfully completes the MDDP (which may last 6-18 months, often involving a series of random urinalysis tests), the charges are essentially vacated, permitting the offender to correctly state that they were not convicted of the crime for which they were arrested. The program costs $250, which is presumably a fraction of what they would have otherwise been fined. This type of program may be particularly attractive to misdemeanor and first-time offenders that were arrested pursuant to non-habitual drug experimentation.

2 Data

From January 2010 through August of 2011 there were 250 misdemeanor drug offenders who were eligible to receive a suspended sentence and enter the MDDP as an alternative. The DA's office provided us with data on all 250 offenders offered the program. The data set includes the age, sex, and race of the offender. It also includes information on the misdemeanor drug charge in question, if this was the offender's first offense, how long after the original offense the offender is observed (i.e., measuring opportunity to commit another crime), if an offender re-offended, and when the re-offense(s) occurred. Lastly, and most importantly, the data state whether or not the offender chose the MDDP option and if they completed the program. Descriptive statistics on these 250 offenders appear in table 1.

The available data does have shortcomings. For example, it does not provide information about the offender's education, employment status, martial status, income/wealth, among several other key psychosocial and demographic variables. Additionally, we only observe offenders for these two years; ideally, the data would span a longer time frame. Nevertheless, the data provides us with sufficient information to generally evaluate the effectiveness of the MDDP.

3 Estimation

Our primary estimation strategy is a Cox proportional hazard model, which, in our case, relates the time that passes before a criminal re-offense to one or more characteristics about an individual that may also be associated with a re-offense. Therefore, the proportional hazard approach allows us to determine if a misdemeanor drug offender's hazard rate of criminal re-offense is impacted by their completion of the MDDP.

Our proportional hazard model is defined as follows:

$$H(t|X_i) = H_0(t) \times \exp(\beta MDDP_i + \gamma^\top X_i)$$

Table 1: Descriptive Statistics[†]

Variable	Mean or Proportion	Standard Deviation
Finished MDDP	0.132	0.339
re-offended (yes or no)	0.152	0.360
Number of re-offenses (entire sample)	0.224	0.600
Number of re-offenses (if re-offended)	1.368	0.6334
Age	27.28	10.67
Sex (Female=1)	0.188	0.391
White	0.876	0.330
Black	0.088	0.284
Hispanic	0.020	0.140
Other Race	0.016	0.126
First Offense	0.688	0.464
Drug Charge: THC	0.788	0.410
Drug Charge: Rx	0.180	0.384
Drug Charge: Cocaine	0.032	0.176

[†]Sample Size = 250.

where $H(t|X)$ is the hazard of re-offense at time t, $MDDP_i$ indicates completion of the program by offender i, X_i represents a vector of predictor variables, including dummy variables for sex, race, drug offense, and first offense, as well as the offender's age, and $H_0(t)$ is the baseline hazard at time t, which represents the hazard for a individual with the value 0 for all the predictor variables. The baseline hazard corrects for baseline differences between offenders that might influence the likelihood of re-offense over time.[1] The Cox proportional hazard regression model produces an adjusted hazard ratio that accounts for baseline differences between offenders that may influence the choice to re-offend that are not already captured in the predictor variables. Within this hazard analysis framework, the goal is to capture the factors that lead to re-offense as best as possible.

3.1 Proportional Hazard Results

Table 2 reports the hazard ratios for re-offense from the Cox proportional hazard analysis. The statistical significance of the hazard ratios (and their underlying parameter estimates)

[1]This issue is less of a concern if randomization has rendered both groups similar in terms of their baseline characteristics, but, obviously, this is not the case in studies of public policy adoption.

are calculated using robust standard errors that have been clustered at the state-level to accommodate non-independence of observations from the same state over time (Bertrand, Duflo, and Mullainathan, 2004). A hazard ratio greater than one would indicate that re-offense was higher after completing MDDP, while a ratio below one indicates the opposite.

Our estimates indicate that, after controlling for age, race, sex, type of original drug offense, and if the individual was a first time offender, completion of the MDDP has a negative and statistically significant effect on the likelihood of re-offense. Specifically, results suggest that individuals who completed the MDDP were about 75% less likely to criminally re-offend on any given day than those who did not complete the program. We emphasize that this does *not* mean they are 75% less likely to re-offend in general. *More precisely, this means that a MDDP completer who has not re-offended by a certain time has a 75% lower chance of offending in the next day compared to someone who did not complete the MDDP.*

While not statistically significant, the results also suggest that first-time offenders and women are less likely to re-offend, controlling for the aforementioned characteristics. We suspect these effects would be statistically significant were the sample size larger, hence permitting more precise standard errors.

Table 2: Proportional Hazard Model Estimates[†]

Variable	Hazard Ratio Estimates	p-value
Complete MDDP	0.2536*	[0.050]
Age	0.9871	[0.430])
Sex	0.5467	[0.228]
White	@	@
Black	1.1986	[0.728]
Hispanic	1.0836	[0.938]
Other Race	0.0001*	[0.000]
First Offense	0.5861	[0.126]
Drug Charge: THC	@	@
Drug Charge: Rx	1.4352	[0.384]
Drug Charge: Cocaine	0.8131	[0.801]

[†]Sample Size = 266 individuals (due to multiple re-offenders counting multiple times). The "@" sign indicates a reference group; e.g., all race estimates are effects relative to the reference group White. Robust standard errors were used to calculate all p-values. A "*" indicates statistical significance.

It is often difficult for non-statisticians to get an intuitive understanding of hazard rates and hazard ratios. Recognizing this, in the appendix we present a somewhat simpler and more illustrative example, which will hopefully help make these results more accessible.

3.2 Propensity Score Matching

To estimate the effects of a treatment program we would, ideally, have a randomized controlled experiment wherein the treatment group (those who completed the MDDP) and control group (those who qualified but did not choose or complete the MDDP) differ in expectation only by the treatment effect. However, we cannot control the treatment allocation process in observational data. Indeed, subjects assign themselves to treatment and control, resulting in selection bias. While estimation of a baseline hazard in a proportional hazard model helps mitigate these concerns by including covariates, we can utilize propensity score matching (PSM) to verify the estimates of the proportional hazard model. We caution, however, that these estimation strategies are gauging somewhat different quantities, thus their results require different interpretations.

We conduct PSM by first estimating a propensity score for each offender based on their likelihood of completing the MDDP treatment. We obtain these scores by estimating a logistic regression wherein the completion of the MDDP serves as the dependent variable and individual offender characteristics or characteristics about the initial crime committed are used as predictor variables. Then offenders are divided into two groups: those that actually completed the MDDP treatment and those that did not. The treated and untreated groups are matched based on their propensity scores by using the standard "nearest neighbor within caliper" approach, which matches offenders that have propensity scores within a certain caliper range. We chose a caliper of 5%, which is relatively standard in PSM (Oakes and Johnson, 2006).

Having completed the matching, we can now perform the desired PSM estimation. The average treatment difference of MDDP completion on re-offense between the treatment and control groups equals -0.1818 (bootstrap standard error = 0.1048). While not expressible

as a hazard ratio, the matching results provide us with the difference in the proportion of initial offenders that re-offend at any point in the sample, adjusting for selection bias and opportunity. Specifically, the data suggests that there is about an 18% lower chance of recidivism for offenders that completed the MDDP during the two years we observed. These results are consistent with the hazard estimates in that they both find a statistically significant reduction in recidivism for offenders that completed the MDDP.

4 Discussion and Conclusions

We use data on 250 misdemeanor drug offenders from the 2010-2011 time frame to estimate the effectiveness of the MDDP. Our estimates gauge the rate of recidivism for those offenders that completed the program relative to their counterparts that chose not to complete the MDDP. Our specific findings suggest that the hazard ratio for the MDDP is about 0.25, indicating that recidivism is far lower (about a 75% lower chance of re-offense per day) among those that completed the MDDP. These results translate into an 18% reduction in the likelihood of criminal re-offense over the complete time frame studied, which suggests that completion of the MDDP reduces both the rate and likelihood of criminal behavior.[2]

Of course, the quality and accuracy of these results are limited by the level of detail provided in the data. Specifically, access to information on the offender's education, employment status, martial status, and income/wealth, provided on a greater number of individuals, over a longer period of time, would greatly enhance the accuracy and precision of the estimates.

Nevertheless, these preliminary results are testimony to the effectiveness of the MDDP, suggesting that the Winnebego County DA and the Safe Streets Initiative have created a program that appears to curtail drug-related offenses. Lower recidivism means fewer crimes in the future (all else held constant) and accordingly reduces the demand for current and future DA resources—a partial solution to tight budgets and jail/prison overcrowding.

[2]As further clarification, that hazard ratio of 0.25 simply implies that the relative risk of re-offense in the treatment group is 75% lower *per day*, not across the entire time frame. This difference in the relative risk of re-offense roughly translates into an 18% reduction over the time frame studied. Informally, this 18% figure approximates the difference between the 6% recidivism rate in the treatment group and the 24% recidivism rate in the control group.

Appendix: Hazard Redux

In this appendix we illustrate hazard rates and ratios using a simple medical example.

Imagine a medical trial with 1000 patients, 500 of which receive a new medication thought to reduce heart attacks and 500 of which do not receive the treatment; the former is termed the "treatment group" and the latter is the "control group." Assuming that the subjects are randomly assigned to the treatment and control groups, these two "arms" of the study permit the researcher to assess the efficacy of the new medication. Assume the trial lasts 60 months.

Within this context, there will likely be heart attacks during the trial, some of which will occur in the treatment group and some of which will occur in the control group. If the medication is effective, we would expect to see far fewer heart attacks in the treatment group than in the control; likewise, if the medication is not effective, we would expect to see roughly the same number of heart attacks in each group. This brings us to the doorstep of hazard rates and the hazard ratio.

The hazard rate, almost universally denoted as $h(t)$, is the probability that a subject has a heart attack at time t. Time often means per day, week, or month, depending on the exact nature of the study. For our heart attack example, we use months as our unit of time. So, put slightly differently, the hazard rate indicates the monthly probability that a heart attack will occur, given that a subject has survived to month t.

Suppose that 8 treatment and 10 control subjects suffer heart attacks in the first month. The hazard rate for the treatment group, $h(t = 1 \text{ month})$, would be $8/500 = 0.016$, while the hazard rate for the control group would be $10/500 = 0.02$. If, in the second month, the treatment group suffers 6 more heart attacks and the control suffers 9 more, the respective hazard rates $h(t = 2 \text{ months})$ would be $6/492 = 0.012$ and $9/490 = 0.018$. Upon completing the study, we would have 60 hazard rates for the treatment group and 60 hazard rates for the control group. Researchers often plot these points and connect them with a step function, giving a discrete approximation to the hazard function over the entire time line of the study.

Often researchers are not interested in the hazard functions themselves, but rather by how much the two functions differ. Indeed, the difference in the functions captures the effect of the treatment (heart attack medication) itself, which is usually the effect of primary interest in a randomized clinical trial. One way to gauge the difference is by creating a ratio—a hazard ratio—of the treatment hazard rate at a given time divided by the control hazard rate at that same time.

If the treatment and control groups showed no differences in heart attack frequency during the trial, the hazard ratio would equal one; i.e., the two hazard rates would be equal to each other, indicating that the medication appears ineffective in reducing heart attacks. Alternatively, if the hazard ratio were two, this would mean that a treated patient who has not had a heart attack by a certain time in the study has twice the chance of having a heart attack in the next month compared to someone in the control group. Such a finding would suggest that the medication was most definitely not working, and perhaps even damaging the heart.

This comparison of hazard rates, encapsulated in the single hazard ratio, is often taken one step further by formally testing the hazard ratio using hypothesis testing methods (e.g., the log-rank test, which assumes, under the null, that the hazard ratio is one). The purpose of this is to determine if the hazard ratio estimates are credible, or of they are just the result of sampling error.

With this heart example in mind, it is perhaps now easier to interpret the 0.25 result obtained from our Cox proportional hazard estimation.[3] As before, *this means that a MDDP completer who has not re-offended by a certain time has a 75% lower chance of offending in the next day compared to someone who did not complete the MDDP.*

[3]The "Cox proportional" part means that we are assuming that the treatment and control hazard rates occur in the same proportion throughout the time line; this is a common way to facilitate the interpretation of the results.

References

[1] Bertrand, M., Duflo, E., and S. Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*, 119: 249-75.

[2] Oakes, J. and P. Johnson. 2006. "Chapter fifteen: Propensity score matching for social epidemiology" in *Methods in Social Epidemiology* (editors J. Oakes and J. Kaufman), Jossey-Bass.