

Essays in Applied Economics

by

Benjamin Crost

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division
of the
University of California, Berkeley

Committee in charge:
Professor Alain de Janvry, Co-chair
Professor Elisabeth Sadoulet, Co-chair
Professor Sofia Berto Villas-Boas
Professor Edward Miguel

Spring 2011

Essays in Applied Economics

Copyright 2011
by
Benjamin Crost

Abstract

Essays in Applied Economics

by

Benjamin Crost

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Alain de Janvry, Co-chair

Professor Elisabeth Sadoulet, Co-chair

This dissertation combines research on three topics in applied empirical economics. The first paper, which is based on joint work with Patrick Johnston, examines the effect of development projects on civil conflict. The second paper estimates the effect of subsidized employment on the happiness of the unemployed. The third paper, based on joint work with Santiago Guerrero, analyzes the effect of restrictions to alcohol accessibility on Marijuana use.

The first paper develops a theoretical model of bargaining and conflict in the context of development projects. The model predicts that development projects cause an increase in violent conflict if governments cannot (1) ensure the projects success in the face of insurgent opposition and (2) credibly commit to honoring agreements reached before the start of the project. The model is tested by estimating the causal effect of a large development program on conflict casualties in the Philippines. Identification is based on a regression discontinuity design that exploits an arbitrary poverty threshold used to assign eligibility for the program. Consistent with the models predictions, we find that eligible municipalities suffered a substantial increase in casualties, which lasts only for the duration of the project and is split evenly between government troops and insurgents.

The second paper estimates the causal effect of a type of subsidized employment projects - Germany's *Arbeitsbeschaffungsmassnahmen* - on self-reported happiness. Results from matching and fixed effects estimators suggest that subsidized employment has a large and statistically significant positive effect on the happiness of individuals who would otherwise have been unemployed. Detailed panel data on pre- and post-project happiness suggests that this effect can neither be explained by self-selection of happier individuals into employment nor by the higher incomes of the employed. This suggests that subsidized employment programs are more effective at increasing the happiness of the unemployed than an increase in unemployment benefits.

The third paper estimates the effect of the Minimum Legal Drinking Age of 21 years on Marijuana use. The casual effect of this law is estimated through a regression discontinuity design that compares Marijuana use among individuals just below and just above age 21. We

find a significant drop in Marijuana use at age 21, which suggests that individuals substitute between alcohol and Marijuana. Policies that restrict alcohol accessibility are therefore likely to have the unintended consequence of increasing Marijuana use.

To my parents

Contents

List of Figures	iv
List of Tables	v
Acknowledgments	vi
1 Overview	1
2 Aid Under Fire: Development Projects and Civil Conflict	3
2.1 Introduction	3
2.2 Development Projects, Bargaining and Conflict: A Simple Model	6
2.2.1 Comparison with other models of conflict:	10
2.3 Empirical Setting: Conflict and Development Projects in the Philippines . .	11
2.3.1 Violent conflict in the Philippines	11
2.3.2 The KALAHY-CIDSS Program	13
2.4 Empirical Strategy	16
2.4.1 Regression discontinuity design	16
2.4.2 Data	18
2.4.3 Variables	19
2.5 Results	21
2.5.1 Summary statistics and balance tests	22
2.5.2 The effect of KALAHY-CIDSS on conflict casualties	24
2.5.3 Robustness tests	28
2.5.4 The effect of project size on conflict casualties	31
2.5.5 Who suffers and who initiates the violence?	32
2.6 Conclusion	34
3 The Effect of Subsidized Employment on Happiness	36
3.1 Introduction	36
3.2 Institutional Background: Subsidized Employment Projects in Germany . . .	39
3.3 Empirical Strategy	40
3.3.1 A simple model of happiness and (un)employment	40

3.3.2	A Matching Estimator for the Effect of Subsidized Employment on Happiness	41
3.3.3	Disentangling the effects of income and employment	43
3.4	Data and Summary Statistics	45
3.5	Results	47
3.5.1	The effect of subsidized employment on happiness	47
3.5.2	Does the conditional independence assumption hold?	48
3.5.3	Disentangling the effects of employment and income	51
3.5.4	Can changes in consumption or expected future income explain the results?	60
3.5.5	Can misreporting of happiness explain the results?	60
3.6	Conclusion	61
4	The Effect of Alcohol Availability on Marijuana Use: Evidence from the Minimum Legal Drinking Age	63
4.1	Introduction	63
4.2	Literature Review	65
4.3	Empirical Strategy	66
4.4	Data and Results	67
4.5	Robustness Tests	70
4.5.1	Placebo Tests for Location of the Discontinuity	70
4.5.2	Robustness Tests for Choice of Bandwidth	70
4.5.3	Results by Gender	74
4.6	Conclusions	74
	Bibliography	76

List of Figures

2.1	Timeline of KALAH-CIDSS Implementation	15
2.2	Probability of KALAH-CIDSS Participation by Distance from Poverty Threshold	17
2.3	Time Trend of Conflict in Treatment and Control Municipalities	25
3.1	Prevalence of Subsidized Employment Projects	39
3.2	Percentage of Unemployed in Subsidized Employment Projects	46
3.3	Trends of Happiness around the Start of Subsidized Employment projects . . .	49
3.4	Happiness, unemployment and income before and after the start of SEPs - all participants	52
4.1	Alcohol and Marijuana Use Around Age 21	68
4.2	Placebo Tests: RD Estimates by Location of RD Threshold	70
4.3	Robustness of Results to Choice of Bandwidth	71
4.4	Alcohol and Marijuana Use Around Age 21: Men	72
4.5	Alcohol and Marijuana Use Around Age 21: Women	73

List of Tables

2.1	Variables Used to Determine KALAHY-CIDSS Eligibility	14
2.2	Timetable of KALAHY-CIDSS	15
2.3	Summary Statistics of Conflict Outcomes, 2001-08	23
2.4	Balance of Observed Variables Across Eligibility Threshold	24
2.5	The Effect of KALAHY-CIDSS on Conflict Casualties	27
2.6	Robustness Tests: Choice of Bandwidth	29
2.7	Robustness Tests: Robustness to Outliers	30
2.8	Effect of Project Size on Casualties	31
2.9	Who Suffers and Who Initiates the Violence? Conflict Casualties by Actor .	33
3.1	Summary Statistics	46
3.2	Effect of Subsidized Employment Projects on Happiness: Matching Estimators	47
3.3	Employment vs. Income: Fixed Effects Estimates	53
3.4	Instrumental Variables Estimates: First Stage	54
3.5	Instrumental Variables Estimates: Second Stage	55
3.6	Robustness tests: changes in the effect of employment over time	58
3.7	Robustness tests: adaptation to unemployment	59
4.1	Effect of the MLDA on Alcohol and Marijuana Use: Regression Discontinuity estimates	69
4.2	Effect of the MLDA on Alcohol and Marijuana Use: Men	72
4.3	Effect of the MLDA on Alcohol and Marijuana Use: Women	73

Acknowledgments

I am indebted to many people who contributed to this dissertation, either directly or indirectly. First of all, I would like to thank Betty Sadoulet and Alain de Janvry for being great advisers. They have held me to a very high standard and my work has become much better because of it. I also want to thank the other members of my dissertation committee, Sofia Villas-Boas and Ted Miguel for their advice and encouragement.

I was lucky to have two very kind and productive co-authors, Patrick Johnston and Santiago Guerrero, who have made invaluable contributions to the first and third paper in this dissertation. I have also benefited greatly from conversations and seminar discussions with many people at UC Berkeley, both faculty and students. Among them are Christian Traeger, Jeremy Magruder, Ethan Ligon, Michael Anderson, Brian Wright, Max Aufhammer, Matt Rabin, Erick Gong, Leslie Martin, Steve Buck, Gianmarco Leon, Andrew Dustan, Anna Spurlock and Charles Seguin. I could go on.

Thanks also to my advisers and teachers at the University of Reading, most notably Bhavani Shankar, Garth Holloway and Chittur Srinivasan. They were the ones who sparked my interest in economics and without their encouragement and support I would never have started working on this dissertation.

Finally, thanks to my family and friends, both in Berkeley and abroad, for their support and for making the past five years memorable and enjoyable. You guys are the best!

Chapter 1

Overview

This dissertation combines research on three topics in applied empirical economics. Each of the three papers uses econometric methods to attempt to answer a question that is relevant for public policy.

The first paper, which is based on joint work with Patrick Johnston, examines the effect of development projects on civil conflict. In the last decade donors and governments have targeted an increasing amount of development aid to conflict-affected areas, some of it in the hope that aid will reduce conflict by weakening popular support for insurgent movements. This paper offers a simple but frequently overlooked explanation for why the opposite can be the case: if insurgents know that development projects will weaken their position, they have an incentive to oppose them, which may exacerbate conflict. To formalize this intuition, we draw on previous work by Powell (2004, 2006), to develop a theoretical model of bargaining and conflict in the context of development projects. The model predicts that development projects cause an increase in violent conflict if governments cannot (1) ensure the projects success in the face of insurgent opposition and (2) credibly commit to honoring agreements reached before the start of the project. To test the model, we estimate the causal effect of a large development program on conflict casualties in the Philippines. Identification is based on a regression discontinuity design that exploits an arbitrary poverty threshold used to assign eligibility for the program. Consistent with the models predictions, we find that eligible municipalities suffered a substantial increase in casualties, which lasts only for the duration of the project and is split evenly between government troops and insurgents.

The second paper estimates the effect of subsidized employment on happiness. A large body of previous research shows that the unemployed report significantly lower levels of average happiness. The negative correlation between unemployment and happiness, both across individuals and over time, remains significant after controlling for a wide range of observable characteristics, including income (Clark and Oswald 1994, Winkelmann and Winkelmann 1998, Marks and Fleming 1999, Clark 2003, Carroll 2007). But it is not clear whether this reflects a causal effect of unemployment on happiness, or whether unhappy individuals are merely less likely to find jobs. In addition it is unclear whether the kind of jobs that are created by employment subsidies, which are often menial and have low pay, can increase the

happiness of the unemployed. To close this gap in the literature, I estimate the causal effect of a type of subsidized employment projects - Germany's *Arbeitsbeschaffungsmassnahmen* - on self-reported happiness. Results from matching and fixed effects estimators suggest that subsidized employment has a large and statistically significant positive effect on the happiness of individuals who would otherwise have been unemployed. Detailed data on pre- and post-project happiness from the German Socio-Economic Panel suggests that this effect can neither be explained by self-selection of happier individuals into employment nor by the higher incomes of the employed. These results suggest that subsidized employment programs are more effective at increasing the happiness of the unemployed than an increase in unemployment benefits.

The third paper, based on joint work with Santiago Guerrero, analyzes the effect of restrictions to alcohol accessibility on Marijuana use. Economic theory suggests that when the cost of consuming a good increases, people will consume more of its substitutes and less of its complements. In the case of alcohol, these substitutes and complements are likely to include other intoxicating substances. The Minimum Legal Drinking Age (MLDA), which restricts access to alcohol for those under 21, is therefore likely to affect the consumption of other drugs among that age group, as it sharply decreases the cost of consuming alcohol for individuals just over the MLDA. We estimate the causal effect of the MLDA on Marijuana use through a regression discontinuity design that compares Marijuana use among individuals just below and just above age 21. We find a significant drop in Marijuana use at age 21, which suggests that individuals substitute between alcohol and Marijuana. Policies that restrict alcohol accessibility are therefore likely to have the unintended consequence of increasing Marijuana use.

Overall, the three papers highlight the usefulness of econometric research in addressing questions of relevance for public policy.

Chapter 2

Aid Under Fire: Development Projects and Civil Conflict

with Patrick Johnston

Abstract

An increasing amount of development aid is targeted to areas affected by civil conflict; some of it in the hope that aid will reduce conflict by weakening popular support for insurgent movements. But if insurgents know that development projects will weaken their position, they have an incentive to oppose them, which may exacerbate conflict and derail projects. We formalize this intuition in a theoretical model of bargaining and conflict in the context of development projects. Our model predicts that development projects cause an increase in violent conflict if governments cannot (1) ensure the project's success in the face of insurgent opposition and (2) credibly commit to honoring agreements reached before the start of the project. To test the model, we estimate the causal effect of a large development program on conflict casualties in the Philippines. Identification is based on a regression discontinuity design that exploits an arbitrary poverty threshold used to assign eligibility for the program. Consistent with the model's predictions, we find that eligible municipalities suffered a substantial increase in casualties, which lasts only for the duration of the project and is split evenly between government troops and insurgents.

2.1 Introduction

Over the last six decades, civil conflict has led to the deaths of more than 16 million people and the destruction of immense amounts of physical and human capital (Fearon and Laitin 2003). It has been associated with the spread of pandemics (Elbe 2002; Murray et al. 2002; Ghobarah et al 2004); the degradation of the rule of law (Reno 1998; Collier et al 2004; Fearon 2004; Ross 2004 ; Angrist and Kugler 2008) and the forced displacement of hundreds

of thousands of people (Salehyan 2007; Ibanez and Velez 2008). Unfortunately, civil conflict is widespread. Since the end of World War II over half of all countries have suffered at least one incidence of civil conflict, (Blattman and Miguel 2010). The urgency of civil conflict has never been greater: The proportion of conflict-affected countries, which increased steadily from 1945 through the mid-1990s, is once again on the rise (Harbom & Wallensteen 2010).

In recent years, donors and governments have targeted an increasing amount of development aid to conflict-affected areas, some of it in the hope that aid will reduce conflict by “winning the hearts and minds” of the population. This idea is at the heart of the US Armed Forces’ current counterinsurgency strategy (U.S. Army/Marine Corps 2007, Gompert et al. 2009) and is supported by experimental evidence, which shows that development aid increases support for the government in the Afghan population Beath et al. (2010). If increased support for the government translates into an increased willingness to supply the government with intelligence about the insurgents, a successful project would make it harder for insurgents to launch successful attacks on the government Berman et al. (2008). Another way in which development aid may reduce conflict is by increasing individuals’ economic opportunities, making them less likely to join insurgent groups and participate in conflict (Collier and Hoeffler Collier & Hoeffler (2004)). Some recent empirical evidence suggests that conflict can indeed be reduced through these mechanisms. Berman et al. (2008) show that increased spending on reconstruction programs was correlated with reductions in violence against coalition forces in Iraq. Miguel et al. (2004) find that positive shocks to economic growth - in the form of good rainfall - caused a decrease in conflict in Africa. Dube and Vargas (2007) show that increases in world market prices of agricultural goods reduced conflict in Colombia, presumably because they lead to increases in the return to peaceful economic activities.

Based on these empirical findings, one might expect that development projects can reduce conflict, either by winning the “hearts and minds” of the population or by increasing individuals’ returns to peaceful activities. This paper offers a simple but frequently overlooked explanation for why the opposite can be the case: if insurgents know that development projects will weaken their position, they have an incentive to oppose them, which may exacerbate conflict. This hypothesis is supported by a large body of anecdotal evidence which suggests that insurgents in many countries frequently attack aid workers and infrastructure projects. A recent report on civil counterinsurgency strategies¹ warns that “insurgents strategically target government efforts to win over the population. Indeed, the frequency with which insurgents attack schools, government offices, courthouses, pipelines, electric grids, and the like is evidence that civil [counterinsurgency] threatens them” (Gompert et al. (2009)).

To formalize the mechanism behind these anecdotes, this paper develops a simple theoretical model of bargaining and conflict around development projects. Based on the work of Fearon (1995) and Powell (2004, 2006), the model shows that development projects can

¹The term civil counterinsurgency is used to describe efforts to weaken popular support for insurgent movements by raising the living-standards of the population.

cause conflict if (1) a successful project changes the future balance of power in favor of the government, (2) the insurgents have the ability to hinder the project's successful implementation by violent means, and (3) governments cannot commit to honoring agreements reached before the start of the project. The first two conditions ensure that the insurgents have an incentive to use violence to derail the project, the third condition ensures that governments cannot pay off insurgents in return for allowing the project's peaceful implementation. We explore this logic in more detail in Section 2.

To test our theoretical model, we estimate the causal effect of a large development program - the Philippines' KALAHI-CIDSS - on casualties in armed civil conflict. During the period 2003-08, KALAHI-CIDSS was the Philippines' flagship anti-poverty program with a budget of \$180 million, financed through a loan from the World Bank. Two fundamental challenges have limited previous efforts to identify the relationship between development aid and conflict. First, aid allocation is usually non-random, making it hard to pin down the direction of causality. This is especially problematic in conflict-affected areas because aid assignment could either be positively or negatively correlated with unobserved determinants of conflict. On the one hand, agencies that allocate aid based on need are likely to target conflict-affected areas, since these are home to the poorest and most vulnerable populations. On the other, development agencies might assign programs to more peaceful areas out of concern about the safety of their staff. Second, lack of high quality micro-data has made it difficult for researchers to study the relationship between aid and conflict. The most commonly-used indicators of conflict come from cross-national estimates of battle deaths per country-year, which do not have a high enough resolution to identify the causal effects of micro-level interventions.²

Our empirical analysis addresses these challenges and contributes to the literature in two ways. First, we employ a Regression Discontinuity Design (RDD) to cleanly identify the causal effect of the KALAHI-CIDSS program on conflict violence. Eligibility for the program was restricted to the poorest 25 percent of municipalities in participating Philippine provinces³. This eligibility threshold created a discontinuity in the assignment of aid, which allows us to identify its causal effect by comparing municipalities just above and just below the threshold.⁴ Our second contribution lies in the scope and precision of our data, which provide information on all conflict incidents that involved units of the Armed Forces of the Philippines (AFP) between 2001 and 2008. The data contain information on the dates, location, participating units, and measurable outcomes of each incident, including which party

²Examples of commonly used conflict measures are found in the Correlates of War Intra-State War Dataset and UCDP/PRIO Battle Deaths Dataset. On these data, see Sarkees and Schafer 2000 and Lacina and Gleditsch 2005.

³Municipal poverty rankings were derived from comprehensive local poverty indices based on pre-existing census and survey data, which we describe in detail below. For a full description of the poverty indices, see Balisacan et al. 2002; Balisacan and Edillon 2003. For more on poverty mapping methodology, see Elbers et al 2003.

⁴See Imbens and Lemieux (2008) for a primer on the theory and practice of Regression Discontinuity Designs.

was the initiator and how many government, insurgent, and civilian casualties occurred.⁵

Consistent with the predictions of our model, our empirical analysis finds that the KALAH-CIDSS program exacerbated violent conflict in eligible municipalities. Municipalities just above the eligibility threshold suffered a large and statistically significant increase in violence compared to municipalities just below the threshold. This effect cannot be explained by differences in pre-program violence or other observable characteristics. Our model is further supported by the finding that the increase in violence only lasted for the duration of the program - while insurgent attacks could still affect its success - and was stronger for municipalities that received larger amounts of aid. The majority of casualties was suffered by insurgents and government troops, while civilians appear to have suffered less. We further find that the program caused similar increases in violence initiated by insurgents and government troops, suggesting that the effect is not the result of a one-sided offensive by either party.

The remainder of the paper proceeds as follows. In Section 2, we develop a formal theoretical model of bargaining and conflict in the context of development projects. In Section 3, we give brief overviews of conflict in the Philippines and of the KALAH-CIDSS program. Section 4 contains a description of the data and of the empirical strategy we use to identify the causal effect of the KALAH-CIDSS program on conflict. In Section 5, we present our main empirical results and the results of robustness tests. Section 6 concludes by highlighting the implications of our results for future research and policy.

2.2 Development Projects, Bargaining and Conflict: A Simple Model

This section develops a simple theoretical model of bargaining between an insurgent organization and a local government in a municipality that is scheduled to receive aid in the form of a development project. The model draws heavily on the work of Fearon (1995) and Powell (2004, 2006), which shows that sudden shifts in expected power between conflicting parties can lead to a breakdown of bargaining. While the model is an over-simplified abstraction from the complex reality of interactions between local governments and insurgent groups, we believe that it captures a fundamental mechanism through which development projects - like the Philippines' KALAH-CIDSS program that we analyze empirically - can increase violence in an ongoing civil conflict. There are two main reasons for modelling the interaction between insurgents and local governments as a bargaining game. First, bargaining failures are thought to be a central cause of civil conflict in many contexts (see Blattman & Miguel 2010, for a recent review of the conflict literature). Second, there is strong anecdotal evidence of negotiations between local governments and insurgents over the implementation

⁵For a fuller description of this dataset, see Felter 2005.

of the KALAH-CIDSS program⁶, so that a bargaining model is well suited to describe the context of our empirical analysis.

The intuition behind our model is the following: The central government, in our case the Philippines' Department for Social Welfare and Development (DSWD), plans to implement a development project in a municipality. If the program is successfully implemented, it will shift the balance of power towards the local government and away from insurgents. One possible explanation for this shift in power is that a successful project "wins the hearts and minds" of the population. This idea is at the heart of the US Armed Forces' current counterinsurgency strategy (U.S. Army/Marine Corps 2007, Gompert et al. 2009). It is supported by experimental evidence, which shows that development aid, in the form of community-driven development projects, increases support for the government in the Afghan population Beath et al. (2010). If increased support for the government translates into an increased willingness to supply the government with intelligence about the insurgents, a successful project would make it harder for insurgents to launch successful attacks on the government Berman et al. (2008). Another possible mechanism is that a successful project will decrease poverty and increase returns to peaceful economic activities, making it harder for insurgents to find recruits (this mechanism is suggested by the model of Dal Bo and Dal Bo, forthcoming, and the empirical findings of Dube & Vargas 2007). Regardless of the precise mechanism, insurgents are aware that a successful project will decrease their ability to inflict damage on the government and thus decrease their bargaining power in future negotiations. They therefore have an incentive to launch attacks in order to hinder the project's implementation. This alone would not be enough to explain an increase in violence if the government could pay off insurgents in return for allowing the project's peaceful implementation. However, this may not be possible since the government cannot credibly commit to honoring bargaining agreements reached before the project's start. Since a successful project increases the government's bargaining power, it will have an incentive to renege on any existing agreement after the project is completed, in order to reach an agreement with more favorable terms. As described by Fearon (1995) and Powell (2004, 2006), this inability to credibly commit to a bargaining agreement can lead to conflict by making it impossible to reach a mutually acceptable agreement before the project's start.

To describe the mechanism behind this intuition more formally, consider a simple two-party sequential bargaining model with a finite number of rounds. In each round the two parties, which we call insurgents and government, negotiate over the division of a "pie" of payoffs, the size of which is normalized to one. These payoffs can be thought of as taking either the form of material rewards or the realization of political goals.

In each bargaining round the insurgents demand to receive a portion m_t of the pie⁷. If the government accepts this demand, it receives a payoff of $(1 - m_t)$ from the current round and the insurgents receive a payoff of m_t . If the government rejects the demand, conflict

⁶Authors' interview with KALAH-CIDSS Program Manager Camilo Gudmalin, Department for Social Welfare and Development, Quezon City, Philippines, May 28, 2010.

⁷It will become clear later that the same intuition holds in a slightly modified model in which the government proposes the split

occurs and the insurgents launch attacks on government facilities (these could include but are not limited to individuals and infrastructure associated with the project). In the case of conflict, the government receives a payoff of c_t , and the insurgents a payoff of d_t . We assume that the attacks are costly and destroy a part of the aggregate pie, but that there is a lower bound for the damages each party can incur so that $c_t + d_t < 1$ and both c_t and d_t are greater than zero. By allowing insurgents to conduct costly attacks after their demands have been rejected, the model implicitly assumes that they are able to overcome the commitment problem and make credible threats. While we do not model the precise mechanism by which insurgents commit to attacking, anecdotal evidence suggests that insurgent organizations are often able to follow up their threats with violent attacks when extorting individuals and companies (e.g. Lobjo et al. 2005, Holden & Jacobson 2007).

The timing of the game is as follows. At the beginning of period 1, it becomes known that the municipality is eligible for the project. The insurgents choose m_1 , which the government either accepts or rejects. At the end of period 1, a move of nature decides whether the project is successfully implemented or fails. In period 2 bargaining takes place as in period 1, but there are no more moves of nature.

The model's first key assumption is that conflict in the first period affects the probability that the project is successfully implemented. Anecdotally, there are at least two potential mechanisms for this. First, insurgents can use violent attacks to disrupt the preparations for the project and threaten the security of project staff, leading the implementing agency to withdraw. Second, even if the project continues, insurgents can hinder its successful implementation by attacking project staff and destroying project infrastructure. In the case of KALAH-CIDSS there is anecdotal evidence for both mechanisms. In some municipalities, insurgents launched attacks during the program's preparation phase. In four initially eligible municipalities, insurgent attacks caused the program's implementing agency to abort implementation due to concerns about the safety of its staff⁸. In other municipalities, insurgents attacked construction work that was being funded through the project (DSWD, 2009). For the purpose of the model, we define p^c as the probability that the project is successful if conflict occurs in period 1 and p^p as the probability that the project is successful if there is no conflict in period 1. To keep things simple, we do not model the precise mechanism through which conflict affects the project's implementation, but merely assume that $p^c \leq p^p$.

The model's second key assumption is that the government's cost of conflict in later rounds depends on whether the program was successfully implemented. We thus define the government's payoff in the case of conflict in later periods as $c_t(K)$, where $K = 1$ if the project was successful and $K = 0$ if it failed. As mentioned above, there are two possible mechanisms to explain the program's effect on the government's cost of conflict. First, a successful project may increase the population's support for the government. This makes individuals in the project area more likely to supply the government with information about insurgent's plans and whereabouts, making it easier for the government to defend itself

⁸Authors' interview with KALAH-CIDSS Program Manager Camilo Gudmalin, Department for Social Welfare and Development, Quezon City, Philippines, May 28, 2010.

against attacks (Berman et al. 2008). Second, a successful project may decrease poverty and increase the return to peaceful activities, making it harder for insurgents to find recruits to carry out risky attacks (Dal Bo and Dal Bo, forthcoming, and Dube & Vargas 2007). We remain agnostic about which of these mechanisms causes the change in the government's cost of conflict and merely assume that $c_t(1) < c_t(0)$, so that a successful project reduces the government's cost of conflict.

To see how bargaining can break down in this model, note that in rounds 2 to N , conflict does not affect the government's future cost of conflict and the government will accept any split in which they receive $c_t(K)$ or more. Since $c_t(K)$ is known to the insurgents, they maximize their payoff by demanding m_t such that $m_t = 1 - c_t(K)$, which the government accepts. Thus, in any subgame perfect equilibrium, both parties' payoffs from rounds 2 to N only depend on whether the project was successfully implemented in round 1. The insurgents' payoff from rounds 2 to N is $C(K) = \sum_2^N \beta^{t-1}(1 - c_t(K))$. If bargaining fails in the first period, the insurgents therefore receive an expected payoff of $U^p = d_1 + p^p C(1) + (1 - p^p)C(0)$. On the other hand, the highest possible concession the government can make in the first round is to give the entire pie of value 1 to the insurgents. The highest possible payoff to the insurgents from a peaceful bargaining solution is therefore $U^c = 1 + p^c C(1) + (1 - p^c)C(0)$. The insurgents will only accept a peaceful solution if $U^p \geq U^c$, so that a condition for successful bargaining is:

$$(p^p - p^c)(C(0) - C(1)) \leq 1 - d$$

If this condition fails, conflict occurs in the first round since the highest transfer the government can make in exchange for peace cannot compensate the insurgents for their loss of bargaining power in future periods. The condition shows that conflict is more likely if the project causes a larger reduction in insurgents' future bargaining power, $C(\cdot)$. If the project has no effect on $C(\cdot)$, or if conflict has no effect on the probability of the project's implementation, bargaining will always be successful, since we assume that conflict is costly, so that $1 - d_1 > 0$. As in the models of Fearon (1995) and Powell (2004, 2006), bargaining fails if the (potential) shift in bargaining power from one round to the next is large compared to the inefficiency of conflict. The bargaining condition also shows that conflict is more likely if insurgents can credibly threaten the program's implementation, i.e. if first-round conflict has a large effect on the program's probability of success. As a consequence, the model predicts that conflict only lasts as long as it affects the project's probability of success, which is only in round 1. In later rounds, after the project has been implemented (or failed), neither party has an incentive to engage in conflict, so that bargaining is always successful. A related prediction is that being eligible for a project can lead to conflict even if the project is eventually not implemented.

At this point it should be noted that bargaining only fails because the government cannot credibly commit to not using the increase in bargaining power it gains from the project. If the government could fix its second-round offer in the first round, the game would collapse into a single-round bargaining game with a peaceful outcome. In addition, bargaining failure

depends on the assumption of discrete time. If bargaining took place in continuous time (in other words, if bargaining rounds become infinitely short), the government would be able to devise a continuous stream of payments that the insurgents prefer to conflict (e.g. Schwarz and Sonin 2007). While it is thus possible to devise a continuous-time model in which development projects do not lead to a breakdown of bargaining, we believe that the discrete time assumption is better suited for the present context because negotiations with insurgents pose considerable logistical challenges so that there are likely to be substantial lags between successive rounds of bargaining.

2.2.1 Comparison with other models of conflict:

This section reviews the predictions that other models make about the effect of development projects on conflict and compare these predictions to those of our model. The “hearts and minds” model of conflict, as described by Berman et al. (2008), predicts that development projects cause a decrease in conflict. The model’s key assumption is that development projects increase the population’s support for the government, which leads individuals to be more willing to share information about insurgents with the armed forces. Insurgents therefore find it more difficult to launch attacks in areas affected by development projects, which leads to a decrease in conflict.

A related model is that of Dal Bo and Dal Bo (forthcoming), who take a general equilibrium approach to modeling conflict. In their model, conflict is a consequence of low returns in the peaceful economy and high returns in the “conflict economy”, i.e. in appropriating the economy’s output by violent means. Assuming that conflict is a labor intensive activity, the authors find that increases in the returns to labor cause a decrease in conflict, because fewer individuals are willing to participate in conflict, making it harder for insurgents to find recruits. On the other hand, increases in the returns to capital cause an increase in conflict, because they increase the value of the total output to be fought over. For the case of development projects, their model’s predictions therefore depend on whether the project increases the return to labor or the return to capital. Other general equilibrium models, like the one of Grossman (1999), predict that economic growth can increase conflict by increasing the amount of resources to be fought over, regardless of whether growth favors labor or capital. Regardless of the sign of the effect, general equilibrium models predict that a development project’s effect on conflict materializes only after the project has started and persists as long as the project affects the economy. These predictions differ from those of our model, which predicts that a development project can cause conflict even before its implementation has begun and that conflict only lasts as long as it is being implemented. Our model follows those of Berman et al. (2008), and Dal Bo and Dal Bo (forthcoming), in assuming that development projects can reduce insurgents’ capacity to launch successful attacks, either by making it harder to find recruits or making it harder to operate clandestinely. The point of departure is that our model explicitly incorporates the strategic interaction between insurgents and the government, which can lead to conflict over a project’s implementation.

Finally, models of bargaining with asymmetric information make predictions that are

similar to those of our model. Suppose, for example, that local governments know the exact benefits they will receive from successfully implementing a development project while the insurgents do not. This means that insurgents do not know the government's willingness to pay to avoid conflict, which may make it optimal to make demands that are rejected with positive probability. In a dynamic game, the government will have an additional incentive to reject high demands in order to affect insurgents' beliefs about its willingness to pay to avoid conflict in later rounds. Multiple-round games of asymmetric information have been described by Fudenberg & Tirole (1991). Translated to the present context, their results suggest that asymmetric information can cause conflict over the implementation of development projects, but that conflict decreases over time as insurgents learn about the government's true willingness to pay to avoid conflict. While these predictions are similar to those of our model - conflict initially increases in municipalities eligible for a development project, but returns to baseline levels over time - we believe that asymmetric information models cannot plausibly explain conflict around development projects, especially in countries with long-running conflicts. For example, insurgents in the Philippines have been attacking road construction and other infrastructure projects for over 30 years. While asymmetric information might explain this type of attack for the first few years of the conflict, the information asymmetry should disappear over time as insurgents learn about the government's willingness to pay to avoid attacks. The fact that insurgents keep attacking infrastructure projects even after years of conflict is difficult to explain by asymmetric information and suggests that a different mechanism is at work.

2.3 Empirical Setting: Conflict and Development Projects in the Philippines

2.3.1 Violent conflict in the Philippines

Civil conflict in the Philippines has been ongoing for over four decades, caused more than 120,000 deaths, and cost the country an estimated \$2-3 billion (Schiavo-Campo and Judd 2005). During the period we study in this paper, 2001-2008, the two largest insurgent organizations active in the country were the New People's Army (NPA) and the Moro Islamic Liberation Front (MILF). Below we briefly describe these organizations and a third category of insurgents, the so-called "lawless elements".

New People's Army (NPA)

As the armed wing of the Communist Party of the Philippines (CPP), the New People's Army is a class-based movement that seeks to replace the Philippine government with a communist system. Since taking up arms in 1969, the NPA has relied on pinprick ambushes and harassment tactics rather than conventional battlefield confrontations against government armed forces. The NPA's current strength is estimated at 8000 armed insurgents, down from

a 1986 peak of approximately 25,000 insurgents, who exerted influence in 63 of the (then) 73 Philippine provinces (Felter 2005). The NPA operates mostly in rural areas - its military strategy relies on small guerrilla fronts that are deployed in and around villages. Due to its guerilla tactics and lack of a broad ethnic or religious constituency, the NPA's activities require significant support from the Philippines' rural poor who supply most of the group's recruits and logistical support⁹. According to our dataset, the NPA is by far the most active insurgent organization in the Philippines. During the period 2001-08, the NPA was involved in 65% of all incidents in our data for which the enemy organization was reported (about 10% of incident reports do not report an enemy organization).

Moro Islamic Liberation Front (MILF)

The Moro Islamic Liberation Front is a separatist movement fighting for an independent Muslim state in the *Bangsamoro* region of the southern Philippines. The MILF was formed in 1981, when the group's founders defected from the Moro National Liberation Front (MNLF), another longstanding southern Philippines insurgent movement, due to disagreement about the means by which to pursue independence. After the split, the MILF escalated armed conflict against the government while the MNLF signed a peace accord in 1996 that created the Autonomous Region of Muslim Mindanao (ARRM). The MILF's core grievances stem from government efforts to retitle lands considered by the southern Muslim population to be part of their ancestral homeland and the group reportedly enjoys broad support in the Muslim population. (Kreuzer and Werning 2007).

With an estimated 10,500 fighters under arms, the MILF is larger than the NPA. Furthermore, its tactics are more manpower intensive. While the NPA relies mainly on small unit guerrilla tactics, it is not uncommon for MILF commanders to mass their forces into larger units to fight semi-conventional battles against government forces (Felter 2005). However, because of its narrow geographic focus, the MILF not a major cause of conflict in our data, being involved in only 10% of all reported incidents.

Lawless Elements (LE)

The term "lawless elements" refers to small, loosely-allied bands of guerrilla and criminal groups operating across the Philippines. Some of these groups are local manifestations of the NPA, the MILF, or the Abu Sayyaf Group (ASG)¹⁰. Many others are criminal organizations that employ guerrilla-like tactics but use violence primarily as part of criminal activities such as kidnapping-for-ransom rather than to pursue political objectives. During the period

⁹Chapman (1987) and Jones (1989) provide detailed histories of the NPA. For an insider view on the war from an AFP officer's perspective, see Corpus (1989). From an NPA leadership perspective, see Sison and Werning (1989).

¹⁰The ASG, a high profile southern Philippine terrorist organization with well established links to al-Qaeda, does not figure in our analysis because the provinces in which the ASG operates are not eligible for KALAHY-CIDSS.

2001-2008, Lawless Elements were involved in roughly 25% of all conflict incidents involving AFP units

2.3.2 The KALAHY-CIDSS Program

KALAHY-CIDSS is a major development program in the Philippines. Designed to enhance local infrastructure, governance, participation, and social cohesion, KALAHY-CIDSS has been the Philippines' flagship development program since 2003. As of mid-2009, more than 4000 villages in 184 municipalities across 40 provinces had received KALAHY-CIDSS aid.¹¹ Plans to expand KALAHY-CIDSS are currently being made, with the aim of doubling the number of recipient municipalities during the program's next phase.

Run by the Philippine government's Department of Social Welfare and Development and funded through World Bank loans, KALAHY-CIDSS aims to promote local governance reform and development by supporting bottom-up infrastructure and institution-building processes. As a community-driven development (CDD) program, KALAHY-CIDSS is representative of a common type of development intervention. The World Bank lends more than two billion dollars annually for CDD projects (Mansuri and Rao 2004) and donors are increasingly making use of CDD programs in conflict-affected countries. Over the last decade, for example, CDD programs have been launched in Afghanistan, Angola, Colombia, Indonesia, Nepal, Rwanda, and Sudan.¹²

KALAHY-CIDSS follows a standard CDD template. First, each participating municipality receives a block grant for small-scale infrastructure projects. Within the municipality, each village (*barangay* in Tagalog) holds a series of meetings in which community members draft project proposals. Villages then send democratically elected representatives to participate in municipal inter-*barangay* fora, in which proposals are evaluated and funding is allocated. Proposals are funded until each municipality's block grant has been exhausted. Once funding has been allocated, community members are encouraged to monitor or participate in project implementation.¹³

The amount of aid distributed through KALAHY-CIDSS is substantial. Participating municipalities receive PhP300,000, or approximately \$6000, per village in their municipality. The average municipality has approximately 25 villages, making the average grant approximately \$150,000, or about 15% of an average municipality's annual budget. Over the course of the program, the project cycle is repeated three times—occasionally four—meaning that on average, participating municipalities receive a total of between \$450,000 and \$600,000 dollars.

¹¹As of March 2010, there were 80 provinces and 1496 municipalities in the Philippines. A complete list of all Philippine administrative units is available from the National Statistical Coordination Board.

¹²For an overview, see World Bank 2006. . See also Mansuri and Rao 2004. For an assessment of the impact of Indonesia's CDD program, called KDP, on local corruption and public goods provision, see Olken 2007; 2010. On KALAHY-CIDSS and social capital, see Labonne and Chase 2009.

¹³See Parker 2005 for a detailed overview of the KALAHY-CIDSS process.

Targeting

KALAHY-CIDSS was designed in the early 2000s as a nationwide anti-poverty program that would target aid to the poorest populations in the Philippines. Aid was targeted following a two-stage approach. First, 42 eligible provinces were selected, among them the 40 poorest based on estimates from the Family Income and Expenditure Survey (FIES). To identify the poorest municipalities within eligible provinces, a team of economists was hired to estimate municipal poverty levels using data from the 2000 National Census, the Family Income and Expenditure Survey (FIES) and rural accessibility surveys (Balisacan et al. 2002; Balisacan and Edillon 2003). Poverty levels were estimated following the poverty mapping method of Elbers et al (2003).¹⁴ The first step of this method is to estimate the relationship between measures of household expenditure, which are only available for a subset of municipalities, and variables from census data and accessibility surveys, which are available for all municipalities. The estimated relationship between census and accessibility variables and poverty in this subset of municipalities was then used to predict poverty levels for all municipalities. Based on the estimated poverty levels, only the poorest 25% of municipalities within each participating province were eligible for aid from KALAHY-CIDSS. The arbitrary nature of this eligibility cutoff enables us to identify the program's casual effect through a regression discontinuity design.

Table 2.1: Variables Used to Determine KALAHY-CIDSS Eligibility

Variable	Weight
Proportion of Households with Electricity	4.41
Proportion of Households with Water-Sealed Toilets	2.83
Proportion of Households with Access to Level III Water Systems	4.56
Proportion of Houses with Roofs Made of Strong Material	4.27
Proportion of Houses with Walls Made of Strong Material	7.47
Proportion of Population Aged 0-6	23.7
Proportion of Population Aged 7-14	18.05
Proportion of Population Aged 15-25	5.96
Proportion of Population Aged >25	0.08
Educational Attainment of All Family Members Relative to Potential	8.28
Density of Good Barangay Roads that are Passable Year-Round	10
Road Distance to Provincial Center of Trade	10

Source: Balisacan and Edillon 2003

Table 2.3.2 shows which variables were used in calculating the poverty index and the weights that they were assigned. For the first ten variables, the weights were determined by the regression of the poverty mapping approach, the weights of the last two variables were chosen by the researchers.

¹⁴Details can be found in Balisacan et al. 2002; Balisacan and Edillon 2003.

Timeline

Figure 2.1: Timeline of KALAHY-CIDSS Implementation

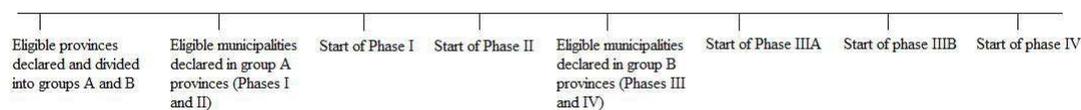


Figure 2.1 shows the timeline of the program, which was rolled out on a staggered schedule. Participating provinces were first divided into two groups, Group A and Group B. Eligible municipalities in Group A and Group B provinces were then divided into phases with different start dates.

Table 2.2: Timetable of KALAHY-CIDSS

Phase / Duration	Duration	Municipalities	Barangays
I	Jan 2003 - June 2006	11	201
II	June 2003 - Dec 2006	56	1291
III A	Oct 2004 - Dec 2007	34	883
III B	Jan 2006 - Dec 2008	29	727
IV	Aug 2006 - July 2009	54	1127
Total	Jan 2003 - July 2009	184	4229

Source: Department of Social Welfare and Development

Table 2.3.2 displays the start dates and the number of municipalities that participated in each phase of the program. Group A municipalities learned their eligibility status in December 2002 and began receiving project aid in either January 2003 (Phase I) or June 2003 (Phase II).¹⁵ Group B municipalities were informed of their eligibility status in October 2003 and implementation began in October 2004 (Phase IIIA), January 2006 (Phase IIIB), or August 2006 (Phase IV).

Eligibility and participation

In each eligible municipality, implementation of the program was preceded by a “social preparation phase” in which the program was introduced to the public and preparations were made for its implementation. During this time, eligible municipalities were required to ratify a memorandum of understanding and put in place basic institutional mechanisms required for implementation. If an eligible municipality failed to meet these conditions by the time KALAHY was scheduled to be launched, it was declared ineligible for the program. There were some cases in which eligible municipalities failed to comply with program requirements and were replaced by municipalities that were not initially eligible. In some other cases,

¹⁵Phase I was a pilot phase whose municipalities were outside of the bottom quartile of poverty.

initially eligible municipalities were dropped from the program because of concerns about the security of program staff¹⁶.

2.4 Empirical Strategy

2.4.1 Regression discontinuity design

To test our theoretical model of bargaining and conflict in the context of the Philippines, we use a regression discontinuity (RD) design to estimate the causal effect of KALAH-CIDSS - the country's flagship anti-poverty program in the period 2003-2008 - on the intensity of violent conflict. The RD approach is made possible by the arbitrary eligibility threshold used to target the program. Targeting followed a two-staged approach. First, 42 eligible provinces were selected, among them the 40 poorest based on estimates from the Family Income and Expenditure Survey (FIES). The poverty levels of all municipalities within the eligible provinces were estimated using a poverty mapping methodology based on a combination of data from FIES and the 2000 Census of the Philippines (Balisacan et al. 2002, 2003). In each eligible province, municipalities were ranked according to their poverty level and only the bottom quartile was eligible for KALAH-CIDSS. The arbitrary cutoff at the 25th percentile of poverty created a discontinuity that we exploit to identify the program's causal effect on violent conflict. In essence, we estimate the causal effect by comparing the outcomes of municipalities just below the eligibility threshold with those of municipalities just above it. The identification assumption of the RD design is that municipalities close to the threshold on either side do not differ in unobserved variables that affect conflict, so that any change in conflict across the threshold can be attributed to the KALAH-CIDSS program¹⁷.

The running variable in our RD regressions is the distance of the municipality's poverty rank from the provincial eligibility threshold¹⁸. Since only municipalities in the poorest quartile were eligible, the provincial threshold was calculated by dividing the number of municipalities in each province by four and then rounding to the nearest integer¹⁹. This

¹⁶Authors' interview with KALAH-CIDSS Program Manager Camilo Gudmalin, Department for Social Welfare and Development, Quezon City, Philippines, May 28, 2010.

¹⁷Robustness tests for this identification assumption, using pre-treatment conflict and other observable variables, are presented in the Results section of the paper.

¹⁸Unfortunately, data for the last two variables used for the poverty ranking - density of good barangay roads and road distance to the provincial center of trade, both in 2000 - are no longer available, so that we are unable to reproduce the poverty index that formed the basis of the ranking. However, our regressions control for the remaining ten Census variables used for the ranking. Balance tests show that there are no discontinuous breaks in these variables or other observable municipal characteristics at the eligibility threshold. To avoid bias from omitting the road density and road distance variables (or any other unobserved municipal characteristics) some of the regressions presented in the Results section include municipality fixed effects.

¹⁹In cases where dividing a province's number of municipalities by 4 ended on .5, the number of eligible municipalities was rounded down more often than up, so in calculating municipalities' normalized poverty rankings, we follow suit and round down at .5. Doing so improves the accuracy with which the normalized

threshold number was then subtracted from the municipality's actual poverty rank to obtain the municipality normalized poverty rank. For each participating province, the richest eligible municipality has a normalized poverty rank of zero and the poorest ineligible municipality has a normalized poverty rank of one. Formally, the RD estimator of the causal effect of eligibility for KALAH-CIDSS is

$$\tau_{RDD} = \lim_{x \downarrow c} [Y_i | X_i = x] - \lim_{x \uparrow c} [Y_i | X_i = x]$$

where Y_i is municipality i 's outcome, X_i is the municipality's normalized poverty rank and c is the threshold that determines assignment (i.e. the 25th percentile of each municipality's poverty index). Verbally, the estimated causal effect is the difference in the limits of the expected outcome as we approach the eligibility threshold from above and below. In practice, linear regressions are fitted on both sides of the threshold and the limits are estimated by extrapolating the regression lines²⁰

Figure 2.2: Probability of KALAH-CIDSS Participation by Distance from Poverty Threshold

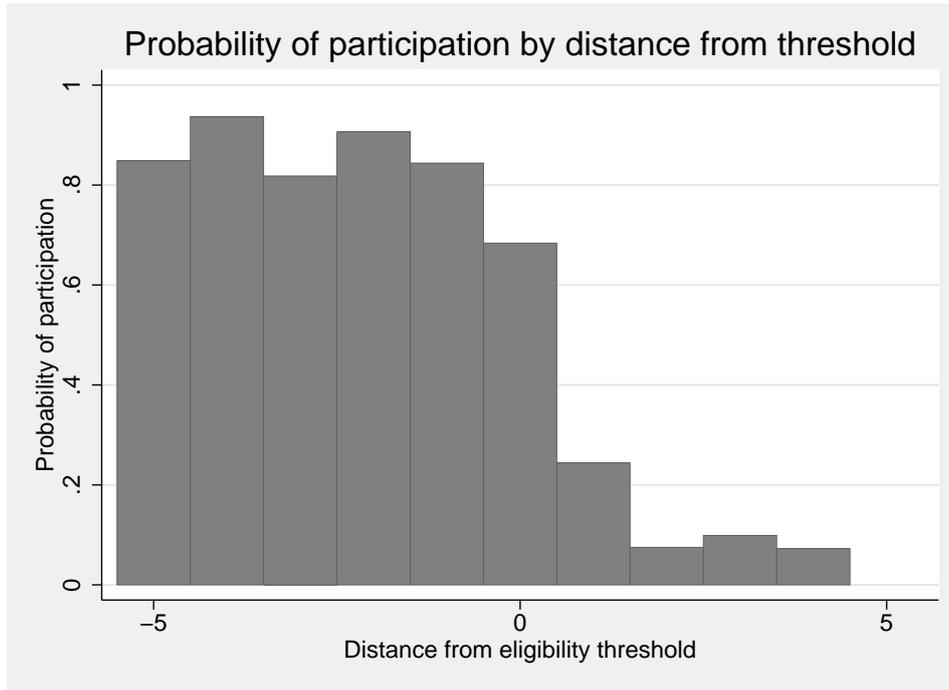


Figure 2.2 plots the observed probability of participating in KALAH-CIDSS against the normalized poverty rank. The graph shows that the probability of participation decreases sharply at the eligibility threshold, though some eligible municipalities did not participate and were replaced by municipalities above the threshold. The probability of participation is

poverty rank predicts participation in KALAH-CIDSS but otherwise does not affect the empirical results.

²⁰For further details, see Imbens and Lemieux 2008.

somewhat lower for municipalities at the eligibility threshold, i.e., those with a normalized poverty rank of zero. A possible explanation is that the implementing agency had room for discretion on the margins when calculating the number of eligible municipalities per province. The standard procedure for determining the number of eligible municipalities per province was to divide the number of municipalities in each province by four and then to round to the nearest integer, but in some cases the number was rounded down due to budget constraints, particularly if the municipality at the threshold did not express a strong interest in participating in KALAH-CIDSS. The fact that not all eligible municipalities participated in the program might suggest the use of a “fuzzy” RD design that uses eligibility as an instrument for participation. However, our theoretical model suggests that eligibility itself affects conflict and that participation is an endogenous outcome. We therefore estimate the “intention to treat” effect - the effect of eligibility regardless of later participation status.

2.4.2 Data

Three types of data are used in this paper: Program data from the Department for Social Welfare and Development (DSWD), the government agency responsible for implementing KALAH-CIDSS; armed conflict data from the Armed Forces of the Philippines (AFP); and population data from the Philippines 2000 National Census. All variables in our analysis are measured at the municipality level.

Program data

We use KALAH-CIDSS program data from the Philippines Department for Social Welfare and Development. These data include information on municipalities’ eligibility for KALAH-CIDSS, whether or not eligible municipalities participated in the program, and the phase or timing of the program’s roll out. These data are available from 2003 through 2009, the full duration of the program to date.

Conflict data

Our data on civil conflict and violence come from the Armed Forces of the Philippines’ (AFP) records of civil conflict-related incidents. The data were derived from the original incident reports of deployed AFP units that operated across the country from 2001 through 2008. With authorization from the AFP’s Chief of Staff, researchers were hired and trained to compile and code the field reports to an unclassified database. The incident-level data contains information on the date, location, the involved insurgent group or groups, the initiating party, and the total number of casualties suffered by government troops, insurgents, and civilians (see Felter 2005). The data are comprehensive, covering every conflict-related incident reported to the AFP’s Joint Operations Center by units deployed across the country. In total, the database documents more than 21,000 unique incidents during this period, which led to just under 10,000 casualties. The depth, breadth, and overall quality of the AFP’s

database makes it a unique resource for conflict researchers and enables credible assessment of the average impact of KALAH-CIDSS on the dynamics of insurgent and counterinsurgent violence.

The outcome of interest for our analysis is the number of casualties in conflict incidents. We believe that this outcome best captures the true intensity of conflict; better than other outcomes such as the number of conflict incidents. In particular, there is reason to believe that the number of incidents does not tell us much about the actual intensity of conflict. Even in municipalities where local governments have negotiated peace agreements with insurgents, AFP units still have an incentive to conduct patrols and other operations, if only to convince their superiors that their deployment in the municipality serves a useful purpose. It is therefore likely that AFP units will encounter insurgents on a regular basis regardless of whether a local peace agreement is in place or not, making the number of incidents a weak measure of conflict. However, the intensity with which AFP and insurgent units engage each other in combat, and as a consequence the resulting number of casualties, clearly depends on whether a (possibly informal) peace agreement is in place or not. We therefore believe that using incidents as the outcome of interest is likely to understate the effect of KALAH-CIDSS on the intensity of conflict, and instead use the number of casualties as the outcome of interest.

Other data

Data from the Philippines' 2000 National Census are also used. The primary purposes for using these data are to test the plausibility of the RD identifying assumption and to check the sensitivity of the results to alternative specifications. These variables are described in more detail below.

2.4.3 Variables

Our main dependent variable is *total casualties*. This variable measures the total number of people killed and wounded in conflict-related incidents per municipality-year from 2001-2008 as documented in the AFP's field reports. The total casualties variable is calculated as the sum of government casualties, insurgent casualties, and civilian casualties.²¹

To study the dynamics of civil conflict—who suffered and inflicted the casualties—we break down the total casualties variable by individual parties to the conflict.

To this end, the first variable we use is *government casualties*. *Government casualties* measures the number of government-affiliated troops killed and wounded in action, per municipality-year, from 2001-2008, as documented in the AFP's field reports. The variable counts the casualties suffered by all Philippine government armed forces conducting internal security operations during the study period, including “elite” units such as Special Forces

²¹These data were originally made available by the AFP's Chief of Staff and the staff in the Office of the Deputy Chief of Staff for Operations J3 in their unclassified form. For a full description of the conflict data, see Felter 2005, 48-67.

and Scout Rangers units; conventional, or “regular,” units such as infantry battalions; and local auxiliary units, such as Citizen Armed Force Geographical Units (CAFGUs), that were administered by the AFP.

The second variable we use is *insurgent casualties*. *Insurgent casualties* measures the number of insurgents killed and wounded in action, per municipality-year, from 2001 to 2008, as documented in the AFP’s field reports. The variable counts the casualties suffered by all insurgent movements operating in a given municipality.

The third variable we use is *civilian casualties*. *Civilian casualties* measures the total number of civilians killed and wounded in conflict-related incidents, per municipality-year, from 2001 to 2008, as documented in the AFP’s field reports. The variable counts the total number of casualties suffered by civilians in a given municipality but does not distinguish between insurgent-inflicted civilian casualties and government-inflicted civilian casualties.

To test whether insurgencies with differing aims and organizational structures behaved differently in response to the aid intervention, we measure conflict intensity by insurgency. These variables measure the total number of people killed and wounded—government, insurgent, and civilian—per municipality-year, from 2001-2008, in conflict-related incidents involving the communist guerrilla movement the New People’s Army and the Muslim separatist movement the Moro Islamic Liberation Front (MILF). These variables are named *Casualties - NPA incidents*, and *Casualties - MILF incidents*.

We also use a number of municipality characteristics as controls. *Municipality Population* measures the total number of residents per municipality in year 2000 as measured by the Philippines’ 2000 National Census. As Table 2.5.1 below shows, the average population of control and treatment municipalities was 29,578.²² *Highway access* captures the percentage of villages per municipality with access to a national highway. Taken from the *barangay* characteristics section of the Philippines’ 2000 National Census, the data show that 68 percent of the villages in municipalities included in our analysis were recorded as having national highway access in 2000. *Timber* measures the amount of land per municipality, in squared kilometers, covered with timber. Data on timber come from the Philippines’ National Statistics Coordination Board. *Affected by NPA* is an indicator for whether the NPA reportedly had a local presence. These estimates were made in 2001, two years before the beginning of the KALAHYON-CIDSS treatment. Importantly, the affectation data are based on intelligence estimates of insurgent *presence*, rather than *violence*, providing a separate measure of insurgent activity.²³

²²This is the average population of municipalities in eligible provinces using a normalized poverty rank bandwidth of three. The average population for all Philippine municipalities in 2000 was 47,043. The lower average population for municipalities’ covered in this study reflects conventional wisdom on poverty in the Philippines—rural areas tend to be the most stricken with poverty.

²³These data were originally made available in unclassified form by the AFP’s Office of the J2. See Felter 2005, 39. The primary limitation of these data is a lack of comparable data on MILF presence. While having the same kind of data on MILF presence would be ideal, the NPA data are extremely useful. They provide a measure, however crude, of the density of insurgent “control” within municipalities—a variable posited to be important in previous theoretical work but that is nearly always unobserved in empirical studies of conflict. (Kalyvas 2006). Of all of the Philippines’ insurgent movements, however, the NPA provides the

As additional controls, we include municipality-level pre-treatment demographic characteristics. The first is an index of *ethnic fractionalization*. Computed using microdata from the 2000 National Census, this variable gives the probability that two individuals drawn randomly from a municipality are from different ethnic groups.²⁴ The second is a similar index measuring religious differences, also based on year 2000 census microdata, which we call *religious fractionalization*. We also include a control for *percentage Muslim* that measures the percentage share of Muslims, by municipality, based on 2000 census data.²⁵

Finally, we control for most of the variables that were used to calculate the municipal poverty index used to determine eligibility for KALAH-CIDSS. The variables used to calculate the poverty index are shown in Table 2.3.2. We control for the first ten of these variables, which come from the 2000 Census: *Age 0-6*, *Age 7-14*, *Age 15-25* and *Age 25+*, denote the proportion of the municipal population that falls into the respective age range. *Electricity*, *Water-sealed toilet* and *Level III water system* denote the proportion of households that have access to the respective facilities. *Strong walls* and *Strong roof* denote the proportion of households whose dwelling has walls or a roof made of “strong” materials²⁶ Unfortunately, data for the last two variables used for the poverty ranking—density of good barangay roads and road distance to the provincial center of trade, both in 2000—are no longer available and consequently cannot be controlled for in our regressions or used to reproduce the poverty indices. However, the balance tests in the next section show that there are no discontinuous breaks at the eligibility threshold in any observable municipal characteristics including pre-treatment conflict, which suggests that the identifying assumption of the RD design holds. To avoid bias from omitting these variables (or any other unobserved municipal characteristics) some of the regressions presented in the next section include municipality fixed effects.

2.5 Results

In this section, we present results from the regression discontinuity approach based on the poverty threshold that determined eligibility for the Philippines’ KALAH-CIDSS program. As mentioned above, aid from KALAH-CIDSS was targeted following a two-staged approach. First, 42 eligible provinces were selected, among them the 40 poorest

most leverage empirically since it operates nationwide. Despite the incompleteness of the insurgent presence data, the pre-treatment NPA presence variable can consequently help us determine whether pre-existing insurgent presence influences either the intervention or the outcomes of interest.

²⁴This variable is similar to the commonly employed ethnolinguistic fractionalization (ELF) index used in Fearon and Laitin’s (2003) study of civil war onset, which was based on 1964 country-level data from *Atlas Narodov Mira*.

²⁵Area experts have suggested that predominantly Muslim areas have unique conflict dynamics due to clan-based social and political structures and exogenous historical circumstances. See, e.g., Abinales 2000; Kreuzer 2005; 2009.

²⁶Materials counted as “strong” are galvanized iron or aluminium, concrete, clay tiles and asbestos for roofs and concrete, brick, stone, wood, galvanized iron or aluminium, asbestos and glass for walls

based on estimates from the Family Income and Expenditure Survey (FIES). Next, the poverty levels of all municipalities within the eligible provinces were estimated using a poverty mapping approach based on a combination of data from FIES and the 2000 Census (Balisacan et al. 2002, 2003). Within each province, municipalities were ranked according to their poverty level and only the bottom quartile was eligible for KALAH-CIDSS. Our regression discontinuity design compares eligible and ineligible municipalities close to the eligibility threshold to estimate the causal effect of the KALAH-CIDSS program. For most of our analysis, we restrict the sample to contain only the four municipalities closest to the eligibility threshold in each participating province: the two richest municipality that are still eligible for the program and the two poorest that are not. In the language of regression discontinuity design, our running variable is the distance from the provincial eligibility threshold and we choose a rectangular kernel with a bandwidth of two ranks. This restriction makes sure that the identifying assumption of the regression discontinuity design holds: that eligible and ineligible municipalities within our sample do not differ on observed and unobserved characteristics, except for program eligibility. To test this assumption we present evidence that municipalities near the eligibility threshold on either side do not differ significantly on any observable characteristics, including levels of conflict prior to the start of KALAH-CIDSS.

We then present a graphical comparison of trends in conflict casualties in eligible and ineligible municipalities near the eligibility threshold. This comparison shows that the number of casualties in eligible municipalities sharply increases after eligibility for KALAH-CIDSS is announced, while the number of casualties in ineligible municipalities remains virtually unchanged.

To obtain a quantitative estimate of the causal effect of eligibility of KALAH-CIDSS on conflict violence, we present results of negative binomial and linear regressions that exploit the discontinuity described above. In addition to the comparison of eligible and ineligible municipalities across the threshold, the regressions exploit the timing of the program's implementation. Specifically, the program's causal effect is estimated as the "double-difference" of casualties in eligible and ineligible municipalities, before and after the roll-out of KALAH-CIDSS. As is standard for the regression discontinuity design, all regressions control for the running variable, which in our case is the distance from the eligibility threshold in ranks. All standard errors are adjusted for clustering at the province level. To test the robustness of our results to the choice of bandwidth, we also present estimates based on larger and smaller bandwidths.

2.5.1 Summary statistics and balance tests

To get an idea of the intensity of conflict in the Philippines in the period of observation, 2001-08, Table 2.5.1 reports summary statistics of some conflict outcomes. The first column reports average outcomes per municipality per year for the whole of the Philippines, the second column reports outcomes for our restricted sample of municipalities within two ranks of the eligibility threshold for KALAH-CIDSS.

Table 2.3: Summary Statistics of Conflict Outcomes, 2001-08

	All Philippines	RD Sample
Incidents	1.41 (3.92)	2.18 (3.72)
Violent incidents (casualties >0)	0.30 (0.92)	0.49 (1.00)
Casualties	0.66 (3.16)	0.97 (2.57)
AFP casualties	0.32 (1.96)	0.49 (1.62)
Insurgent casualties	0.20 (1.16)	0.31 (1.15)
Civilian casualties	0.17 (1.47)	0.20 (0.89)
Cas. in AFP initiated inc.	0.28 (1.89)	0.40 (1.51)
Cas. in insurgent initiated inc.	0.38 (2.03)	0.56 (1.80)
Observations	14650	1285
Municipalities	1632	160

Reported values are means, standard deviations are in parenthesis

Table 2.4: Balance of Observed Variables Across Eligibility Threshold

	Eligible	Ineligible	Difference
Population ('000)	30.2	28.56	1.6 (2.9)
Area (km ²)	0.029	0.036	7.8×10^{-3} (5.2×10^{-3})
Highway Access (%)	69.1	67.2	1.9 (4.7)
Forest (km ²)	8.4×10^{-3}	8.2×10^{-3}	2×10^{-4} (2.4×10^{-3})
Affected by NPA in 2001 (%)	41.8	42.0	-0.2 (7.9)
Percent Muslim	3.2	4.1	-0.9 (2.2)
Ethnic fractionalization	0.32	0.29	0.03 (0.05)
Religious fractionalization	0.32	0.30	0.03 (0.04)
Municipalities	81	79	160

Standard errors of differences in parentheses

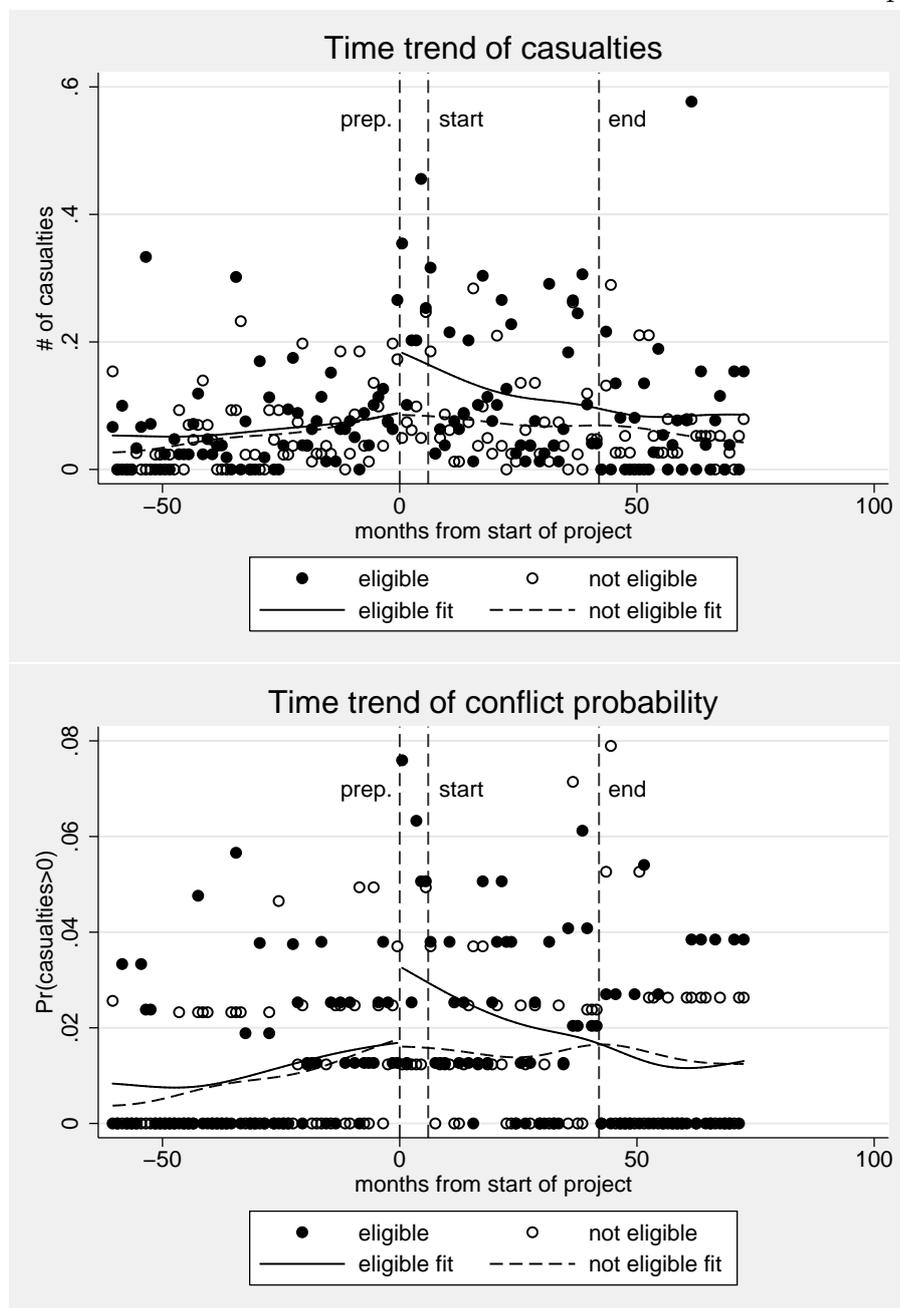
As a robustness test of the identifying assumption of the regression discontinuity design, Table 2.5.1 compares observable variables of eligible and ineligible municipalities in our restricted sample. All of these variables were measured in 2000 or 2001, at least two years before the start of the KALAH-CIDSS program. Significant differences between eligible and ineligible municipalities would point to a violation of the identifying assumption that no unobserved variable changes discontinuously across the eligibility threshold. To test this, we conduct t-tests for equality of means of eligible and ineligible municipalities in our sample. The results show that none of the variables are significantly different at the 10% level, which increases our confidence in the identifying assumption of the RD design. To further rule out possible bias from discontinuous changes in variables across the threshold, some of the regressions presented below include municipality fixed effects to control for all unobserved time-invariant differences between eligible and ineligible municipalities.

2.5.2 The effect of KALAH-CIDSS on conflict casualties

Graphical evidence

Figure 2.5.2 displays the time trend of conflict in eligible and ineligible municipalities within our restricted sample of municipalities within two ranks of the eligibility threshold. The variable on the x-axis is months from the start of KALAH-CIDSS in the municipality.

Figure 2.3: Time Trend of Conflict in Treatment and Control Municipalities



Scatter dots denote monthly averages of conflict for municipalities within two ranks above and below the eligibility threshold. Fitted lines are estimated by local quadratic regression.

The plots contain three vertical dashed lines that denote the timeline of the KALAHICIDSS program. The first line at $t=0$ denotes the start of the social preparation phase, the second line at $t=6$ denotes the start of the actual program implementation and the third line at $t=42$ denotes the end of the last program cycle. In the top plot, the variable on the y-axis is the number of conflict casualties, in the bottom plot it is the probability of having at least one casualty in a given month. The scatter dots mark monthly averages, the fitted line is obtained by local quadratic regression. To clarify the effect of KALAHICIDSS, the figure displays two fitted lines, one for the pre-program period and one for the period after the program's start. The figure shows that eligible and ineligible municipalities experienced similar levels of conflict in the pre-program period. However, the conflict in eligible municipalities increased sharply in the first month after the start of the social preparation phase - while the number of casualties in ineligible municipalities remained virtually unchanged. The difference in conflict between eligible and ineligible municipalities then becomes smaller and disappears as the program ends.

Quantitative evidence

Table 2.5.2 presents regression estimates of the effect of KALAHICIDSS on conflict casualties. The estimated causal effect of eligibility for KALAHICIDSS on conflict casualties is the regression coefficient associated with the interaction of eligibility and the program time-period. Since the program was scheduled to last for 3 years, we define the program time period as the three years after the start of the program. The program time-period thus depends on which phase of the program a municipality was covered. For municipalities covered in Group A (Phases 1 and 2), the program period is 2003-2005, since implementation began in 2003. For municipalities in Phase Group B, the case is slightly more complicated. Implementation in Phase IIIA began in 2004, so that the program period for municipalities covered in that phase is 2004-2006. For the remaining municipalities, implementation began in 2006, so that the program period is 2006-2008. One difficulty comes from the fact that we do not know when implementation was scheduled to begin in the eligible municipalities on Group B that did not participate in the program. To deal with this issue, we assume that the non-participating municipalities in Group B would have been assigned to phases with the same probability as the participating municipalities. In our sample, out of the 15 participating municipalities in Group B, only 2 (13%) participated in phase IIIA, while the remaining 13 (87%) participated in phases IIIB and IV. We thus assume that the 6 non-participating but eligible municipalities would have participated in phase IIIA with a probability of 13% and in phases IIIB or IV with a probability of 87%. To these municipalities we therefore assign a value of 0.13 to the interaction of eligible and program in the years 2004/05 when implementation had only started in phase IIIA and a value of 0.87 for the years 2007/08 when implementation was only ongoing in phases IIIB and IV. For the year 2006, we assign a value 1 since implementation was ongoing in all participating municipalities in Group B.

Columns 1-3 of Table 2.5.2 report the results of negative binomial regressions with

Table 2.5: The Effect of KALAHI-CIDSS on Conflict Casualties

Outcome: Conflict Casualties (/Year)	Negative Binomial Regression			OLS
	(1)	(2)	(3)	(4)
Eligible \times Program	0.91*** (0.26)	0.66** (0.31)	0.93*** (0.30)	1.31*** (0.44)
Eligible \times Post-Program	0.45 (0.55)	-0.28 (0.46)	-0.09 (0.55)	-0.04 (0.56)
Eligible	0.08 (0.23)	0.07 (0.22)		
Population (/1000)		0.024*** (0.007)		
Area (km ²)		5.9* (3.1)		
Pct. Barangays with Highway Acc.		-0.40 (0.40)		
Timber		0.068 (0.065)		
Affected by NPA in 2001		1.05*** (0.27)		
Ethnic Fractionalization		1.76*** (0.41)		
Religious Fractionalization		-0.77 (0.67)		
Percent Muslim Population		1.93*** (0.51)		
Constant	-0.36 (0.23)	1.5 (5.9)		
Additional Controls	No	Yes	Yes	Yes
Municipality Fixed Effects	No	No	Yes	Yes
Observations	1285	1285	1285	1285

Robust standard errors in parentheses. Standard errors are clustered at the province level. The sample is restricted to municipalities within 2 ranks of the provincial eligibility threshold for KALAHI-CIDSS. All regressions control for the running variable (distance from threshold in ranks), fully interacted with eligibility and the project and post-project time-periods. All regressions include year fixed effects. Asterisks denote statistical significance at the 1% (***) 5% (**) and 10% (*) levels. Additional controls are the ten census variables used to determine eligibility for KALAHI-CIDSS. In the negative binomial regressions, the fixed effects refer to unconditional fixed effects.

conflict casualties as the dependent variable. The results suggest that eligible municipalities were significantly more likely to suffer conflict casualties in the period in which the program was implemented. The point estimate of the causal effect (the coefficient associated with the interaction of Eligible and Program) ranges from 0.66 to 0.93 in the negative binomial regressions. The effect is strongly statistically significant and robust to the inclusion of municipality fixed effects and clustering of standard errors at the province level. Since in the negative binomial regression the mean is an exponential function of the parameters, their size can be approximately interpreted as the effect of a unit change in the explanatory variable on a percentage change in the outcome. This means that, according to our preferred fixed effects specification, eligibility for KALAH-CIDSS caused a 90% increase in the number of conflict casualties, which is clearly a large effect relative to the baseline level of violence. In absolute terms, municipalities barely eligible for the program were likely to experience approximately 0.9 more conflict-related casualty per year than similar municipalities that narrowly missed the cutoff for eligibility. This means that, over the three-year program period, an eligible municipality experienced close to 3 additional casualties. Assuming a constant treatment effect, the 182 municipalities that received KALAH-CIDSS experienced almost 500 excess casualties. Given that leading datasets only require 25 annual battle-deaths for a violent dispute to be coded as a “civil conflict,” the size of the program’s effect is quite large. The results also show that, consistent with our theoretical model, the effect only persists as long as the project is being implemented. The point estimates of the effect in the post-program period (the coefficient associated with the interaction of Eligible and Post-Program) are much smaller than the effect during the program-period. In the models that include control variables and municipality fixed effects, the point estimate of the post-program effect is negative and close to zero. The effect is not statistically significant in any of the models. This finding is consistent with the predictions of our theoretical model, that violence only increases while the project is ongoing, since that is when insurgents can still hinder the project’s successful implementation. The results of the fixed effects linear regression in column 4 demonstrate that robustness of our estimation to different assumptions about the functional form.

2.5.3 Robustness tests

Balance on pre-treatment violence

The crucial identifying assumption of the Regression Discontinuity design is that municipalities on both sides of the eligibility threshold do not differ in unobserved variables that determine the intensity of conflict. This assumption might fail if the poverty mapping exercise that determined eligibility was manipulated in order to target the program to municipalities with higher or lower levels of pre-program conflict. Since there is some discretion about which census variables to use for the poverty mapping, it is possible that the variables were specifically chosen to make sure that a certain set of municipalities become eligible or ineligible. While we are not aware of any anecdotal evidence of manipulation of the poverty

mapping exercise, we present two pieces of evidence in support of the identifying assumption.

First, Figure 2.5.2 shows that eligible and ineligible municipalities experienced similar levels of conflict before eligibility for KALAH-CIDSS was announced. If eligible and ineligible municipalities differed in unobserved characteristics that determine conflict, we would expect them to experience different levels of violence before the program was announced, which does not appear to be the case. To the contrary, the fact that the increase in conflict in eligible municipalities coincides exactly with the start of the program’s roll-out suggests a causal effect of KALAH-CIDSS.

To quantitatively test the hypothesis that eligible and ineligible municipalities experienced equal numbers of casualties in the pre-KALAH-CIDSS period, we turn to the regression results in Table 2.5.2. In columns 1 and 2, the coefficient associated with “Eligible” is an estimate of the difference in annual casualties in eligible and ineligible municipalities before the start of the program. The estimates show that the difference is not statistically significant and close to zero. Thus, there is no evidence to suggest that eligible and ineligible municipalities differed on unobserved determinants of conflict prior to the start of the program, which suggests that our RD estimates measure the causal effect of the program on violence.

Robustness to choice of bandwidth

Table 2.6: Robustness Tests: Choice of Bandwidth

	Negative Binomial Regressions on Casualties (/year)					
	Bandwidth = 3 ranks			Bandwidth = 1 rank		
	(1)	(2)	(3)	(4)	(5)	(6)
Eligible \times Project	0.97*** (0.22)	0.91*** (0.30)	0.73** (0.29)	1.06*** (0.22)	1.07*** (0.29)	0.90*** (0.31)
Eligible \times Post-Project	0.41 (0.40)	0.19 (0.34)	-0.17 (0.43)	0.62 (0.41)	0.13 (0.38)	-0.09 (0.41)
Eligible	0.05 (0.17)	-0.08 (0.43)		-0.14 (0.27)	-0.17 (0.39)	
Constant	-0.48 (0.35)			0.15 (0.28)		
Controls	No	Yes	Yes	Yes	Yes	Yes
Municipality Fixed Effects	No	No	Yes	Yes	No	Yes
Observations	1865	1865	1865	657	657	657

Robust standard errors in parentheses. Standard errors are clustered at the province level. All regressions control for the running variable (distance from threshold in ranks), fully interacted with eligibility and the project and post-project time-periods. All regressions include year fixed effects. Asterisks denote statistical significance at the 1% (***) , 5% (**) and 10% (*) levels. Control variables are the same as in Table 2.5.2.

We now test the robustness of our results to the choice of bandwidth. Our baseline esti-

mates are based on a bandwidth of 2, meaning that the sample only included municipalities within two ranks of the provincial eligibility threshold (i.e. the richest two municipalities that were still eligible and the poorest two that were ineligible). Table 2.5.3 shows results of regressions based on bandwidths of 1 and 3. Overall, the estimates of the program's effect are very robust to changes in the bandwidth. The point estimates range between 0.73 and 1.07, which is comparable to the baseline results which ranged between 0.66 and 0.93, and are statistically significant at the 5% level. Thus, it does not appear that our estimates are strongly influenced by the choice of bandwidth.

Robustness to outliers

Table 2.7: Robustness Tests: Robustness to Outliers

	Probability of cas. >0		Number of incidents with cas. >0		
	(1)	(2)	(3)	(4)	(5)
Eligible × Program	0.123*	0.132**	0.64**	0.60***	0.37*
	(0.069)	(0.062)	(0.26)	(0.19)	(0.21)
Eligible × Post-Program	0.071	-0.050	0.44	0.23	-0.27
	(0.125)	(0.087)	(0.52)	(0.44)	(0.38)
Eligible	-0.019	-0.039	-0.19	-0.07	
	(0.096)	(0.11)	(0.49)	(0.45)	
Mean	0.29		0.49		
Controls	No	Yes	No	Yes	Yes
Municipality FE	No	No	No	No	Yes
Observations	1285	1285	1285	1285	1285

Robust standard errors in parentheses. Standard errors are clustered at the province level. The sample is restricted to municipalities within 2 ranks of the provincial eligibility threshold for KALAHICIDSS. All regressions control for the running variable (distance from threshold in ranks), fully interacted with eligibility and the project and post-project time-periods. All regressions include year fixed effects. In columns (1) and (2), reported values are marginal effects. Asterisks denote statistical significance of the underlying coefficient at the 1% (***), 5% (**) and 10% (*) levels. Control variables are the same as in previous tables.

One concern is that our results are driven by a small number of observations with very large numbers of casualties. To rule this out, the first two columns of Table 2.5.3 report the results of Probit regressions of the probability of having any casualties at all in a given year. The estimated marginal effect of KALAHICIDSS on the probability of having any conflict casualties is approximately 13 percentage points, which is large compared to the observed probability of 29%. To rule out that our results are driven by a small number of incidents with a very high number of casualties, the last three columns of Table 2.5.3 report the results of regressions that use the number of violent incidents (incidents with at least one casualty) as the dependent variable. The results show that eligibility for KALAHICIDSS increases

the number of violent incidents by between 37 and 64 percentage points, though the estimate in the fixed effects specification is only statistically significant at the 10% level. Overall, the results in Table 2.5.3 lead us to conclude that KALAH-CIDSS affects the probability as well as the intensity of conflict and that this result is not entirely driven by a small number of municipalities that experience severe conflict, or a small number of severe incidents.

2.5.4 The effect of project size on conflict casualties

Table 2.8: Effect of Project Size on Casualties

	Negative Binomial Regressions on Casualties (/year)		
	(1)	(2)	(3)
Eligible*Project	0.43 (0.59)	0.78 (0.83)	0.85 (0.96)
Eligible * Project * # of villages	0.041** (0.019)	0.047** (0.019)	0.047** (0.019)
Eligible * Project * area	2.80 (3.10)	8.0 (11.5)	5.5 (19.2)
Eligible * Project * area squared		-24.4 (56.7)	-15.4 (77.4)
Eligible * Project * population	-0.019 (0.012)	-0.052 (0.040)	0.049 (0.045)
Eligible * Project * pop. squared		3.5×10^{-7} (3.6×10^{-7})	3.3×10^{-7} (3.9×10^{-7})
Eligible * Project * pop. density			-5.3×10^{-5} (3.3×10^{-4})
Eligible*Post-Project	-0.24 (0.40)	-0.22 (0.40)	-0.17 (0.43)
Municipality Fixed Effects	Yes	Yes	Yes
Observations	1285	1285	1285

Robust standard errors in parentheses. Standard errors are clustered at the province level. All regressions control for the running variable (distance from threshold in ranks), fully interacted with eligibility and the project and post-project time-periods. All regressions include year fixed effects. Asterisks denote statistical significance at the 1% (***), 5% (**) and 10% (*) levels. Control variables are the same as in Table 2.5.2.

We now analyze the relationship between a project's size and its effect on conflict. To do this, we exploit the fact that the amount of aid an eligible municipality received was a function of the number of villages (called *barangays* in Tagalog) it contains²⁷. As mentioned in Section 3, the amount of aid an eligible municipality received from KALAH-CIDSS was

²⁷Barangays are the smallest administrative unit in the Philippines, with an average of 25 barangays per municipality

determined by multiplying its number of villages by PhP300,000 (about US\$6000). This means that municipalities with more villages received larger amounts of aid. Assuming that larger amounts of aid cause a larger (potential) shift in power between the government and insurgents, our model predicts that eligible municipalities with many villages will experience a larger increase in conflict than municipalities with few villages. The regressions reported in Table 2.5.4 test this hypothesis by including an interaction between eligibility, the program time-period and the number of villages in the municipality. Of course, the number of villages is likely to be correlated with both area and population, which might also affect the size of the program's effect on conflict. To control for this, we also include interactions between eligibility, the program time-period and the municipality's area and population (both linear and squared), as well as its population density. The results suggest that holding area and population constant, municipalities with a larger number of villages experienced a larger increase in conflict from being eligible for KALAH-CIDSS. The point estimates are in the range of 0.04 to 0.05 suggest that having an additional village - which increased the grant size by PhP300,000, or US\$6000 - increased the project's effect on casualties by approximately 4 to 5 percentage points. This result suggests that larger grants caused more conflict, which is consistent with our model's prediction that conflict is more likely when the (potential) shift in power between government's and insurgents is large.

2.5.5 Who suffers and who initiates the violence?

Table 2.5.5 reports estimated effects of KALAH-CIDSS on casualties from the three groups and civilians. The results show that all groups suffered more casualties in KALAH-eligible municipalities than in the similar ineligible municipalities. The largest effect is on insurgents, who suffered an approximate increase of 113% from a baseline of 0.31 casualties. At approximately 63%, the estimated effect is smaller for government troops. However, government troops suffer a higher number of casualties on average, so that the absolute effect is almost as large as for insurgent casualties. Civilians appear to suffer fewer casualties, both overall and as a result of the project. These results are consistent with our model, in which insurgents and government forces engage in conflict over the division of the surplus from the program and civilians are not directly involved.

The results in Table 9 also show which group, government or insurgents, initiates the violence caused by KALAH-CIDSS. Our theoretical model of bargaining and conflict does not make predictions about this - it simply states that conflict occurs if both parties cannot agree on a peaceful bargaining solution and remains agnostic about who initiates the violence. Nevertheless, knowing who initiates the violence may yield insights about whether KALAH-CIDSS gives one party the initiative in the ensuing conflict. The results show that the program causes an increase in violence originating from both groups. The increase of insurgent-initiated violence is slightly larger than that of government-initiated violence, but the difference is fairly small. Overall, the results suggest that the violence around the KALAH-CIDSS program was not the result of a one-sided offensive by either the government or the insurgents

Table 2.9: Who Suffers and Who Initiates the Violence? Conflict Casualties by Actor

Outcome:	Parameters of	
	Neg. Binom. Regression	Mean
	(1)	(2)
	0.61	0.49
	(0.39)	(0.05)
Insurgent Casualties	1.13**	0.31
	(0.45)	(0.03)
Civilian Casualties	0.57	0.20
	(0.51)	(0.02)
Cas. in AFP initiated inc.	0.75	0.40
	(0.58)	(0.04)
Cas. in insurgent initiated inc.	0.93**	0.56
	(0.40)	(0.05)
Control Variables	Yes	
Municipality Fixed Effects	Yes	
Observations	1285	1285

Results in column 1 are parameter estimates associated with eligible*project in a fixed effects negative binomial regression. Robust standard errors in parenthesis, clustered at the province level. Results in column 2 are sample means. The sample is restricted to municipalities within 2 ranks of the provincial eligibility threshold for KALAH-CIDSS. All regressions control for the running variable (distance from threshold in ranks), fully interacted with eligibility and the project and post-project time-periods. Asterisks denote statistical significance at the 1% (***), 5% (**) and 10% (*) levels.

2.6 Conclusion

In recent years, donors and governments have targeted an increasing amount of development aid to areas affected by civil conflict, some of it in the hope that aid will reduce conflict by weakening popular support for insurgent movements. This paper has presented a simple mechanism through which development aid can have the unintended effect of *increasing* conflict: If a successful project will weaken the position of insurgent groups in the future, they have an incentive to oppose it, which may exacerbate conflict. To analyze this mechanism, we have developed a theoretical model of bargaining and conflict in the context of development projects. The model predicts that development projects can cause conflict if (1) a successful project changes the future balance of power in favor of the government, (2) the insurgents have the ability to hinder the project's successful implementation by violent means, and (3) governments cannot commit to honoring agreements reached before the start of the project. The first two conditions ensure that the insurgents have an incentive to use violence to hinder the project's implementation, the third condition ensures that governments cannot pay off insurgents in return for allowing the project's peaceful implementation.

Our empirical analysis tests the model by estimating the causal effect of a large development program - called KALAH-CIDSS - on conflict casualties in the Philippines. During the period 2003-08, KALAH-CIDSS was the Philippines' flagship anti-poverty program with a budget of \$180 million, financed through a loan from the World Bank. To overcome the problem of endogenous targeting of aid, we employ a regression discontinuity design that compares municipalities just above and just below the poverty threshold that determined eligibility for the program. Using detailed data on all conflict incidents involving the Armed Forces of the Philippines between 2001 and 2008, our estimates show that eligibility for KALAH-CIDSS caused a large and statistically significant increase in conflict casualties. Consistent with the predictions of our model, this effect only persists for the duration of the program - while insurgents can still hinder its implementation - and is stronger for municipalities that received larger amounts of aid. We further find that the majority of casualties was suffered by insurgents and government troops, while civilians appear to have suffered less. Eligible municipalities experienced a similar increase in the number of casualties in insurgent-initiated and government-initiated attacks, suggesting that the effect is not due to a one-sided offensive by either party.

Our results have implications for future research and policy. First, they highlight the potential pitfalls of extrapolating from the effect of natural experiments (in the sense of largely uncontrollable phenomena like rainfall and world-market prices) to the effect of local human interventions. Since conflict results from a strategic interaction between (at least) two parties, interventions that can be influenced by either party can have very different effects from interventions that are truly exogenous. For example, the recent literature suggests that shocks to world-market prices of agricultural goods reduce conflict by increasing the population's return to peaceful activities (Dube & Vargas 2007). Based on this finding, it is tempting to conclude that development projects that increase people's economic opportunities will have a similarly conflict-reducing effect. However, while insurgents cannot

affect world-market prices, they *can* use violence to sabotage development projects. Development projects can therefore cause an increase in conflict violence; and our theoretical model outlines the conditions under which they are most likely to do so.

How this insight affects the targeting and design of development projects depends on the projects' goals. If a project's main goal is to reduce poverty and alleviate the suffering of populations in conflict-affected areas, one way of avoiding conflict is to make sure that the project does not affect the balance of power between governments and insurgents. One possible way of doing this is to cooperate with both governments and insurgents in designing the project and delivering the aid. (an example of this is the recent collaboration of Japan's development agency JICA with the MILF in extending aid to parts of Mindanao in the southern Philippines). If the project's goal is to reduce conflict by weakening insurgents, one way of reducing violence is to focus aid on a smaller number of projects but heavily defending these. This would ensure that insurgents have less ability to sabotage project (and face higher costs if they do), which should help deter violent attacks. To ensure that projects are implemented successfully, it may also be desirable to weaken insurgent capacity before the start of the project by military means, following a "clear, hold, build" strategy (US Army2007). Of course, these policy conclusions are speculative, since they are derived from theory and have not been tested empirically. We hope that future research will be able to test these implications of our model and will help design development interventions that can operate in conflict-affected areas without exacerbating violent conflict.

Chapter 3

The Effect of Subsidized Employment on Happiness

Abstract

While a large body of evidence suggests that unemployment and self-reported happiness are negatively correlated, it is not clear whether this reflects a causal effect of unemployment on happiness and whether subsidized employment can increase the happiness of the unemployed. To close this gap, this paper estimates the causal effect of a type of subsidized employment projects - Germany's *Arbeitsbeschaffungsmaßnahmen* - on self-reported happiness. Results from matching and fixed effects estimators suggest that subsidized employment has a large and statistically significant positive effect on the happiness of individuals who would otherwise have been unemployed. Detailed panel data on pre- and post-project happiness suggests that this effect can neither be explained by self-selection of happier individuals into employment nor by the higher incomes of the employed.

3.1 Introduction

A large body of research shows that the unemployed report significantly lower levels of average happiness¹ and higher levels of psychological distress than the employed (see McKee-Ryan, 2005, for a review of the psychological literature). The negative correlation between unemployment and happiness, both across individuals and over time, remains significant after controlling for a wide range of observable characteristics, including income (Clark and Oswald 1994, Winkelmann and Winkelmann 1998, Marks and Fleming 1999, Clark 2003, Carroll 2007). A possible explanation for this finding is that, in addition to income, jobs confer social status, respect and a sense of purpose, competence and efficacy, all of which

¹For the rest of the paper, I use the term “happiness” as meaning “self-reported happiness”. Following Arrow and Dasgupta (2009), I use happiness as synonymous with life-satisfaction and well-being.

are thought to be important contributors to well-being and job-satisfaction (Izard 1991, Ryan and Deci 2000, Ellingsen and Johannesson 2007, Ariely et al. 2008). Involuntary unemployment² may therefore have a psychological cost - a negative effect on well-being that goes beyond its effects on income and consumption (Frey and Stutzer 2002, Carroll 2007). A psychological cost of unemployment would have implications for labor market and welfare policy, implying that the welfare cost of unemployment is greater than the value of lost output and that subsidized employment may be a better way to increase the well-being of the unemployed than direct cash transfers (see, for example, Edlin and Phelps 2009, who cite the psychological benefits of employment as an argument for the introduction of tax credits for employers of low-wage workers).

But the evidence for a negative causal effect of unemployment on happiness is not entirely conclusive. Happiness and unemployment are simultaneously determined, so it is possible that unobserved shocks - for example adverse shocks to (mental) health - simultaneously decrease happiness and increase the probability of becoming or remaining unemployed (e.g. Mastekaasa 1996). There is some evidence that a similar mechanism may explain the negative correlation between self-reported health and unemployment. While the average unemployed person reports a lower health status than the average employed person, individuals who lose their jobs for exogenous reasons, such as the closure of their employer's business, do not experience a decline in health status (Salm 2009). Causality is therefore likely to run from bad health to unemployment and not in the other direction (Bockermann & Ilmakunnas 2009). If the same is true for happiness, the unemployed may be less happy than the employed even if unemployment has no causal effect on happiness.

Even if unemployment causes unhappiness, it is not clear that subsidized employment can increase happiness. It is for example possible that people's happiness is only increased by jobs that have certain desirable characteristics, such as being perceived as meaningful or conferring high social status and respect (Ellingsen 2007, Ariely 2008). Since the jobs created by subsidized employment often have low pay and low social status, it is possible that they do not have the desirable characteristics that cause an increase in happiness. In other words, even if the *average* job increases the happiness of the average employee, the *marginal* job created by an employment subsidy may have no (or even a negative) effect on the happiness of the marginal employee.

This paper contributes to the literature on happiness and unemployment by estimating the effect of subsidized employment on the happiness of the unemployed. To do this I analyze the happiness of participants in a type of public subsidized employment projects (SEPs) - Germany's *Arbeitsbeschaffungsmaßnahmen*. Previous research suggests that these projects have on average had little success in increasing participants' future income and probability of employment (Hujer et al. 2004, Caliendo et al. 2008). But if the goal of public policy is to increase people's happiness, employment subsidies may still be desirable if they prevent the unhappiness of unemployment.

Since participation in the subsidized employment projects is non-random, my identifi-

²For the rest of the paper, I will use the term "unemployment" as meaning "involuntary unemployment".

cation strategy relies on detailed panel data on happiness before and after the start of the project. The data come from the German Socio-Economic Panel which, among other things, collects information on respondents' happiness and employment status, including participation in subsidized employment projects. Using this data I show that, for the duration of the subsidized employment project, the happiness of participants is significantly higher than that of unemployed non-participants with similar observable characteristics. The data further show that participants and similar non-participants have virtually identical levels and trends of happiness in the months before the start of the project, which suggests that the observed effect is not driven by self-selection of happier individuals into the projects. Quantitative estimates from fixed-effects and nearest-neighbor matching estimators suggest that, compared to the counterfactual of remaining unemployed, subsidized employment increases happiness by about 0.4 to 0.6 points on a scale from 0-10. This effect corresponds to about 0.4 within-individual standard deviations of happiness, which is large compared to the effects of other observable characteristics like income and marital status.

It should be noted that this estimate does not reflect the total effect of subsidized employment programs on the happiness of their participants, but only their effect against the counterfactual of remaining unemployed. Some participants would have been employed even without the SEP, so that the overall effect of participating in the project (against the counterfactual of not participating) is most likely smaller. Still, the estimated effect can be useful for evaluating the effect of employment subsidies on happiness. Regardless of whether some participants would have found jobs in the absence of the SEP, economic theory suggests that an employment subsidy creates jobs in equilibrium. One could therefore combine this paper's estimate of the effect of subsidized employment on happiness with an estimate of the number of jobs created by a subsidy in equilibrium to derive an estimate of the subsidy's total effect on happiness.

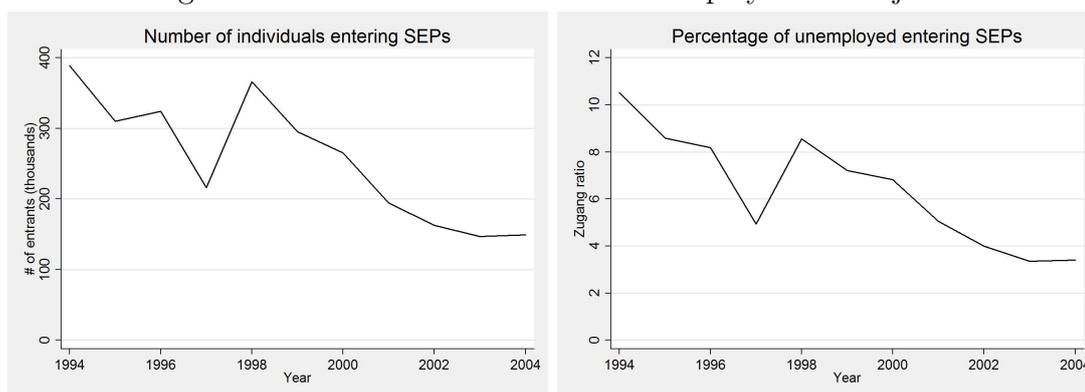
As an additional contribution, I attempt to disentangle whether subsidized employment increases happiness by conferring direct psychological benefits - for example through feelings of competence and efficacy - or by increasing income and consumption. In order to draw policy conclusions it is important to disentangle these two channels. A happiness-based argument for publicly subsidized jobs, similar to the one made by Edlin and Phelps (2009), only holds if employment itself increases happiness. If the employed are merely happier because of their higher incomes, direct income transfers are most likely a more cost-effective way of increasing the happiness of the unemployed. To disentangle the two channels, I exploit the fact that participation in a SEP prolonged individuals' entitlement to public unemployment benefits. Thus the projects' positive effect on income remained even after employment in it had ended. This creates sufficient independent variation in income and employment to allow me to identify the effect of employment while controlling for differences in income. Intuitively, if the positive effect of SEPs on happiness were mainly due to their effect on income, we would expect participants' happiness to remain high as long as the project's effect on income persists, even after employment has ended. But the data show that participants' happiness decreases substantially as employment in the project ends, suggesting that employment has psychological benefits that are independent of its effect on income. To

obtain quantitative estimates of the net effect of employment on happiness (excluding the effect of income), I estimate a fixed effects instrumental variables estimator that exploits the fact that participants' probability of employment drops sharply at the end of the subsidized employment project while their incomes remain nearly unchanged.

The next section briefly describes the institutional details of Germany's subsidized employment projects. Section 3 describes the econometric methods used to identify the causal effect of subsidized employment on happiness. Sections 4 and 5 describe the data and present results, Section 6 concludes.

3.2 Institutional Background: Subsidized Employment Projects in Germany

Figure 3.1: Prevalence of Subsidized Employment Projects



Source: Bundesagentur für Arbeit

Subsidized employment projects (*Arbeitsbeschaffungsmaßnahmen*, SEPs) have been an integral part of Germany's active labor market policy for over 30 years (Bernhard et al. 2008). Figure 3.1 plots the trend of entrants into SEPs over the period of observation, 1994-2004, both in absolute numbers and as a percentage of the unemployed³. At the peak in 1994, approximately 390,000 individuals - slightly over 10% of all unemployed individuals - entered a SEP within a single year. The average annual number of entrants is about 280,000, corresponding to approximately 6% of unemployed individuals.

The institutional rules governing subsidized employment projects have been described in detail by Hujer et al. (2004) and Caliendo et al. (2008) and this section draws on their

³The "Hartz-IV" labor market reforms, which came into force in January 2005, introduced substantial changes to the system of subsidized employment in Germany. *Arbeitsbeschaffungsmaßnahmen* lost in importance and were largely replaced with so-called One-Euro-Jobs. In order to keep the results consistent, I therefore only focus on the period until 2004

descriptions. The two main instruments of German active labor market policy are vocational training and subsidized employment projects. Local job-centers have a large amount of autonomy in allocating their budget to different policies, but subsidized employment appears to be the favored instrument in areas with higher unemployment. To create a subsidized employment project, the potential employer applies to the job-center with a description of the proposed jobs. The job-center assesses the proposal according to a number of criteria, most importantly whether the proposed activity is in the public interest and whether the activity is “additional” in nature, meaning that it would not be undertaken in the absence of a SEP. Up to 2002, SEPs were reserved for employers in the non-profit sector, after 2002 exceptions became possible with the approval of the job-center. After approving the SEP, the job-center assigns some of its unemployed clients to the project and pays between 30 and 75 % of their wages, though in exceptional cases the amount of the subsidy can go up to 100%. In the assignment decision, job-centers are mandated to give priority to individuals whose chances of employment outside of SEPs are small.

Before 2002, participants in SEPs had to be unemployed for at least 6 out of the previous 12 months, though exceptions existed for young people without professional training, the short-term unemployed and people with disabilities. In addition, 5 % of the places in SEPs could be allocated to individuals who did not meet any of these conditions. After 2002, all unemployed individuals could be assigned to SEPs, under the condition that the job-center saw the SEP as their only opportunity for employment. Individuals can refuse to participate in the project, but refusal can be penalized by a reduction in unemployment benefits. The duration of a project is usually 12 months but projects can be extended in special cases.

3.3 Empirical Strategy

3.3.1 A simple model of happiness and (un)employment

This section presents a simple empirical model of happiness and (un)employment. It assumes that the happiness of individual i while being unemployed at time t is a function of her characteristics (X_{it}) at the time,

$$h_{it}(0) = f(X_{it})$$

while the individual’s happiness while being employed in job k is also a function of the job’s characteristics (Z_{kt}),

$$h_{it}(k) = g(X_{it}, Z_{kt})$$

The goal of this paper is to estimate the expected difference in $h_{it}(k)$ and $h_{it}(0)$ for participant/job pairs created by subsidized employment programs (SEPs),

$$\tau = E_{X_{it}, Z_{kt}}[h_{it}(k) - h_{it}(0)], \quad (i, k, t) \in S$$

where $(i, k, t) \in S$ implies that at time t individual i participated in a SEP, through which she was employed in job k . The parameter τ is the expected gain in happiness the average participant in a SEP experiences at a given point in time from being employed in the SEP rather than being unemployed.

As mentioned in the introduction, this is of course not the total effect of subsidized employment programs on the happiness of their participants. First, some participants would have been employed even without the SEP and, second, SEPs may affect future happiness by changing the probability and characteristics of future employment. Still, τ is useful for evaluating the effect of employment subsidies on happiness. Regardless of whether some participants would have found jobs without the SEP, economic theory suggests that an employment subsidy creates jobs in equilibrium. Assuming that the characteristics of the job/employee pairs created by a subsidy in equilibrium are the same as the characteristics of the job/employee pairs in SEPs, τ yields the effect of the average job created by the subsidy. This could be combined with an estimate of the number of jobs created by an employment subsidy to yield an estimate of the subsidy's aggregate effect of happiness.

It should also be noted that τ is an "average effect of treatment on the treated", since it measures the effect of employment in SEPs on individuals who participate in them. Thus τ is the expected effect of employment in SEPs on individuals who are involuntarily unemployed - meaning those who are willing to accept a low-paying job in a SEP - and not the effect on the average person in the population.

3.3.2 A Matching Estimator for the Effect of Subsidized Employment on Happiness

To estimate τ , I use the nearest neighbor matching estimator described by Abadie and Imbens (2002). Since we can observe individuals' happiness while employed in a SEP, $h_{it}(k)$, the matching estimator only needs to estimate their counterfactual happiness while unemployed, $h_{it}(0)$. This is imputed from the outcomes of matched unemployed non-participants with similar observed characteristics:

$$\widehat{h_{it}(0)} = \frac{1}{M} \sum_{(j) \in J_M(i,t)} h_{jt}(0)$$

In this notation $J_M(i, t)$ is the set of matched control observations associated with participant i at time t . Matched controls are selected so that their observed characteristics in the pre-treatment period X_{jt-1} are as similar as possible to the observed characteristics of the participant in the pre-treatment period X_{it-1} ⁴. More precisely, $J_M(i, t)$ is defined as containing the M observations with the smallest distance between X_{jt-1} and X_{it-1} , using a suitable metric, so that observations are matched to their nearest neighbors in the space

⁴Observations are matched on characteristics in the pre-treatment period in order to avoid that participants' characteristics are already affected by the treatment

of observed characteristics. For this paper, I use the standard distance metric $(X_{jt-1} - X_{it-1})' \Sigma^{-1} (X_{jt-1} - X_{it-1})$, where Σ is the covariance matrix of X .

Subsidized employment projects usually last for 12 months, which is the same as the average interval between two interviews for the German Socio-Economic Panel. The majority of participants is therefore observed only once per employment spell in a SEP⁵.

For the baseline estimates, observations are matched on 11 variables: sex, age, years of education, marital status, household size, number of children, unemployment status, household income, income from public unemployment benefits, region⁶ and month of interview. In an extended specification, observations are also matched on pre-treatment happiness in order to control for unobserved determinants of happiness.

Testing the conditional independence assumption

The matching estimator's main identifying assumption is that, conditional on the pre-treatment variables used for matching, X_{it-1} , the counterfactual outcome $h_{it}(0)$ is independent of participation in a SEP. This assumption ensures that the actual outcome of the unemployed matched non-participants $h_{jt}(0)$ is a consistent estimator of the counterfactual outcome of the participants under unemployment, $h_{it}(0)$. It implies that participants in subsidized employment projects would have been as (un)happy being unemployed as the matched non-participants who actually were unemployed.

There are two reasons why this assumption might be violated. First, happier people may be more likely to participate in SEPs, so that participants may have been more happy than non-participants even in the absence of the project. If this were the case, $h_{it}(0)$ would be greater than $h_{jt}(0)$ and the matching estimator of τ would be biased upward. Fortunately, the panel nature of the data allows me to test for this violation by comparing the pre-treatment happiness of participants, h_{it-1} , to the happiness of matched controls in the pre-treatment period, h_{jt-1} . If happier individuals self-select into the project we would expect participants to already be happier than matched controls in the pre-treatment observation, so that h_{it-1} would be greater than h_{jt-1} . On the other hand, observing that h_{it-1} is equal to h_{jt-1} , even for observations close to the the start of the SEP, should increase our confidence that happier individuals do not self-select into the projects.

Second, since participation in SEPs is voluntary⁷, people may self-select into the projects according to how much they benefit from them. Participants and matched controls may therefore differ in how strongly their happiness is affected by unemployment, so that $h_{it}(0)$ may be different from $h_{jt}(0)$ even if participants and controls were equally happy when being employed in the pre-treatment observation. To test for this, I compare the

⁵In some cases SEPs are extended beyond 12 months so that we observe participants more than once during the project. To avoid problems from endogenous duration of employment, the matching estimator only uses participants' first observation during a SEP. Results that use all observations within a SEP are not reported, but are similar to the reported ones.

⁶Western or Eastern Germany

⁷Though repeated refusal to participate can lead to sanctions by the job-center.

pre-treatment happiness of participants and matched controls who were unemployed in the pre-treatment observation - $h_{it-1}(0)$ and $h_{jt-1}(0)$. Finding that participants and matched controls report different levels of happiness when unemployed, or that their happiness during unemployment follows different trends, would indicate that the groups are differently affected by unemployment and that the conditional independence assumption is violated. Finding no difference in pre-treatment levels and trends of happiness between unemployed participants and matched controls should increase our confidence that both groups are equally affected by unemployment and that the conditional independence assumption holds.

3.3.3 Disentangling the effects of income and employment

In principle, there are two ways in which employment in a SEP might affect happiness: by conferring direct psychological benefits - for example feelings of competence and efficacy - and by increasing individuals' incomes. To inform policy, it is useful to disentangle these two channels. A happiness-based argument for publicly subsidized jobs, similar to the one made by Edlin and Phelps (2009), only holds if employment itself increases happiness. If the employed are merely happier because of their higher incomes, direct income transfers are likely to be a more cost-effective way of increasing the happiness of the unemployed.

I therefore present an estimator for the "pure" effect of employment on happiness, net of the effect of increased incomes. Slightly modifying the notation of the previous section, I define $h_{it}(0, k)$ as the happiness that individual i reports at time t if she is unemployed but her income is as high as if she were employed in job k . Using this notation, the net effect of subsidized employment on happiness can be written as

$$\theta = E_{X_{it}, Z_{kt}}[h_{it}(k) - h_{it}(0, k)], (i, k) \in S$$

where, as before, S is the set of job/employee pairs created through SEPs. Unfortunately, θ is not easily identified without additional assumptions. Comparing participants and non-participants with similar post-treatment levels of income - either by matching on post-treatment income, or controlling for it in a regression - would not cleanly identify the effect. Since participation in a SEP has a positive effect on wage income, participants and non-participants can only have identical incomes if they differ in unobserved variables. Comparing participants and matched controls with similar incomes therefore risks introducing omitted variable bias (see, for example, Gelman and Hill 2007, pp 188-194).

To allow identification of θ , I assume that the effect of income (Y_{it}) follows a logarithmic functional form and is linearly separable from the effects of individual and job characteristics (X_{it} and Z_{it}). Thus, the happiness of employed and unemployed individuals is given by:

$$h_{it}(k) = g(X_{it}, Z_{kt}) + \log(Y_{it})\gamma + u_{it}$$

$$h_{it}(0) = f(X_{it}) + \log(Y_{it})\gamma + u_{it}$$

so that θ can be written as:

$$\theta = E_{X_{it}, Z_{kt}} [g(X_{it}, Z_{kt}) - f(X_{it})], (i, k) \in S$$

where X_{it} and Z_{it} now exclude income.

I estimate θ in two ways. First, I estimate a fixed effects regression of happiness that includes an indicator for being employed in a SEP and controls for income. The estimated equation is

$$h_{it} = \delta_1 D_{it}^{reg} + \delta_2 D_{it}^{SEP} + X_{it}\beta + \log(Y_{it})\gamma + \alpha_i + u_{it}$$

where D^{reg} and D^{SEP} are indicators for being employed in a regular job and in a SEP. Under the identifying assumption that u_{it} is uncorrelated with employment in a SEP, $\widehat{\delta}_2$ is an unbiased estimate of θ . However, there are several reasons why this assumption may be violated. First, entry into and exit from SEPs is non-random, so that unobserved shocks may be correlated with employment in a SEP. This concern is similar to the one that was previously discussed in the context of the matching estimator for the aggregate effect of employment in a SEP. A concern that is specific to estimating the net effect of employment - excluding the effect of income and consumption - is that entry into subsidized employment may increase individuals' expectations of future income. This could lead individuals to increase their consumption as they enter a SEP, which may positively affect their happiness (alternatively, expected future income might have a direct effect if individuals receive happiness from anticipating future income). Thus, entry into a SEP may be correlated with unobserved shocks to expected future income and consumption, which would bias the estimate of θ .

As a robustness test, I graphically examine the happiness of SEP participants at the end of the project. SEPs usually last for one year, and most participants go back into unemployment when they exit the project. Thus, one year after the start of the project, there is a sharp drop in participants' probability of employment. But their incomes do not immediately decrease since participation in a SEP extends their entitlement public unemployment benefits. Also, while participants' expectations of future income may increase as they enter a SEP, it is unlikely that their expectations decrease discontinuously exactly one year after the start of the project (since the duration of the project is known in advance). Thus, if employment affects happiness independently of income, happiness should drop one year after the start of a SEP, as employment ends while current and expected future income remain unchanged (or at least do not change discontinuously). If, on the other hand, the effect of SEPs on happiness is only due to their effect on income, we would not expect a drop in happiness one year after the start of the project.

In addition to the graphical test, I calculate a fixed effects instrumental variables estimator. This estimator exploits the fact that participant's probability of employment drops substantially one year after the start of a SEP while their incomes do not decline immediately. For this regression I use only observations of SEP participants after the start of a SEP in order to avoid bias from endogenous entry and from shocks to expected future income that may affect happiness at the start of a SEP. To avoid bias from endogenous exit from SEPs, I instrument employment by an indicator for an individual's first observation after entering

a SEP. The first-stage relationship between this instrument and employment is created by the fact that the usual SEP lasts for one year. This is the same as the average interval between two observations in the German Socio-Economic Panel, so that the probability of employment drops significantly between the first and second observation after entering a SEP. The exclusion restriction rests on the assumption that SEP participants experience no systematic unobserved shocks between their first and second observation after entering a SEP. Additional robustness tests for this assumption are discussed in more detail in Section 5, together with the results.

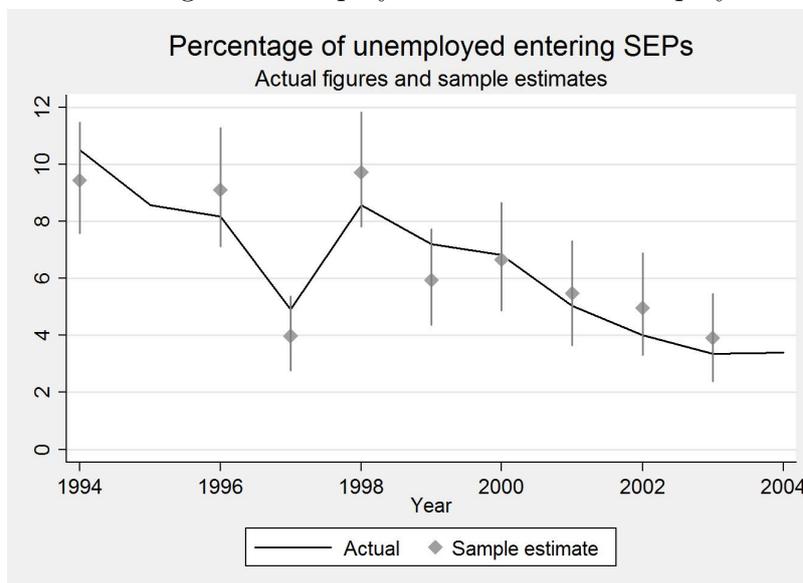
3.4 Data and Summary Statistics

The empirical analysis in this paper uses data from the German Socio-Economic Panel, from the years 1992 to 2004. The sample is restricted to respondents between the ages of 18 and 65. The outcome of interest is respondents' self-reported happiness measured by their answer to the question: "All things considered, on a scale from 0 to 10, how satisfied are you with your life?"⁸. Answers to questions of this type correlate well with more detailed measures of psychological distress (Koivumaa-Honkanen et al. 2004) and physiological indicators of well-being such as blood-pressure (Blanchflower & Oswald 2008). They also predict suicide risk and mortality (e.g. Koivumaa-Honkanen 2001, Chida 2008). The explanatory variable of interest is participation in subsidized employment projects (SEPs). From 1992 onwards, the GSOEP collected information on whether respondents were currently employed in a SEP. Figure 3.2 shows that the sample estimate of the fraction of unemployed individuals who participate in SEPs closely follows the actual time of participation.

Table 3.1 reports summary statistics of the whole sample, of the unemployed and of individuals who participate in subsidized employment projects. For SEP participants, the table reports summary statistics in the observation before the project began, so that the results do not measure the effect of participation itself. To be comparable, the summary statistics for the unemployed are based on the lagged observation in which individuals may still have been employed. Columns 4 and 5 show differences in means between participants in SEPs and non-participants as well as between participants and the unemployed. Compared to the population as a whole, SEP participants live in larger households with lower incomes, are more likely to be female and have a steady partner and report lower levels of happiness. Compared to the unemployed, SEP participants are also younger and better educated. Clearly participation in SEPs is not random, even conditional on being unemployed, so that we should expect participants and non-participants to differ in observed as well as unobserved characteristics.

⁸As mentioned in the introduction, I follow Arrow and Dasgupta (2009) in using the term happiness as synonymous with life-satisfaction. I do this to make the text more readable: saying that employment makes people happy is a briefer way of saying that employment makes people more satisfied with their lives.

Figure 3.2: Percentage of Unemployed in Subsidized Employment Projects



Data source: German Socio-Economic Panel, 1992-2004. Start dates are based on retrospective reports of individuals that have started a SEP since the previous observation. In 1996, respondents were not asked about SEP participation, so that the estimate for the previous year, 1995, is missing.

Table 3.1: Summary Statistics

	Means (in previous observation)			Differences	
	Whole Sample	Unemployed	SEP Participants	Participants to Non-Participants	Participants to Unemployed
Female	0.510 (0.500)	0.521 (0.500)	0.591 (0.492)	0.082 [0.029]**	0.070 [0.027]**
Age	40.6 (12.7)	42.1 (12.4)	40.5 (10.6)	-0.08 [0.60]*	1.65 [0.58]
Steady Partner	0.642 (0.479)	0.602 (0.491)	0.656 (0.476)	0.014 [0.027]	0.055 [0.025]**
Household Size	2.97 (1.20)	2.87 (1.23)	3.19 (1.26)	0.22 [0.07]***	0.312 [0.068]***
HH Income (Euros/month)	2762 (1513)	2112 (1092)	2012 (945)	-753 [52]***	-99.9 [50.5]**
Years of education	11.9 (2.4)	11.2 (1.9)	11.7 (2.0)	-0.17 [0.13]	0.46 [0.12]***
Self-reported happiness	6.86 (1.75)	5.89 (2.00)	5.44 (2.06)	-1.43 [0.11]***	0.44 [0.11]***
Observations	90185	6236	413	90185	6649
Individuals	11366	2605	329	11366	2649

Data source: German Socio-Economic Panel, 1992-2004. Values for SEP participants are from the pre-treatment observation. For comparison, values for the whole sample and the unemployed are from the lagged observation. Standard deviations in parentheses. Standard errors of differences in brackets. *, ** and *** denote statistical significance at the 10, 5 and 1 percent levels.

3.5 Results

3.5.1 The effect of subsidized employment on happiness

Table 3.2: Effect of Subsidized Employment Projects on Happiness: Matching Estimators

	All participants		Unemployed in pre-treatment obs.	
Effects on:	(1)	(2)	(3)	(4)
Happiness	0.485 (0.112)***	0.543 (0.0101)***	0.389 (0.136)***	0.390 (0.119)***
Happiness (bias adjusted)	0.531 (0.112)***	0.620 (0.098)***	0.430 (0.136)***	0.434 (0.113)***
Pre-Treatment Differences:				
Happiness	-0.119 (0.121)	-0.125 (0.053)**	-0.031 (0.141)	-0.032 (0.061)
Happiness (bias adjusted)	-0.092 (0.120)	-	-0.008 (0.140)	-
Matched on happiness in pre-treatment obs.	No	Yes	No	Yes
Number of SEP spells	413	413	296	296

Data source: GSOEP, 1992-2004. Standard errors in parenthesis. ***, ** and * denote statistical significance at the 1, 5 and 10 % level. Estimates are based on individuals' first observation in an employment-spell in a SEP and 3 matched observations. For the baseline matching, observations are matched on: sex, age, years of education, relationship status, household size, number of children, household income, unemployment status, household income from unemployment benefits, region (Western/Eastern Germany) and month of interview. To avoid reverse causality, observations are matched on values in the pre-SEP observation.

Table 3.5.1 reports results of matching estimators of τ , the average effect of subsidized employment projects (SEPs) on the happiness of their participants (which is the average effect of treatment on the treated). The matching procedure is described in detail in Section 3.1. Columns 1 and 2 of Table 3.5.1 report results from the whole sample of participants, while columns 3 and 4 report results from the sub-sample of participants who were unemployed in the pre-treatment observation.

For the baseline estimates, presented in columns 1 and 3, observations are matched on 11 pre-treatment variables: sex, age, years of education, marital status, household size, number of children, unemployment status, household income, income from public unemployment benefits, region⁹ and month of interview. In addition, the estimators presented in columns 2 and 4 match on pre-treatment happiness in order to control for unobserved heterogeneity in factors that affect individuals' happiness. The first row reports the simple nearest neighbor matching estimate, the second row reports the estimate after correcting for potential bias from remaining differences in the control variables.

The results in Table 3.5.1 suggest that employment in SEPs has a large and statistically significant effect on participants' happiness falling in the range between 0.39 and 0.62 on the 0 to 10 scale - equivalent to between 0.3 and 0.5 within-individual standard deviations of self-reported happiness. The simple nearest neighbor estimates do not differ much from the bias adjusted estimates, which suggests that the matching procedure succeeded in selecting

⁹Western or Eastern Germany

controls whose observed characteristics are similar to those of the participants they were matched to.

By looking at the pre-treatment differences in happiness, we can see that the matching procedure appears to work better for the sub-sample of participants that were unemployed in the pre-treatment observation, since their pre-treatment happiness is closer to that of the matched controls. This is most likely because participants who were employed in the pre-treatment observation are unusual in unobserved characteristics. As mentioned in Section 2, one of the formal pre-requisites for entering an SEP is to have been unemployed for 6 out of the preceding 12 months, though there are exceptions for special cases. Participants who were employed in the pre-treatment observation are less likely to fulfill the formal pre-requisite, so they are more likely to be drawn from the special cases that are assigned to SEPs through the discretion of the job-center and therefore more likely to have unusual unobserved characteristics. My preferred specifications are therefore the ones in columns 3 and 4 that are based on participants who were unemployed in the pre-treatment observation. For them, the estimated effect of employment in SEPs is slightly smaller, but still large (at around 0.4) and statistically significant.

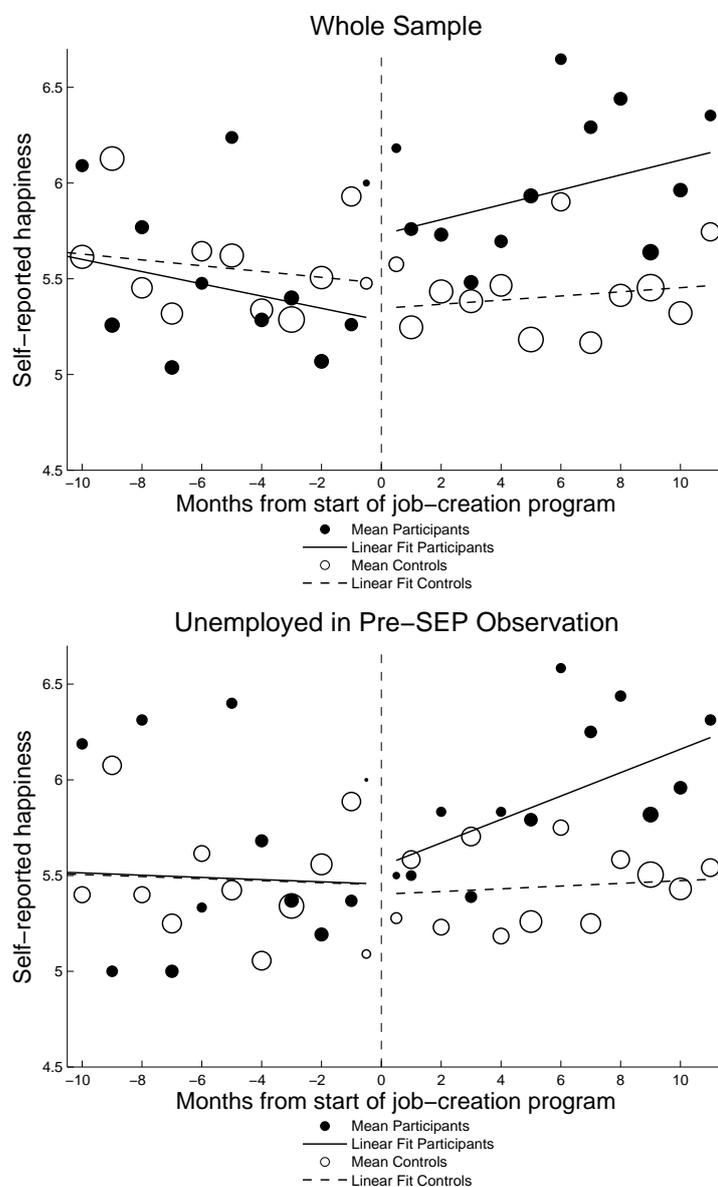
3.5.2 Does the conditional independence assumption hold?

As explained in Section 3.2, the matching estimator's identifying assumption is that individual i 's (possibly counterfactual) happiness when unemployed, $h_{it}(0)$, is independent of participation in a subsidized employment project, conditional on the matching variables. Intuitively, since the estimator uses matched observations to estimate participants' counterfactual outcome if unemployed, the identifying assumption is that participants would have been as (un)happy being unemployed as the matched controls who in fact were unemployed. Since participation in SEPs is non-random, it is not obvious that this assumption holds. I therefore conduct the two robustness tests described in Section 3.2.

I first test whether happier individuals self-select into subsidized employment projects, perhaps because they are more motivated to work or because unobserved shocks - for example to health - affect both happiness and the probability of participation. If this were the case, participants would have been happier than matched controls even if they had remained unemployed and the matching estimator would be biased upward. As a robustness test, I compare the happiness of matched controls and participants in the year before they enter the subsidized employment projects. Column 1 in Table 3.5.1 shows that the average pre-treatment happiness of participants is slightly lower than that of the matched controls and that the difference is not statistically significant, suggesting that there is no self-selection of happier individuals SEPs¹⁰ But average pre-treatment differences are not the only concern. If unobserved shocks increase both happiness and the probability of entering a project, we

¹⁰Surprisingly, the only estimator in which pre-treatment happiness of participants and matched controls differs significantly is the one that matched on pre-treatment happiness. However, this is only due to the fact that matching on pre-treatment happiness decreased the standard error of the difference in happiness, so that the estimate is more precise.

Figure 3.3: Trends of Happiness around the Start of Subsidized Employment projects



Estimates are based on individuals in the German Socio-Economic Panel (GSOEP) that started employment in a SEP in the period 1992-2004 and controls from a nearest neighbor matching procedure. For each participant, the graph plots two observations, one before and one after the start of the project. The horizontal axis plots the time of the interview in months before/after the start of the SEP. Control observations in the “post-treatment” period are plotted at the same time-coordinate as the observation of the matched participant. The time since the control individual’s previous interview is then used to calculate the time-coordinate at which the corresponding pre-treatment control observation is plotted.

would expect the happiness of participants to increase relative to that of non-participants right before the project begins. Thus, despite their slightly lower average happiness in the pre-treatment observation, participants may have been happier than matched controls at the time they entered the project. As a robustness test for this, Figure 3.5.2 plots the average happiness of participants and matched controls in the 12 months before and after the start of employment in a SEP. The plots are constructed as follows: For participants, I use information on the start date of employment in a SEP and the interview date to calculate how many months before or after the beginning of the project an interview took place. For the post-treatment observation, matched controls are plotted at the same time-coordinate as the participants they are matched to. I then use the time since the matched individual's previous interview to calculate the time-coordinate at which her pre-treatment observation is plotted. Since the intervals between two interviews are not fixed, participants and their matches are therefore not necessarily plotted at the same time-coordinate in the pre-treatment period. Still, this procedure makes sure that the pre-treatment time-trend is correctly observed, since the controls' pre-treatment observations are plotted as many months away from the beginning of the project as they would have been if they had entered it at the same time as the participant they are matched to. Consistent with the average difference reported in Table 3.5.1, the top panel in Figure 3.5.2 shows that the pre-treatment happiness of matched controls is slightly higher than that of the participants. Moreover, participants are less happy than matched controls even right before the start of the project. This observation, as well as the fact that participants' happiness is decreasing in the pre-treatment period but starts to increase right at the start of the project, suggests that the results are not driven by self-selection of happier individuals into the projects.

As a second robustness test, I test whether participants and matched controls differ in how strongly their happiness is affected by unemployment. Since participation is largely voluntary, people are likely to self-select into the projects according to how much they benefit from them. Participants' (counterfactual) happiness when unemployed may therefore be different from that of the unemployed matched controls, which would violate the conditional independence assumption. To test for differences in happiness under unemployment, Column 3 in Table 3.5.1 reports differences in the pre-treatment happiness of participants and matched controls who were unemployed in the pre-treatment observation. The point estimate suggests that participants are slightly less happy being unemployed than matched controls, but the difference is very small and not statistically significant. In addition, the bottom panel in Figure 3.5.2 shows that the pre-treatment trends in happiness are virtually identical for unemployed future participants and matched controls, giving no evidence that participants adapt more quickly to unemployment than matched controls.

Taken together these results suggest that there are no substantial violations of the conditional independence assumption, so that the matched controls yield a good counterfactual for the happiness participants would have experienced if they had remained unemployed. This is particularly true for participants who were unemployed in the pre-treatment observation, who are the basis for my preferred specification. The matching estimates therefore suggest a positive causal effect of subsidized employment on happiness.

3.5.3 Disentangling the effects of employment and income

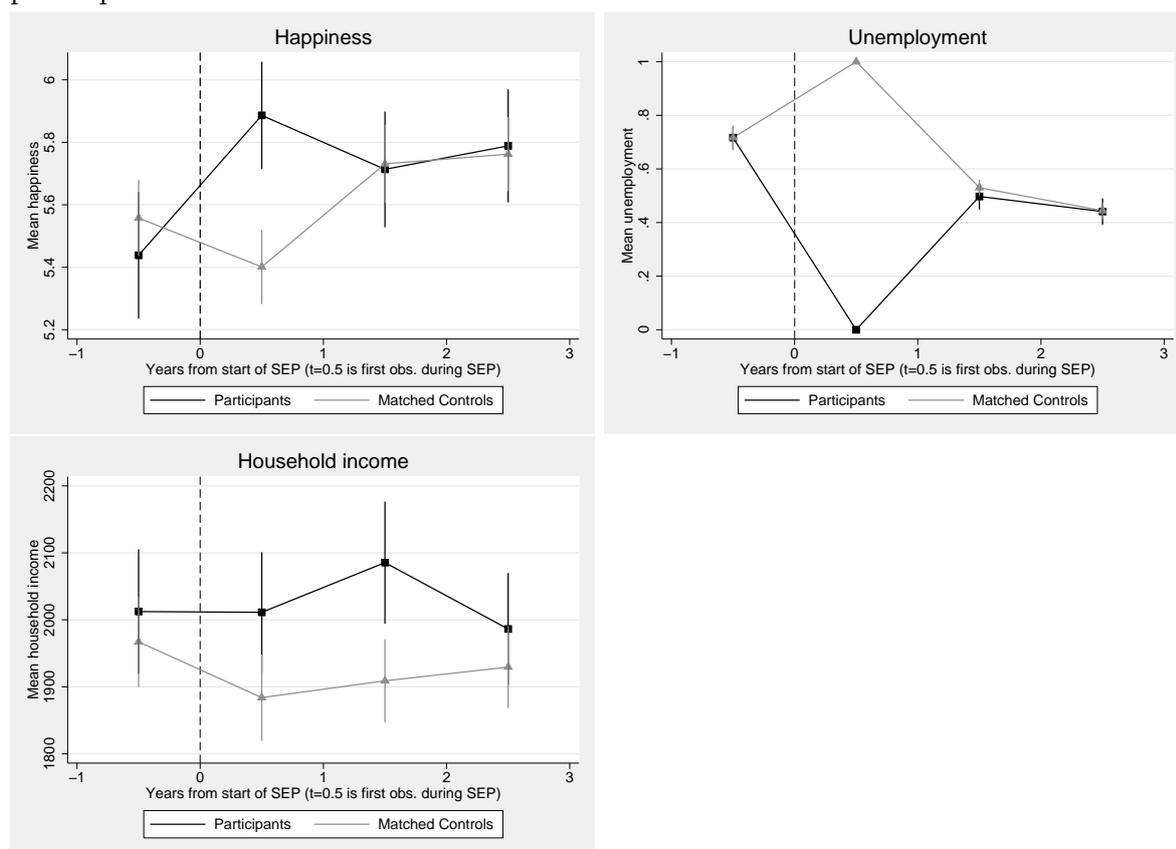
The matching estimators presented in the previous section measure τ , the effect of SEPs on the happiness of individuals who would otherwise have remained unemployed. But as mentioned above, participation in a subsidized employment project has at least two consequences: participants are employed and receive higher incomes. In order to inform policy, it is important to know through which channel - employment or income - SEPs affect happiness. A happiness-based argument for publicly subsidized jobs, similar to the one made by Edlin and Phelps (2009), only holds if employment *per se* increases happiness. If participants in SEPs are only happier because of their higher incomes, increased income transfers would most likely be a more cost-effective way of increasing the happiness of the unemployed. This section presents graphical evidence and quantitative results from fixed effects and instrumental variables estimators, which all suggest that the effect of SEPs on happiness is due to direct psychological benefits and cannot be explained by the higher incomes of the employed alone. The evidence is based on the fact that participation in a SEP prolongs individuals' entitlement to public unemployment benefits. Thus the projects' positive effect on income remains even after employment in it has ended. This creates sufficient independent variation in income and employment to allow me to identify the effect of employment while controlling for differences in income.

Figure 3.4 plots the trends of employment, income and happiness around the start of SEPs. In the first year after entering an SEP, participants are employed in the project and are substantially happier than the unemployed matched controls. Since the duration of SEPs is usually limited to 12 months, most participants leave employment in the subsequent year and are as likely to be employed as matched controls. But since participation in a SEP prolongs individuals' entitlement to unemployment benefits, participants' average income remains higher than that of matched controls. If income were responsible for the projects' effect, we would expect participants to be significantly happier than matched controls until the difference in incomes disappears. But the plot shows that the projects' positive effect on happiness disappears in the second year after their start, at the same time as employment in the project ends for most participants, suggesting that the effect is due to the psychological benefits of employment *per se* and not due to participants' increased incomes.

To obtain quantitative estimates of θ - the pure effect of subsidized employment net of the effect of increased income - I use the estimators described in Section 3.3. The simple fixed effects estimator reported in Table 3.5.3 shows that the correlation between participation in SEPs and happiness remains strong and significant even after controlling for income (both current and future) and unobserved fixed characteristics. The estimated effect is similar in size to the matching estimates reported in Table 3.5.1. The estimated effect of employment in SEPs is slightly smaller than that of employment in a regular job, which is likely due to unobserved heterogeneity in job characteristics.

Tables 3.5.3 and 3.5.3 report results from the fixed effects instrumental variables estimator described in Section 3.3. As described in that section, the estimator uses an indicator for an individual's first observation in a SEP as an instrument for employment in order to control

Figure 3.4: Happiness, unemployment and income before and after the start of SEPs - all participants



Estimates are based on individuals in the German Socio-Economic Panel (GSOEP) that started employment in a SEP in the period 1992-2004. For the time axis, $t=0.5$ is defined as the first observation after the start of employment in the SEP. The average interval between two observations of the same individual in the GSOEP - one year - is used to calculate the other values of t . Happiness is measured on a scale from 0 to 10.

Table 3.3: Employment vs. Income: Fixed Effects Estimates

	Dependent Variable:Happiness		
	(1)	(2)	(3)
Employed in regular job	0.56 (0.03)***	0.51 (0.03)***	0.49 (0.03)**
Employed in SEP	0.45 (0.08)***	0.43 (0.08)***	0.40 (0.08)***
Not seeking employment	0.28 (0.04)***	0.24 (0.04)***	0.26 (0.04)***
Log Household Income		0.23 (0.03)***	0.21 (0.03)***
Log Inc. from Unemp. Benefits		-0.021 (0.003)***	-0.020 (0.005)***
Log Avg. Future Income			0.40 (0.07)***
Age		-0.033 (0.004)***	-0.023 (0.005)
Education (years)		0.005 (0.016)	0.006 (0.017)
Lives with partner		0.18 (0.05)***	0.20 (0.05)
Household Size		-0.061 (0.023)***	-0.078 (0.025)***
Number of children		0.095 (0.027)***	0.090 (0.029)***
Eastern Germany		-0.37 (0.11)***	-0.29 (0.12)**
Constant	6.03 (0.04)	5.78 (0.34)	2.48 (0.66)
Number of observations	34911	34911	30352
Number of individuals	4892	4892	4462

Data source: GSOEP, 1992-2004. *, ** and *** denote statistical significance at the 10, 5 and 1 % levels. All models include individual and year fixed effects. Standard errors are clustered at the individual level. The baseline employment status is unemployed and looking for work.

Table 3.4: Instrumental Variables Estimates: First Stage
 Dependent Variable: Employed

	(1)	(2)	(3)	(4)
First observation after start of SEP	0.52 (0.02)***	0.54 (0.02)***	0.58 (0.03)***	0.59 (0.03)**
Log Household Income		0.070 (0.035)**	0.073 (0.036)**	0.106 (0.039)
Log Inc. from Unemp. Benefits		-0.024 (0.003)***	-0.024 (0.003)***	-0.023 (0.004)***
Log Avg. Future Income				0.24 (0.11)**
Age		-0.0003 (0.0063)	-0.0066 (0.0066)	0.00042 (0.0066)
Education (years)		0.017 (0.031)	0.015 (0.031)	0.009 (0.036)**
Lives with partner		0.055 (0.069)	0.056 (0.069)	0.061 (0.076)
Household Size		-0.063 (0.027)**	-0.063 (0.028)**	-0.072 (0.030)**
Number of children		0.013 (0.039)	0.008 (0.040)	0.017 (0.040)
Eastern Germany		-0.50 (0.15)***	-0.49 (0.15)***	-0.48 (0.17)***
t (years after start of SEP spell)			0.023 (0.023)	0.024 (0.008)***
Constant	0.41 (0.05)	0.38 (0.54)	0.59 (0.53)	-1.66 (1.09)
Number of observations	2493	2493	2493	2216
Number of individuals	406	406	406	371

Data source: German Socio-Economic Panel, 1992-2004. *, ** and *** denote statistical significance at the 10, 5 and 1 % levels. All models include individual and year fixed effects. Standard errors are clustered at the individual level (using a clustered bootstrap with 500 replications). The sample only contains SEP participants and is restricted to individuals' first 6 observations after the start of a SEP.

Table 3.5: Instrumental Variables Estimates: Second Stage
Dependent Variable: Happiness

	(1)	(2)	(3)	(4)
Employed	0.44 (0.15)***	0.53 (0.14)***	0.45 (0.15)***	0.40 (0.16)**
Log Household Income		-0.04 (0.15)	-0.04 (0.15)	-0.04 (0.18)
Log Inc. from Unemp. Benefits		-0.030 (0.011)***	-0.032 (0.011)***	-0.032 (0.011)***
Log Avg. Future Income				0.53 (0.34)
Age		0.002 (0.025)	0.008 (0.027)	0.009 (0.028)
Education (years)		0.22 (0.10)**	0.22 (0.10)**	0.23 (0.11)**
Lives with partner		-0.04 (0.26)	-0.04 (0.26)	0.09 (0.34)
Household Size		0.07 (0.10)	0.06 (0.10)	0.002 (0.10)
Number of children		0.17 (0.13)	0.17 (0.13)	0.19 (0.13)
Eastern Germany		-1.44 (0.74)*	-1.48 (0.74)**	-1.14 (0.71)
t (years after start of SEP spell)			-0.021 (0.023)	-0.023 (0.024)
Constant	5.31 (0.17)	4.12 (2.12)	3.95 (2.13)	-0.28 (3.81)
Number of observations	2493	2493	2493	2216
Number of individuals	406	406	406	371

Data source: German Socio-Economic Panel, 1992-2004. *, ** and *** denote statistical significance at the 10, 5 and 1 % levels. All models include individual and year fixed effects. Standard errors are clustered at the individual level (using a clustered bootstrap with 500 replications). Employment is instrumented by an indicator for the first observation after the start of an employment spell in a subsidized employment project (SEP). The sample only contains SEP participants and is restricted to individuals' first 6 observations after the start of an SEP.

for endogenous exit from SEPs. The instrument exploits the fact that the usual duration of SEPs is 12 months, so that participants' probability of employment drops significantly between their first and second observation after entering a SEP ¹¹.

However, as shown in Figure 3.4, participants' expected incomes do not immediately decrease as employment ends, because participation prolongs their entitlement to payments through the public unemployment insurance. This creates independent variation in employment and income, which makes it possible to use the decrease in the probability of employment after 12 months as an instrument for employment while still controlling for income.

As explained in Section 3.3, I restrict the sample for the fixed effects IV estimator to SEP participants and use only observations made after the start of a SEP. Observations before the start of a SEP are dropped in order to avoid endogeneity bias stemming from unobserved shocks that simultaneously increase happiness and the probability of entering a SEP. To reduce noise from unobserved time-trends, I limit the sample to the first observation after the project's start and the 5 subsequent ones. For participants with multiple spells of employment in a SEP, each spell is treated separately. That is, the first observation in a SEP spell is used as an instrument for employment and the 5 subsequent observations are included in the analysis, regardless of whether the individual enters another SEP during that time. This makes sure that the estimates are not affected by repeated endogenous entry into SEPs. It does, however, have the consequence that some observations are "double-counted", if an individual enters more than one SEP in a 5 year period. To make sure that this double-counting does not lead me to over-state the precision of the estimates, the reported standard errors are clustered at the individual level.

The first-stage results, reported in Table 3.5.3 show that the probability of employment drops between 52 and 59 percentage points between the first observation after entering a SEP and later observations, an effect that is large and statistically significant. The 2-stage least squares estimates in Table 3.5.3 show that employment in SEPs has a large and statistically significant effect on happiness, even after controlling for income, both current and future. Ranging between 0.39 and 0.50, the estimated effect is large compared to the within-individual standard deviation of 1.32 and compared to the "effects" of the control variables. The next subsection discusses the identifying assumptions of the instrumental variables estimator in more detail and presents robustness tests for them. The subsequent sections discuss whether increased consumption or misreporting of happiness can explain the results.

Robustness tests for the fixed effects IV estimator

The identifying assumption for the IV estimator is that the instrument is uncorrelated with the error term. In the present context, this means that there can be no systematic unobserved shocks that affect happiness between individuals' first and second observation

¹¹The average interval between observations in the German Socio-Economic Panel is 12 months.

after entering a SEP. This assumption is likely to hold, since unobserved shocks that occur after the start of a SEP are likely to be evenly distributed over time, so there is no reason to believe that they would affect happiness in the first year differently than in the following years. One concern is that systematic unobserved shocks occur before the start of the program (perhaps because these shocks increase the probability of participation) whose effect persists in the first year of the project and wear off in later years. If this were the case, the instrumental variables estimate would be biased. Reassuringly, the results in Section 5.2 suggest that there are no systematic unobserved shocks to happiness in the run-up to entering a SEP, since the time trends of happiness of participants and matched controls are almost identical in the year before entering the SEP (shown in Figure 3.4). To further control for persistent pre-project shocks, the models in columns 3 and 4 in Table 3.5.3 include a time-trend that begins with the start of the SEP. If pre-project shocks increase the happiness of participants at the beginning of the project, their effects should wear off over time, so that we expect happiness to decrease after the start of the project. Assuming that these (potential) shocks wear off gradually, and not discontinuously between the first and second year after entering a SEP, their effects can be controlled for by a time-trend. The estimates in columns 3, and 4 of Table 3.5.3 should therefore identify the causal effect of employment even in the presence of systematic unobserved shocks to happiness before the start of the program. Additional robustness tests for the exclusion restriction are discussed below.

The exclusion restriction also implies that individuals have to be as happy being employed in SEPs as they are in the jobs they hold in subsequent years. A concern is that regular jobs and jobs in SEPs differ in their effect on happiness due to differences in unobserved characteristics, which would violate the exclusion restriction. As a robustness test, I test whether the individuals in the sample are as happy when employed in the first year after the start of a SEP as they are when employed in subsequent years. The regression results in Table 3.5.3 support the exclusion restriction since there are only small differences between the effect of employment in the first year after the start of a SEP and later years.

The second implication of the exclusion restriction is that individuals would have been as happy being unemployed in the first year after the start of a SEP as they are being unemployed in later periods. Unfortunately, I cannot test this condition in the same way that I tested equality of outcomes under employment since all individuals are employed in the first observation after the start of a SEP. However, it is less likely that this condition is violated since there is less heterogeneity in the situation of the unemployed than in the situation of the employed. One potential violation would occur if individuals adapt to unemployment, so that they are happier being unemployed in later years. As a robustness test, I test for a time trend in the happiness of individuals who are unemployed following a spell in a SEP ($t > 1$). Finding a time-trend would suggest that the effect of unemployment on happiness is changing over time so that the exclusion restriction would be violated. The results in Table 3.5.3 show that this is not the case. The interaction of unemployment and time has only a very small and statistically insignificant effect on happiness, which suggests that the effect of unemployment is stable over time, so that the second condition of the exclusion restriction is satisfied.

Table 3.6: Robustness tests: changes in the effect of employment over time

	Dependent variable: Happiness			
	(1)	(2)	(3)	(4)
Employed	0.45 (0.09)***	0.47 (0.09)***	0.44 (0.10)***	0.40 (0.11)***
Employed in periods $t > 1$	0.017 (0.08)	-0.061 (0.085)	-0.018 (0.11)	-1.8*10-05 (0.11)
Log Household Income		-0.033 (0.15)	-0.036 (0.15)	-0.045 (0.17)
Log Inc. from Unemp. Benefits		-0.032 (0.011)***	-0.033 (0.010)***	-0.032 (0.011)***
Log Avg. Future Income				0.53 (0.35)
Age		0.002 (0.025)	0.008 (0.027)	0.009 (0.029)
Education (years)		0.22 (0.12)*	0.22 (0.11)*	0.23 (0.10)
Lives with partner		-0.034 (0.27)	-0.035 (0.28)	0.089 (0.31)
Household Size		0.061 (0.10)	0.062 (0.10)	0.002 (0.10)
Number of children		0.17 (0.12)	0.17 (0.11)	0.19 (0.13)
Eastern Germany		-1.49 (0.71)**	-1.49 (0.74)**	-1.14 (0.63)*
t (years after start of SEP spell)			-0.021 (0.025)	-0.023 (0.026)
Constant	5.29 (0.17)	4.16 (2.18)	3.97 (2.06)	-0.28 (3.63)
Number of observations	2493	2493	2493	2216
Number of individuals	406	406	406	371

Data source: German Socio-Economic Panel, 1992-2004. *, ** and *** denote statistical significance at the 10, 5 and 1 % levels. All models include individual and year fixed effects. Standard errors are clustered at the individual level (using a clustered bootstrap with 500 replications). The sample only contains SEP participants and is restricted to individuals' first 6 observations after the start of a SEP.

Table 3.7: Robustness tests: adaptation to unemployment

	Dependent variable: Happiness			
	(1)	(2)	(3)	(4)
Not employed	-0.42 (0.11)***	-0.38 (0.11)***	-0.41 (0.12)***	-0.38 (0.12)***
Not employed \times t	-0.016 (0.033)	-0.022 (0.032)	-0.006 (0.039)	-0.009 (0.040)
Log Household Income		-0.039 (0.14)	-0.038 (0.14)	-0.046 (0.17)
Log Inc. from Unemp. Benefits		-0.031 (0.010)***	-0.032 (0.010)***	-0.032 (0.011)***
Log Avg. Future Income				0.53 (0.32)*
Age		-9.6*10 ⁻⁴ (0.025)	0.007 (0.027)	0.008 (0.030)
Education (years)		0.22 (0.11)*	0.22 (0.12)*	0.23 (0.11)**
Lives with partner		-0.032 (0.27)	-0.034 (0.27)	0.090 (0.32)
Household Size		0.063 (0.097)	0.062 (0.10)	0.002 (0.11)
Number of children		0.17 (0.12)	0.17 (0.12)	0.19 (0.12)
Eastern Germany		-1.45 (0.76)*	-1.48 (0.81)*	-1.14 (0.68)*
t (years after start of SEP spell)			-0.021 (0.024)	-0.021 (0.025)
Constant	5.78 (0.17)	4.78 (2.18)	4.44 (2.16)	0.15 (3.50)
Number of observations	2493	2493	2493	2216
Number of individuals	406	406	406	371

Data source: German Socio-Economic Panel, 1992-2004. *, ** and *** denote statistical significance at the 10, 5 and 1 % levels. All models include individual and year fixed effects. Standard errors are clustered at the individual level (using a clustered bootstrap with 500 replications). The sample only contains SEP participants and is restricted to individuals' first 6 observations after the start of a SEP.

Taken together, the results of the robustness tests suggest that the instrumental variables estimator is an unbiased estimator of the Local Average Treatment Effect - the effect of subsidized employment (net of the effect of income) on individuals who participate in a subsidized employment project and are unemployed at some point within 5 years after the project's start.

3.5.4 Can changes in consumption or expected future income explain the results?

The instrumental variables estimators presented in columns 1 through 3 of Table 4 estimate the effect of employment on happiness while controlling for the effect of current income. But this may not be enough to isolate the pure psychological effect of employment. If individuals rationally maximize lifetime utility, their current consumption is a function of their expected lifetime income (Friedman 1957). Thus, if employment in SEPs increases expected lifetime income, it may affect individuals' happiness by increasing their consumption. An increase in expected lifetime income might also increase happiness directly if individuals gain happiness from anticipating future income. To rule these channels out, the model in column 4 of Table 4 adds respondents' average income in all future observations as an additional control variable. In addition, the time trend from the start of the SEP should control for the shock to consumption that comes with starting employment in a SEP. As mentioned in the previous sub-section, this trend controls for shocks that occur at (or before) the start of the SEP and wear off gradually. If individuals conform to the Permanent Income Hypothesis, their consumption should increase discontinuously as they are offered a job in a SEP, since this constitutes a shock to their expected future income. However, in later periods, their consumption should decline gradually¹², so that the time trend should control for the effect of declining consumption. The results in column 4 of Table 4 show that estimated effect of future income is strongly positive, and that the time trend from the start of a SEP is negative, though neither of them is statistically significant. These results are consistent with the hypothesis that future income affects happiness either through consumption or anticipation and that part of the effect of SEPs operates through this channel. But even after controlling for this channel, the remaining effect of employment is large and statistically significant. This result suggests that employment has psychological benefits that are independent of its effects on income and consumption.

3.5.5 Can misreporting of happiness explain the results?

A vital concern when studying self-reported happiness is whether answers to questions like "how satisfied are you with your life?" measure well-being in a meaningful way. One reassuring finding is that self-reported life-satisfaction correlates well with more detailed

¹²In fact, if individuals have quadratic utility over consumption, their expected consumption should follow a linear trend (Hall 1978).

measures of psychological distress (Koivumaa-Honkanen et al. 2004) and predicts objective outcomes like suicide and mortality (e.g. Koivumaa-Honkanen 2001, Chida 2008). Still, in specific cases there could be systematic misreporting of life-satisfaction due to social norms. In many cultures work is seen as a valuable and central aspect of life, so that respondents may be reluctant to admit being happy while unemployed. It is therefore possible that the unemployed under-report their happiness compared to the employed, which would bias the estimated effect of employment upward.

While I cannot fully rule out that unemployed individuals misreport their happiness relative to those in subsidized employment projects, there are several reasons to believe that the effect of misreporting is small. First, the life-satisfaction question is the last question in a long multi-purpose survey (the German Socio-Economic Panel), while the questions about employment are asked in the first half of the survey. Respondents are therefore not “primed” on their employment status when answering the life-satisfaction question. In addition, respondents are not aware that their answers will be used to study the effect of employment on happiness, which should further reduce misreporting due to social norms. Further evidence against misreporting comes from the data. If the unemployed underreport their happiness for reasons of social acceptability, we would expect to see a sharp increase in reported happiness at the start of the SEP. But as shown in Figure 3.5.2, happiness initially remains low and increases over the course of the project; a pattern that is not easily explained by misreporting due to socially preferred answers ¹³.

3.6 Conclusion

This paper tries to answer two questions: does unemployment make people unhappy and, if yes, can subsidized employment increase people’s happiness? Its findings, based on data from the German Socio-Economic Panel, suggest that the answer to both questions is “yes”. A matching estimator suggests that participants in subsidized employment projects (SEPs) are substantially happier than they would have been if they had remained unemployed. Panel data on pre-project happiness suggests that this effect is not due to self-selection of happier individuals into the projects. The data further suggest that the increase in income that comes with subsidized employment does not explain the effect. In the German context, participation in a subsidized employment project prolongs participants’ entitlement to public unemployment benefits, so that their average income does not decrease after the project ends, even though 60% of participants become unemployed. Yet happiness sharply decreases after the project ends, suggesting that most of the previous increase in happiness was due to the projects’ effect on employment and that only a small fraction, if any, can be explained by their effect on income. Taken together, the results presented in the paper suggest that subsidized employment can have a large positive effect on the happiness of individuals who

¹³The upward trend is more plausibly explained by a gradual and cumulative effect of employment on happiness. For example, if part of the psychological benefit of employment comes from the social ties to ones co-workers, we would expect happiness to increase as these social ties strengthen over time.

would otherwise be unemployed.

The paper's results are relevant for two reasons. First, they constitute conclusive evidence for a causal effect of unemployment on happiness. While previous studies (e.g. Clark and Oswald 1994, Winkelmann and Winkelmann 1998, Marks and Fleming 1999, Clark 2003, Carroll 2007) found correlations between (changes in) unemployment and happiness, they were unable to rule out that the correlation was due to reverse causality from happiness to unemployment, or caused by unobserved shocks - for example to health - that simultaneously decrease happiness and increase the probability of unemployment. By showing that the effect of subsidized employment projects on happiness is not due to self-selection of happier individuals into the projects, the current paper provides strong evidence that the effect of unemployment on happiness is causal. Second, the results have implications for labor market policy. For some time, economists (e.g. Edlin and Phelps 2009, Phelps 1994, Katz 1996) have argued that subsidies for low-wage jobs should replace traditional transfer-based welfare policy and several countries (most notably France, but also the Netherlands and the UK) have introduced subsidies of this type. Recently, Edlin and Phelps (2009) have cited potential psychological benefits of employment as an additional argument for subsidising low-wage jobs. The main finding of this paper - that subsidized employment can increase people's happiness directly and not just by increasing their incomes - gives empirical support to their argument.

Chapter 4

The Effect of Alcohol Availability on Marijuana Use: Evidence from the Minimum Legal Drinking Age

Abstract

This paper exploits the Minimum Legal Drinking Age of 21 years to estimate the causal effect of increased availability of alcohol on marijuana use. We find that consumption of marijuana decreases sharply at age 21, while consumption of alcohol increases, suggesting that marijuana and alcohol are substitutes. We further find that the substitution effect between alcohol and marijuana is stronger for women than for men. Our results suggest that policies designed to limit alcohol use have the unintended consequence of increasing marijuana use.

4.1 Introduction

Economic theory suggests that when the cost of consuming a good increases, people will consume more of its substitutes and less of its complements. In the case of alcohol, the substitutes are likely to include other intoxicating substances. The Minimum Legal Drinking Age (MLDA), which restricts access to alcohol for those under 21, is therefore likely to affect the consumption of other drugs among that age group, as it sharply decreases the cost of consuming alcohol for individuals just over the MLDA. When assessing the costs and benefits of policies that aim to reduce alcohol consumption - like the MLDA or alcohol taxes - we need to take possible substitution behavior into account.

For example, proponents of the MLDA at age 21 argue that alcohol consumption in children and adolescents can cause long term and, sometimes, irreversible damages to the brain (Association 2008). In particular, adolescents who drink are more likely to develop smaller

hippocampi, a part of the brain that controls learning and memory, and are more likely to show alterations in their prefrontal cortex (Association 2008). Alcohol consumption has also been shown to induce suicides and car accidents (Carpenter & Dobkin 2009). However, if restricting access to alcohol causes people to switch to substitutes, such as marijuana or other illegal drugs, the benefits of reduced alcohol consumption need to be weighed against the cost of increased consumption of alcohol's substitutes. The potential alcohol substitute we analyze in this paper is marijuana, a substance made of a mixture of flowers, seeds and leaves of the hemp plant. The hemp plant contains tetrahydrocannabinol or THC, a psychoactive chemical that produces most of the intoxicating effects. Consumption of THC has been associated with cognitive deficits and changes in brain morphology and psychiatric disorders (Wilson et al. (2000), Pope et al. (2003), Hall & Degenhardt (2009)). In this paper we study the effects of an increase in the availability of alcohol on the consumption of marijuana.

Most previous studies of substitution between alcohol and marijuana (e.g. DiNardo & Lemieux (2001), Chaloupka & Laixuthai (1997), Pacula (1998), J. et al. (2004), Saffer & Chaloupka (1999), and Farrelly et al. (1999)) are based on cross-sectional (usually between-state) variation in the prices of alcohol and marijuana, the MLDA, alcohol taxes, or laws that partially decriminalize marijuana. A problem for these approaches is that state-level prices of alcohol and marijuana and the policies governing their consumption are likely to be correlated with unobserved characteristics of the population living in those states, making it difficult to infer causality from cross-sectional comparisons (Carpenter & Dobkin 2009).

We address the problem of causal identification that has plagued previous research through a regression discontinuity design. This approach exploits the sharply discontinuous nature of the minimum legal drinking age - the fact that a person cannot legally purchase alcohol up until the day before her 21st birthday, but can do so from her 21st birthday onwards. By comparing substance use in individuals just below and just above the age of 21, we can therefore isolate the causal effect of the MLDA on alcohol and marijuana consumption. The identifying assumption is that, apart from the ability to legally purchase alcohol, individuals just above and just below the age of 21 are similar in all characteristics that determine substance use. The regression discontinuity approach allows us to estimate the extent of substitution between alcohol and marijuana and identify the causal effect of changes in the MLDA on individuals close to 21 years of age.

Our results show that alcohol and marijuana are substitutes. At age 21, we observe a sharp increase in alcohol consumption but a decrease in the marijuana consumption. This suggests that policies that restrict access to alcohol cause an increase in marijuana consumption. Our estimates suggest that the MLDA at age 21 decreases the probability of having consumed alcohol in the past 30 days by 16% and increases the probability of having consumed marijuana by 10%, representing an elasticity of substitution of 0.6. The elasticity of substitution of the frequency of consumption (defined as the number of days in which a substance was consumed) is 0.3. We further find that the substitution effect is substantially stronger for women than for men. Our results suggest that by restricting the age at which people can legally purchase alcohol, the MLDA causes an increase in the consumption of illicit drugs.

The next section reviews the existing literature on the MLDA and marijuana use. Section 3 describes the empirical strategy in more detail. Section 4 presents the data and results, Section 5 shows the robustness of our estimates and Section 6 concludes.

4.2 Literature Review

Most of the previous literature on substitution between marijuana and alcohol exploits between-state variation in the Minimum Legal Drinking Age (MLDA) and marijuana decriminalization during the 1970s and 1980s. DiNardo and Lemieux (2001) estimate a structural model of alcohol and marijuana consumption to test the effect of increases in the MLDA. They analyze state-level percentages of high school seniors that reported having consumed alcohol/marijuana from the Monitoring the Future Surveys (MFS) during the period 1980-1989. Their find that alcohol and marijuana are substitutes and that increases in the MLDA lead to a decrease in alcohol consumption and an increase in marijuana consumption. In a similar study, Chaloupka and Laixuthai (1997) find that youths living in states where marijuana was decriminalized report having consumed less alcohol, providing some evidence of substitution between marijuana and alcohol consumption.

Other studies either find no evidence of substitution or evidence of complementarity. Using data from the National Longitudinal Survey of Youth (NLSY), Thies and Register (1993) do not find statistically significant evidence that state-level marijuana decriminalization affects consumption of alcohol or marijuana. Also using data from the NLSY, Pacula (1998) finds that state beer taxes and the MLDA are positively correlated with marijuana consumption. Focusing on college students, Williams et. al (2004) analyze alcohol and marijuana consumption reported in the Harvard School of Public Health College Alcohol Study. They find that campus regulations banning the consumption of alcohol, and to a lesser extent state policies that restrict alcohol consumption, are negatively correlated with marijuana use. Using data from the National Household Surveys on Drug Abuse (NHSDA), Saffer and Chaloupka (1999) find that, controlling for the price of marijuana, county-level alcohol prices are negatively correlated with marijuana consumption. Also using data from the NHSDA, Farrelly et al. (1999) find that increases in state-level beer prices are negatively correlated with marijuana consumption for youths aged 12 to 20, but not for young adults aged 21 to 30.

In summary, the literature on substitution between alcohol and marijuana finds contradicting results. DiNardo and Lemieux (2001) and Chaloupka and Laixuthai (1997) interpret their findings as reflecting substitution between alcohol and marijuana, while Pacula (1998), Williams et. al (2004), Saffer and Chaloupka (1999) and Farrelly et al. (1999) interpret their findings as reflecting complementarity. One possible reason for these mixed results is that different studies use different surveys and time periods, which prevents comparability. Another reason, perhaps more important, is that many of the previous studies are based on state-level (or in the case of Williams, 2004, campus-level) variations in prices of alcohol and marijuana and policies governing their consumption. While this approach can establish

correlations between substance use, prices and policies, the correlations do not necessarily reflect causal effects, since state-level prices and policies governing alcohol and marijuana are likely to be correlated with unobserved population characteristics that determine alcohol and marijuana consumption (Carpenter & Dobkin 2009). In this paper, we overcome this problem by exploiting the discontinuous nature of the MLDA, which creates an abrupt change in individuals' ability to legally purchase alcohol at age 21. The empirical approach - known as a regression discontinuity design - is described in detail in the next section.

4.3 Empirical Strategy

This paper uses a regression discontinuity design (RDD) to identify the effect of the legal minimum drinking age on alcohol and marijuana use. The RDD approach exploits the sharply discontinuous nature of the minimum legal drinking age - the fact that a person cannot legally purchase alcohol up until the day before her 21st birthday, but can do so from her 21st birthday onwards. Individuals therefore switch from the control regime - being legally prohibited from buying alcohol - to the treatment regime - being allowed to do so - from one day to the next. We can therefore estimate the causal effect of the minimum legal drinking age by comparing individuals who have just turned 21 and individuals who are about to turn 21. Our identifying assumption is that, apart from the ability to legally purchase alcohol, individuals just below and just above the age of 21 are similar in all characteristics that determine substance use, so that differences between the two groups can only be explained by the effect of the minimum drinking age.

Our estimates are based on the standard regression discontinuity estimator described by Imbens and Lemieux (2008):

$$\tau_{RD} = \lim_{x \uparrow 21} [Y_i | X_i = x] - \lim_{x \downarrow 21} [Y_i | X_i = x]$$

where Y_i and X_i denote individual i 's substance use and age, respectively. That is, we estimate the limit of substance use on both sides of the age of 21. The difference between the limits is the regression discontinuity estimate of the effect of the minimum legal drinking age. We follow Carpenter and Dobkin (2009) and estimate the limits by local linear regression on both sides of the age of 21. In practice, this is equivalent to estimating a kernel-weighted regression of the following model (Imbens & Lemieux 2008):

$$Y_i = \beta_0 + X_{it}\beta_1 + X_{it} * D_{it}\beta_2 + D_{it}\tau_{RD} + \gamma_t + \varepsilon_{it}$$

As before, Y_i and X_i denote individual i 's substance use and age, respectively. D_i is an indicator that takes the value 1 if individual i is 21 years old or older. The estimated coefficient τ_{RD} yields the causal effect of the MLDA at age 21 on alcohol/marijuana consumption.

In the following section we describe the data and the results from the graphical analysis and the statistical analysis of equation (1).

4.4 Data and Results

Data on alcohol and marijuana use was obtained from the National Survey of Drug Use and Health (NSDUH), which is administered annually by the U.S. Department of Health and Human Services' Substance Abuse and Mental Health Services Administration (SAMHSA) and conducted by the Research Triangle Institute. The NSDUH provides estimates of alcohol and illicit substance use among persons aged 12 and older at the national and state-level using a randomly selected sample of approximately 70,000 people.

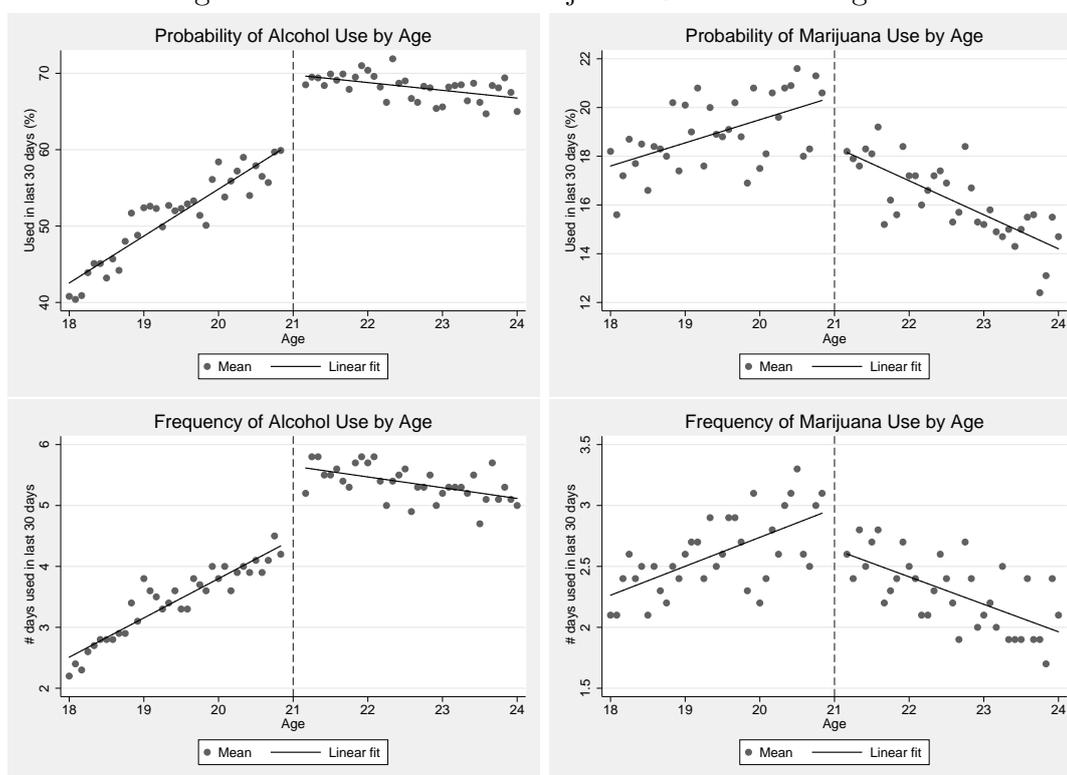
The period of observation for our analysis is 2002-2007. The NSDUH uses two measures of substance use, whether the respondent has used the substance within the past 30 days and the number of days on which the respondent has used it. For alcohol consumption, the question the survey asks is "Think specifically about the past 30 days, from [30 days before the interview date], up to and including today. During the past 30 days, on how many days did you drink one or more drinks of an alcoholic beverage?" For marijuana use, the question it asks is "Think specifically about the past 30 days, from [30 days before the interview date] up to and including today. What is your best estimate of the number of days you used marijuana or hashish during the past 30 days?". Since respondents' precise age is not available in the NSDUH's public-use files, we obtained data on the averages of the substance use measures by month of age from SAMHSA. We obtained these averages for the whole sample and separately for men and women. To maintain confidentiality of the data, SAMHSA only provided us with the average response by month of age but could not provide us with the number of individual responses that were used to calculate the average. For our baseline regressions we use a bandwidth of 3 years around age 21, so that we use data on individuals between the ages of 18 and 24. For this age-group there are 71 month-of-age cells, leading to 71 observations. To avoid measuring the effect of the (anticipated) birthday celebration itself, we drop the observations for the month of the 21st birthday as well as the preceding and following month¹. We use a triangular kernel to estimate local linear regressions on each side of age 21.

Figure 4.1 displays average alcohol and marijuana use between the ages of 18 and 24. The individual observations are averages by month of age; the fitted lines are estimated by linear regressions of substance use on age on both sides of age 21. The top panels show that alcohol consumption increases drastically at age 21. The probability of having consumed alcohol in the last 30 days increases by about 10 percentage points from a baseline just under 60%. A similar result has previously been found by Carpenter and Dobkin and is consistent with the hypothesis that the cost of consuming alcohol decreases significantly at age 21. The frequency of alcohol consumption increases as well, from 4 to 5.5 days drinking out of the previous 30 days.

For marijuana, the effect goes in the opposite direction, though its size is smaller. At age

¹The NSDUH uses a recall period of 30 days, so that the observation in the month following the 21st birthday could still be affected by the birthday celebration. We drop the observation in the preceding month because the anticipation of the birthday celebration may lead people to consume fewer drugs than they normally would.

Figure 4.1: Alcohol and Marijuana Use Around Age 21



Data source: NSDUH 2002-2007. Scatter points denote averages by month of age. Lines are linear fits, estimated separately on both sides of age 21

21, the probability of marijuana use decreases by about 2 percentage points from a baseline of about 20%. The frequency of marijuana use decreases by about 0.3 days out of a 30 day period, from a baseline of about 2.3 days.

Table 4.1: Effect of the MLDA on Alcohol and Marijuana Use: Regression Discontinuity estimates

	Used in last 30 days (%)		# of days used in last 30 days	
	Alcohol	Marijuana	Alcohol	Marijuana
Over 21 (τ_{RD})	9.83 (0.79)***	-2.01 (0.54)***	1.30 (0.10)***	-0.31 (0.11)***
Age	5.23 (0.44)***	0.83 (0.30)***	0.55 (0.06)***	0.23 (0.06)***
Age x Over 21	-6.25 (0.62)**	-2.15 (0.42)***	-0.73 (0.08)***	-0.47 (0.09)***
Constant	60.0 (0.56)***	20.3 (0.38)***	4.34 (0.07)***	2.97 (0.08)***
Number of observations	68	68	68	68

Data source: NSDUH 2002-2007. Each observation is the average of substance use over a month-of-age cell. All estimates are from local linear regressions using a triangular kernel with bandwidth of 3 years, centered at age 21. Standard errors in parenthesis. ***, ** and * denote statistical significance at the 1, 5 and 10 % levels.

Table 4.1 presents quantitative estimates from the local linear regression approach described in the previous section. The results reported in the table are for local linear regressions with a triangular kernel and a bandwidth of 3 years. Each observation is the mean of alcohol/marijuana consumption in the month of age. Robustness tests for different bandwidths are reported in Section 5. The regression results reinforce the visual impression gained from the graphs. There is a strong increase in consumption of alcohol (both probability and frequency of use) at age 21, while consumption of marijuana decreases. The changes are statistically significant and their sizes are similar to the changes visible on the graphs.

The results indicate that alcohol and marijuana are substitutes, at both the extensive and the intensive margin. The decrease in the probability of marijuana use at age 21 suggests that some individuals who use marijuana before the age of 21 stop using it (or at least use it less regularly) once they turn 21 and are able to legally consume alcohol. In absolute terms, the substitution effect on the probability of marijuana use is not very large - a 9.8 percentage point increase in the probability of alcohol consumption leads to a 2 percentage point decrease in marijuana use. However, the 9.8 percentage point increase in alcohol consumption constitutes a 16% increase from the estimated baseline consumption of 60% just below age 21, and the 2 percentage point decrease in marijuana use constitutes an 10% decrease from baseline use, resulting in an estimated elasticity of 0.6.

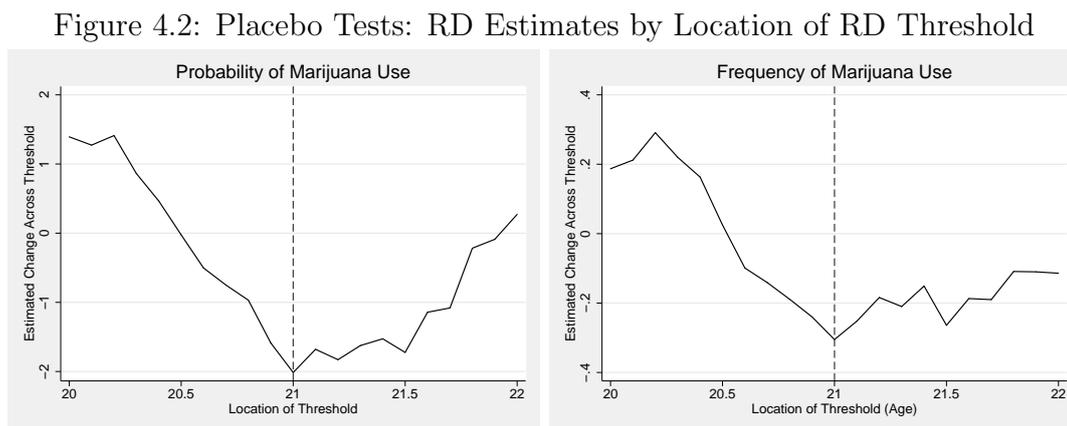
The estimated decrease in the frequency of marijuana use is 0.3 days per month which constitutes a 10% decline from the baseline of 3 days at age 21. Since the decline in the frequency of marijuana use is of similar size (in percentage terms) as the decline in the probability of use, it is unlikely that the estimated drop in the frequency of use is solely driven by the extensive margin, since people who stop using marijuana altogether probably

used it with lower frequency than the average user to begin with. If the entire decline came from these ‘lighter’ users, we would therefore expect that the decline in the average frequency would be lower (in terms of percentage) than the decline in the probability of use. Hence, our estimates suggest that the decline in marijuana use at age 21 occurs on both the extensive and intensive margins.

4.5 Robustness Tests

4.5.1 Placebo Tests for Location of the Discontinuity

In order to make sure that our estimated decline in marijuana use is driven by the change in alcohol accessibility at age 21 and not than merely due to a time trend that follows an inverted u-shape, we conduct placebo tests for the location of the discontinuity. For these tests, we estimate the same regression as in Table 4.1, but vary the location of the threshold. If there really is a discontinuity at age 21 the estimated change in substance use should be largest if the regression’s threshold is located at 21.



Data source: NSDUH 2002-2007. The vertical axis plots RD estimates of the change in Marijuana use across the age-threshold specified on the horizontal axis. For details on the RD estimation, see the description in Table 4.1.

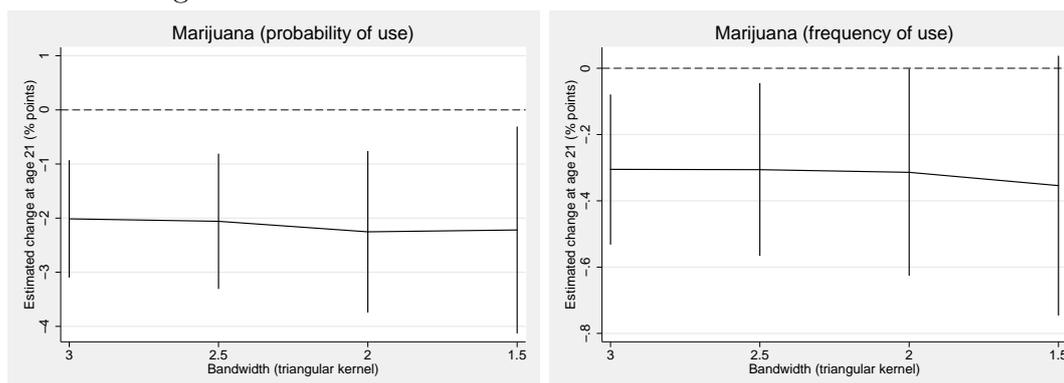
Figure 4.2 displays the estimated change across the threshold as a function of the threshold’s location. The figure shows that for both probability and frequency of marijuana use, the estimated decrease in use across the regression threshold is largest if the threshold is located at exactly 21 years. This increases our confidence that our results are driven by the discontinuous change in alcohol accessibility at age 21 rather than by a time trend that follows an inverted u-shape.

4.5.2 Robustness Tests for Choice of Bandwidth

A crucial parameter for local linear regressions like the ones reported above is the choice of bandwidth. By choosing a bandwidth that is too small we reduce the effective sample size

and obtain estimates of low precision. By choosing a bandwidth that is too large we increase the risk of mis-specification if the relationship between age and substance use is non-linear. Though some authors have suggested rules-of thumb for bandwidth choice (e.g. Fan and Gijbels 1996), no rule-of-thumb guarantees an optimal choice of bandwidth. Imbens and Lemieux (2008) therefore suggest robustness tests for different choices of bandwidth.

Figure 4.3: Robustness of Results to Choice of Bandwidth

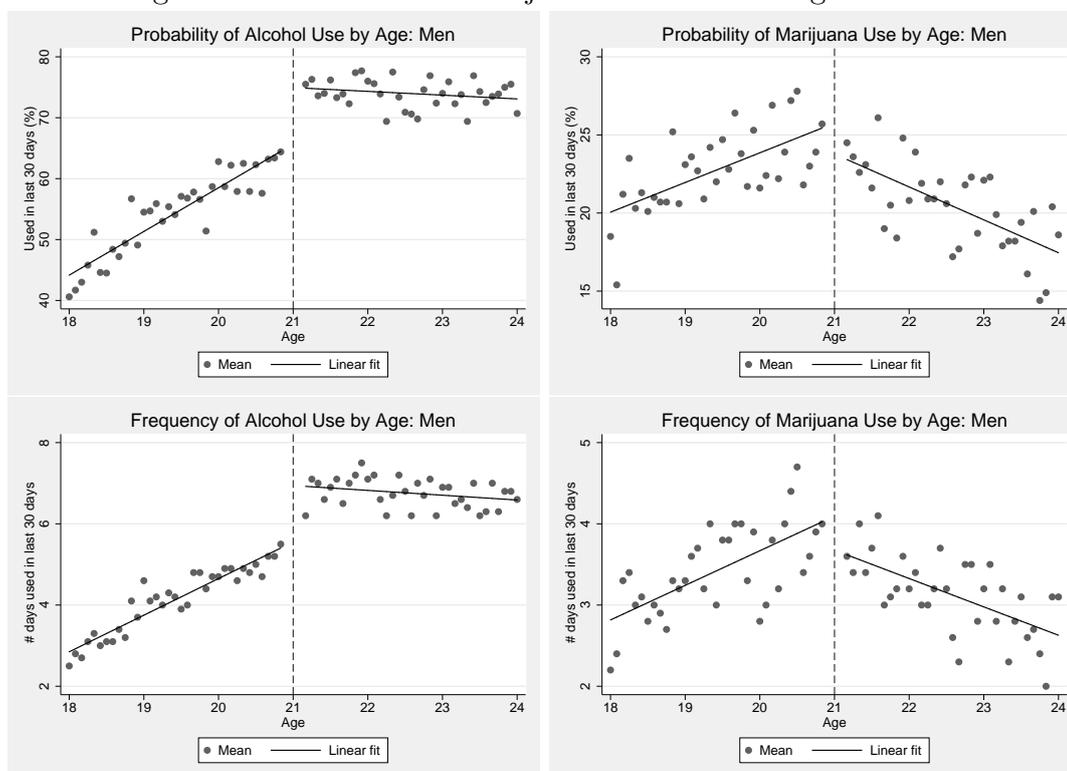


Data source: NSDUH 2002-2007. The vertical axis plots RD estimates of the effect of the MLDA on Marijuana use, using the bandwidth specified on the horizontal axis. Intervals are 95% confidence intervals. For details on the RD estimation, see the description in Table 4.1.

The results of these tests are reported in Figure 4.3. For the tests, we decrease the bandwidth in steps of 6 months until we reach half of the initial bandwidth of 3 years. The point estimates and 95% confidence intervals are plotted on the vertical axes of the graphs, against the bandwidth on the horizontal axes. If the estimates based on larger bandwidths suffer from specification bias due to a non-linear relationship between age and substance use, we would expect the point estimates to change substantially as the bandwidth becomes smaller, since the linear functional form better approximates the true relationship over smaller intervals. If there is no specification bias, we would expect the point estimates to have small fluctuations and hence be robust to the choice of bandwidth. Since the estimates for smaller bandwidths are based on smaller effective samples, we naturally expect the confidence intervals to increase as the bandwidth becomes smaller.

The results in Figure 4.3 show that the estimates are robust to the choice of bandwidth. The point estimates differ very little for bandwidths between 18 months and 3 years. As expected, the size of the confidence intervals increases for smaller bandwidths, since fewer observations are used for estimation. Nevertheless, the estimated effect on the probability of marijuana use is statistically significant for all tested bandwidths; the effect on the frequency of marijuana use is significant at the 10% level for all tested bandwidths (not shown) and at the 5% level for bandwidths of 2 years and larger.

Figure 4.4: Alcohol and Marijuana Use Around Age 21: Men



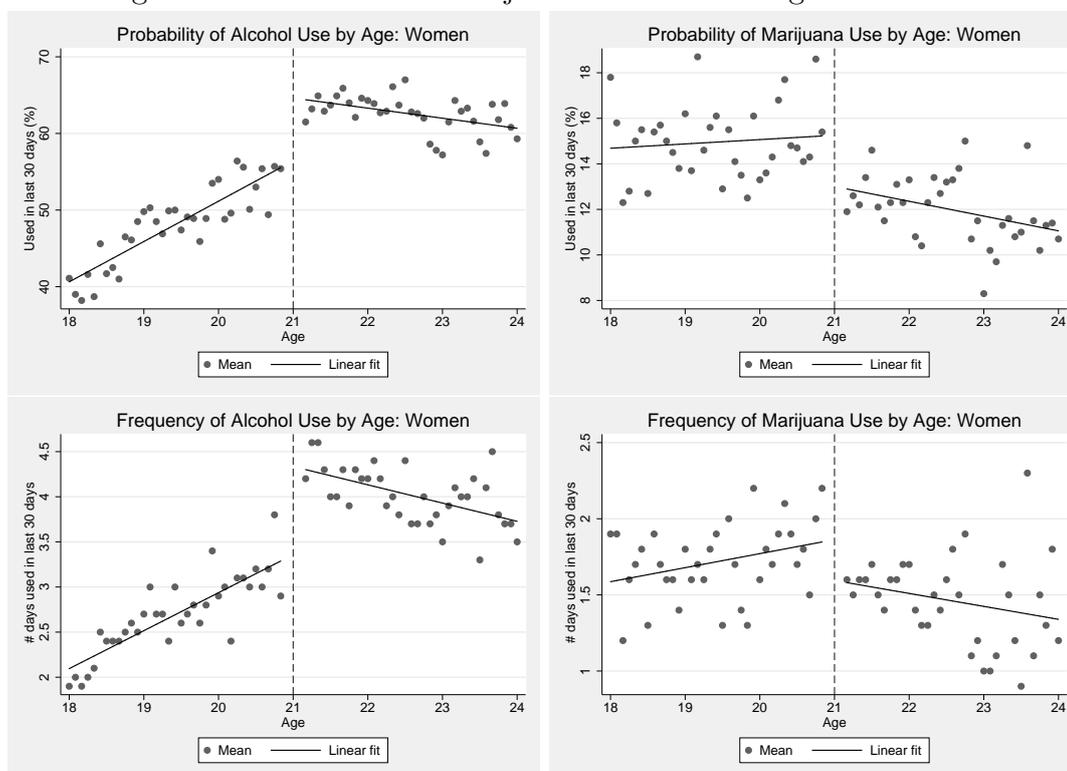
Data source: NSDUH 2002-2007, subsample of male respondents. Scatter points denote averages by month of age. Lines are linear fits, estimated separately on both sides of age 21

Table 4.2: Effect of the MLDA on Alcohol and Marijuana Use: Men

	Used in last 30 days (%)		# of days used in last 30 days	
	Alcohol	Marijuana	Alcohol	Marijuana
Over 21 (τ_{RD})	10.76 (1.11)***	-1.47 (0.92)	1.48 (0.15)***	-0.31 (0.19)*
Age	6.11 (0.62)***	1.39 (0.51)***	0.78 (0.08)***	0.35 (0.10)***
Age x Over 21	-7.04 (0.87)***	-3.45 (0.72)***	-0.87 (0.11)***	-0.73 (0.15)***
Constant	64.5 (0.79)***	25.2 (0.65)***	5.43 (0.10)***	4.01 (0.13)***
Number of observations	68	68	68	68

Data source: NSDUH 2002-2007, subsample of male respondents. Each observation is the average of substance use over a month-of-age cell. All estimates are from local linear regressions using a triangular kernel with bandwidth of 3 years, centered at age 21. Standard errors in parenthesis. ***,** and * denote statistical significance at the 1, 5 and 10 % levels.

Figure 4.5: Alcohol and Marijuana Use Around Age 21: Women



Data source: NSDUH 2002-2007, subsample of female respondents. Scatter points denote averages by month of age. Lines are linear fits, estimated separately on both sides of age 21

Table 4.3: Effect of the MLDA on Alcohol and Marijuana Use: Women

	Used in last 30 days (%)		# of days used in last 30 days	
	Alcohol	Marijuana	Alcohol	Marijuana
Over 21 (τ_{RD})	8.85 (1.06)***	-2.62 (0.68)***	1.08 (0.12)***	-0.29 (0.10)*
Age	4.50 (0.59)***	0.45 (0.38)	0.37 (0.06)***	0.15 (0.06)**
Age x Over 21	-5.65 (0.83)***	-1.02 (0.53)*	-0.61 (0.09)***	-0.27 (0.08)***
Constant	55.9 (0.75)***	15.5 (0.48)***	3.30 (0.08)***	1.92 (0.07)***
Number of observations	68	68	68	68

Data source: NSDUH 2002-2007, subsample of female respondents. Each observation is the average of substance use over a month-of-age cell. All estimates are from local linear regressions using a triangular kernel with bandwidth of 3 years, centered at age 21. Standard errors in parenthesis. ***,** and * denote statistical significance at the 1, 5 and 10 % levels.

4.5.3 Results by Gender

In order to address the possibility of differentiated substitution effects by gender, we replicate the analysis for men and women separately. Figures 4.4 and 4.5 show the trends in alcohol and marijuana use around age 21 for men and women, respectively. Tables 4.2 and 4.3 show the corresponding RD estimates. The results show that men have higher baseline levels of consumption of both alcohol and marijuana. However, the effect of the MLDA on marijuana use is larger for women. For men aged 21 or older the probability of having consumed alcohol is about 11 percentage points higher than younger men and the probability of having consumed marijuana is 1.5 percentage points lower. Compared to a baseline probability of consumption of 65%, the increase of 11 percentage points corresponds to a 17% increase in the probability of consuming alcohol. The decrease of 1.5 percentage points in the probability of consuming marijuana corresponds to a 6% decrease. The estimated elasticity of substitution of the probability of consuming alcohol and marijuana is 0.35. For women, the 16% increase in the probability of consuming alcohol compared to baseline consumption induces a 17% decrease in the probability of consuming marijuana, which represents an elasticity of substitution of about 1.05. The substitution effect on the frequency of substance use is also stronger for women. Women aged 21 or older consume alcohol 1.1 more days/month than younger women and consume marijuana 0.29 days/month less than younger women. Compared to baseline consumption, these effects represent a 33% increase in the frequency of alcohol consumption and a 15% decrease in the frequency of marijuana use. The implied elasticity of substitution of the frequency of alcohol and marijuana use is 0.45. For men, the MLDA induces an increase of 1.5 days/month in the frequency of alcohol consumption and a decrease of 0.3 days/month in the frequency of marijuana consumption. Since men consume alcohol 5.4 days/month and marijuana 4 days/month at baseline, the MLDA induces a 28% increase in the frequency of alcohol consumption and a 7.5% decrease in the frequency of marijuana consumption. For men, the estimated elasticity of substitution of the frequency of alcohol and marijuana consumption is 0.27, which is substantially smaller than the one estimated for women.

4.6 Conclusions

By exploiting the sharp decrease in the effective cost of alcohol consumption induced by the Minimum Legal Drinking Age (MLDA) at age 21, this paper estimates the causal effect of legal access to alcohol on marijuana consumption. Our identifying assumption is that, apart from the ability to legally purchase alcohol, individuals just above and just below the age of 21 are similar in all characteristics that determine substance use. Compared to previous research (e.g. DiNardo & Lemieux (2001), Chaloupka & Laixuthai (1997), Pacula (1998), J. et al. (2004), Saffer & Chaloupka (1999), and Farrelly et al. (1999)), this approach has the advantage of not having to rely on cross-sectional (often state-level) variation in alcohol and marijuana prices and related policies, which are likely to be correlated with unobserved

characteristics of the population. This allows us to cleanly identify the causal effect of the MLDA in a way that is not afflicted by omitted variable bias.

Our results show that legal access to alcohol causes a significant decrease in marijuana use among young adults close to the age of 21. The point estimates suggest that marginally lowering the MLDA would decrease the probability of marijuana consumption in the affected age group by about 10%. The substitution effect is substantially larger for women than for men. Our results suggest that marijuana and alcohol are substitutes, so that a decrease in the ‘full’ price of alcohol (including the cost of access) leads to a decrease in marijuana use. The main implication of our study is that policies - such as the MLDA - that are aimed at restricting alcohol consumption among young adults are likely to have the unintended consequence of increasing the use of illegal drugs such as marijuana. When assessing the net benefits of alcohol-related policies these substitution effects need to be taken into account in order to assess the trade-off between the positive health effects from reduced alcohol consumption and the negative effects of increased use of other substances.

Bibliography

- Abadie, A. & Imbens, G. 2002, Unpublished Working Paper
- Abinales, P. N. 2000, *Making Mindanao : Cotabato and Davao in the formation of the Philippine nation-state* (Quezon City, Philippines: Ateneo de Manila University Press)
- Angrist, J. D. & Kugler, A. D. 2008, *Review of Economics and Statistics*, 90, 191
- Ariely, D., Kamenica, E., & Prelec, D. 2008, *Journal of Economic Behavior and Organization*, 67, 671
- Arrow, K. J. & Dasgupta, P. S. 2009, *Economic Journal*, 119, F497
- Association, A. M. 2008, *Harmful Consequences of Alcohol Use on the Brains of Children, Adolescents, and College Students* (American Medical Association)
- Balisacan, A. M. & Edillon, R. G. 2003, *Second Poverty Mapping and Targeting Study for Phases III and IV of KALAHY-CIDSS*, Tech. rep., Asia-Pacific Policy Center
- Balisacan, A. M., Edillon, R. G., & Ducanes, G. M. 2002, *Poverty Mapping and Targeting for KALAHY-CIDSS*, Tech. rep., Asia-Pacific Policy Center
- Beath, A., Fotini, C., & Enikolopov, R. 2010, Unpublished Working Paper
- Berman, E., Shapiro, J. N., & Felter, J. H. 2008, *National Bureau of Economic Research Working Paper Series*, No. 14606
- Bernhard, S., Hohmeyer, K., Jozwiak, E., Koch, S., Kruppe, T., Stephan, G., & Wolff, J. 2008, *IAB-Forschungsbericht*, Vol. 2/2008, *Aktive Arbeitsmarktpolitik in Deutschland und ihre Wirkungen* (Institut fuer Arbeitsmarkt- und Berufsforschung)
- Blanchflower, D. G. & Oswald, A. J. 2008, *Journal of Health Economics*, 27, 218
- Blattman, C. & Miguel, E. 2010, *Journal of Economic Literature*, 48, 3
- Bockermann, P. & Ilmakunnas, P. 2009, *Health Economics*, 18, 161
- Caliendo, M., Hujer, R., & L., T. S. 2008, *Applied Economics*, 40, 1101
- Carpenter, C. & Dobkin, C. 2009, *American Economic Journal: Applied Economics*, 1, 162
- Carroll, N. 2007, *Economic Record*, 83, 287
- Chaloupka, F. J. & Laixuthai, A. 1997, *Eastern Economics Journal*, 23, 253
- Chapman, W. 1987, *Inside the Philippine Revolution*, 1st edn. (New York: W.W. Norton)
- Chida, Y. & Steptoe, A. 2008, *Psychosomatic Medicine*, 70, 741
- Clark, A. & Oswald, A. J. 1994, *Economic Journal*, 104, 648
- Clark, A. E. 2003, *Journal of Labor Economics*, 21, 323
- Collier, P. & Hoeffler, A. 2004, *Oxford Economic Papers*, 56, 563
- Collier, P., Hoeffler, A., & Soderbom, M. 2004, *Journal of Peace Research*, 41, 253

- Corpus, V. N. 1989, *Silent War* (Quezon City, Philippines: VNC Enterprises)
- Dal Bó, E. & Dal Bó, P. forthcoming, *Journal of the European Economic Association*
- Department for Social Welfare and Development. 2009, *Trials and Triumphs: Communities Fighting Poverty through KALAHÍ-CIDSS* (Quezon City, Philippines)
- DiNardo, J. & Lemieux, T. 2001, *Journal of Health Economics*, 20, 991
- Dube, O. & Vargas, J. F. 2007, CERAC Working Paper No. 2
- Edlin, A. & Phelps, E. 2009, *The Economist's Voice*, 6, 1
- Elbe, S. 2002, *International Security*, 27, 159
- Elbers, C., Lanjouw, J. O., & Lanjouw, P. 2003, *Econometrica*, 71, 355
- Ellingsen, T. & Johannesson, M. 2007, *Journal of Economic Perspectives*, 21, 135
- Fan, J. & Gijbels, I. 1996, *Local Polynomial Modeling and Its Applications* (London: Chapman and Hall)
- Farrelly, M. C., Bray, J. W., Zarkin, G. A., Wendling, B. W., & Pacula, R. L. 1999, NBER Working Paper Series, 6940
- Fearon, J. D. 1995, *International Organization*, 49, 379
- . 2004, *Journal of Peace Research*, 41, 275
- Fearon, J. D. & Laitin, D. D. 2003, *American Political Science Review*, 97, 75
- Felter, J. H. 2005, PhD thesis, Stanford University, Stanford, CA
- Frey, B. & Stutzer, A. 2002, *Journal of Economic Literature*, 40, 402
- Friedman, M. 1957, in *A Theory of the Consumption Function* (Princeton University Press)
- Fudenberg, D. & Tirole, J. 1991, *Journal of Economic Theory*, 53, 236
- Gelman, A. & Hill, J. 2007, *Data Analysis Using Regression and Multilevel/Hierarchical Models* (New York, USA: Cambridge University Press)
- Ghobarah, H., Huth, P., & Russett, B. 2004, *Social Science & Medicine*, 59, 869
- Gompert, D. C., Kelly, T. K., Lawson, B. S., Parker, M., & Colloton, K. 2009, *Reconstruction Under Fire: Unifying Civil and Military Counterinsurgency* (RAND Corporation)
- Grossman, H. I. 1999, *Oxford Economic Papers*, 51, 267
- Hall, R. E. 1978, *Journal of Political Economy*, 86, 971
- Hall, W. & Degenhardt, L. 2009, *Lancet*, 374, 1383
- Harbom, L. & Wallensteen, P. 2010, *Journal of Peace Research*, 47, 501
- Holden, W. N. & Jacobson, R. D. 2007, *Canadian Geographer*, 51, 475
- Hujer, R., Caliendo, M., & L., T. S. 2004, *Research in Economics*, 58, 257
- Ibanez, A. M. & Velez, C. E. 2008, *World Development*, 36, 659
- Imbens, G. W. & Lemieux, T. 2008, *Journal of Econometrics*, 142, 615
- Izard, C. E. 1991, *The Psychology of Emotions* (New York: Plenum Press)
- J., W., Pacula, R. L., & Chaloupka, F. J. 2004, *Health Economics*, 13, 825
- Jones, G. R. 1989, *Red Revolution: Inside the Philippine Guerrilla Movement* (Boulder, CO: Westview Press)
- Kalyvas, S. N. 2006, *The Logic of Violence in Civil War*, Cambridge studies in comparative politics (Cambridge: Cambridge University Press)
- Katz, L. F. 1996, NBER Working Paper Series, 5679
- Koivumaa-Honkanen, H., Honkanen, R., Viinamaki, H., Heikkila, K., Kaprio, J., & Kosken-

- vujo, M. 2001, *American Journal of Psychiatry*, 158, 433
- Koivumaa-Honkanen, H., Kaprio, J., Honkanen, R., Viinamaki, H., & Koskenvuo, M. 2004, *Social Psychiatry and Psychiatric Epidemiology*, 39, 994
- Kreuzer, P. 2005, *Political Clans in the Southern Philippines*, Tech. rep., Peace Research Institute Frankfurt, Frankfurt
- . 2009, *Behemoth*, 1, 47
- Kreuzer, P. & Werning, R. 2007, *Voices from Moro Land: Perspective from Stakeholders and Observers on the Conflict in the Southern Philippines (Petaling Jaya: Strategic Information and Research Development Centre)*
- Labonne, J. & Chase, R. S. 2009, *World Development*, 37, 219
- Lacina, B. & Gleditsch, N. P. 2005, *European Journal of Population / Revue Européenne de Démographie*, 21, 145
- Lobrigo, J. E., Imperial, S., & Rafer, N. 2005, *Philippine Human Development Report*
- Mansuri, G. & Rao, V. 2004, *The World Bank Research Observer*, 19, 1
- Marks, G. N. & Fleming, N. 1999, *Social Indicators Research*, 46, 301
- Mastekaasa, A. 1996, *Journal of Community & Applied Social Psychology*, 6, 189
- McKee-Ryan, F. M., Kinicki, A. J., Song, Z., & Wanberg, C. R. 2005, *Journal of Applied Psychology*, 90, 53
- Miguel, E., Satyanath, S., & Sergenti, E. 2004, *The Journal of Political Economy*, 112, 725
- Murray, C., King, G., Lopez, A. D., Tomijima, N., & Krug, E. 2002, *British Medical Journal*, 324, 346
- Olken, B. A. 2007, *Journal of Political Economy*, 115, 200
- . 2010, *American Political Science Review*, 104, 243
- Pacula, R. L. 1998, *Journal of Health Economics*, 17, 557
- Parker, A. 2005, *Empowering the Poor: The KALAHI-CIDSS Community-Driven Development Project* (Washington, DC: World Bank)
- Phelps, E. S. 1994, *American Economic Review*, 84, 54
- Pope, H. G., Gruber, A. J., Hudson, J. I., Cohane, G., Huestis, M. A., & Yurgelun-Todd, D. 2003, *Drug and Alcohol Dependence*, 69, 303
- Powell, R. 2004, *The American Political Science Review*, 98, 231
- . 2006, *International Organization*, 60, 169
- Reno, W. 1998, *Warlord Politics and African States* (Boulder, CO: Lynne Rienner)
- Ross, M. L. 2004, *Journal of Peace Research*, 41, 337
- Ryan, R. M. & Deci, E. L. 2000, *American Psychologist*, 55, 68
- Saffer, H. & Chaloupka, F. J. 1999, in
- Salehyan, I. 2007, *Civil Wars*, 9, 127
- Salm, M. 2009, *Health Economics*, 18, 1075
- Sarkees, M. R. & Schafer, P. 2000, *Conflict Management and Peace Science*, 18, 123
- Schiavo-Campo, S. & Judd, M. 2005, *The Mindanao Conflict in the Philippines : Roots, Costs, and Potential Peace Dividend*, *Social Development Papers, Conflict Prevention and Reconstruction* 24, World Bank, Washington, D.C. :
- Sison, J. M. & Werning, R. 1989, *The Philippine Revolution: The Leader's View* (New York:

- Crane Russak)
- Thies, C. F. & Register, C. A. 1993, *The Social Science Journal*, 30, 385
- U.S. Army/Marine Corps. 2007, *The U.S. Army/Marine Corps Counterinsurgency Field Manual: U.S. Army Field Manual No. 3-24: Marine Corps Warfighting Publication No. 3-33.5* (Chicago: University of Chicago Press)
- Wilson, W., Mathew, R., Turkington, T., Hawk, T., Coleman, R. E., & Provenzale, J. 2000, *Journal of Addictive Diseases*, 19, 1
- Winkelmann, L. & Winkelmann, R. 1998, *Economica*, 65, 1