When Regulations Backfire:
The Case of the Community Reinvestment Act
JOB MARKET PAPER

Konstantin Golyaev*

September 3, 2010

Abstract

The Community Reinvestment Act (CRA) encourages depositary institutions to lend in low- and moderate-income areas of their markets. This regulation had recently been blamed for contributing to the latest mortgage default crisis. I exploit the discountinuities in the banks' definitions of markets to identify the causal impact of the CRA on the residential mortgage loan approval decisions. I employ a nonlinear Bayesian IV method to quantify the above effect, allowing for the unobserved heterogeneity among loan applicants. I find that, other things equal, the loans that satisfy the CRA criteria have an average of 33% higher chance of being approved, implying almost $500,000$ extra loans originated in 2005 in California. I then demonstrate that in 2010 the foreclosure rates were 5.36 times higher in the CRA-eligible areas. This suggests that 1 out of 6 CRA-induced loans had failed to perform and thus justifies the blame.

* I am indebted to my advisor Patrick Bajari for his constant support and encouragement. I would like to thank my committee members Amil Petrin, Thomas Holmes and Jeffrey McCullough for their suggestions that greatly improved the paper. Special thanks go to Richard Todd, without whom this project would have never been possible. I am grateful to David Wiczer for his help with the C code for some of my computations. I thank Michael Grover, Ron Feldman, Joseph Schneberger, Jason Schmidt and Daniel Rozycki for their assistance with the HMDA and CRA data, Paul Calem, Glenn Canner and Robert Avery for their suggestions, Gautam Gowrisankaran, Robert Town, Minjung Park, Kyoo Il Kim, Thomas Youle, Brian Adams, Nadezda Golyaeva, Michael Reszler, Alexander Stepanov, Christina Marsh, Andrea Szabo, Sunil Kishore, Lydia Cheung, Mahmoud El Amin, Kevin Williams and all participants of the University of Minnesota Applied Microeconomics Workshop. The usual disclaimer applies.

1 Introduction

For the average American a mortgage loan represents their only chance of purchasing a home. Home mortgage lending is thus a sizeable industry subject to a number of federal regulations. The Community Reinvestment Act (CRA) is a mortgage lending regulation that has recently received considerable attention. The CRA encourages depositary institutions[1] to expand what is considered "safe and sound" lending (including mortgage lending) in the lower-income areas (see Avery, Bostic, and Canner (2000) for more information).

The last decade had seen an unprecedented increase in mortgage originations, particularly with the widespread proliferation of so-called subprime mortgages. Such loans were usually made to people with less than stellar credit histories. These mortgages were often subject to higher annual percentage rates. Starting in early 2007, the mortgage industry saw an abrupt increase in delinquencies and, later, foreclosures, particularly in its subprime segment. The ensuing housing market turbulence triggered a broader economic turmoil, which is linked to the latest worldwide recession. These events have become a focus of several recent studies such as Bajari, Chu, and Park (2009).

Some speculated that CRA-induced incentives forced financial institutions to weaken loan evaluation standards, extending too much credit to high-risk individuals (Liebowitz, 2009). Others denounced the above claim, arguing that the CRA had been around for the last thirty years, and as such, it is naïve to expect that banks only started responding to CRA incentives during the last decade (Bair, 2008). To establish the causal relationship between the CRA and excessive mortgage originations, one really has to address two questions. First, is there any link between the CRA and extra mortgage loan approvals? And second, in case the CRA did cause the banks to approve more loans, how did those extra loans perform?

I address these two questions in this paper. First, I identify the causal effect of the CRA by employing the ideas developed by the regression discontinuity literature. The CRA requires its subjects to define *assessment areas* that serve as proxies for markets in which banks carry out the majority of their business. When regulators evaluate the bank's operation under the CRA,

[1] Most common types of depositary institutions are banks and thrifts, and I will use the term "banks" in this paper to refer to all of them.

substantially higher weight is placed on its performance within its assessment areas. At the end of the year, regulators assign CRA ratings to each bank. These ratings are of considerable importance to their subjects: a bank whose performance under the CRA had been deemed unsatisfactory may be prohibited from any future expansion in terms of either acquiring another bank or even opening a new branch.

I exploit the data on boundaries of assessment areas as a source of the identifying discontinuity. Following the approach of Black (1999), I construct a sample of census tracts that are in close proximity and are extremely similar to each other in terms of all observable characteristics. These include demographic variables on neighborhood racial and ethnic composition, median incomes, home ownership costs, house values, poverty levels, crime rates and credit scores. I then pick a subset of these similar census tracts such that roughly half of them belong to the assessment areas, and the others do not. After this, I focus on the lower-income group of census tracts. By comparing the loan approval rates inside and outside the assessment areas in this selected sample, I am able to interpret the observable difference (see Figure 1) as the causal impact of the CRA.

I then address the potential concerns that banks draw borders of their assessment areas in a strategic and nonrandom fashion. This is done by modeling the CRA impact on loan approval incentives as a function of the distance to the nearest bank branch. Drawing analogies with the classic measurement error problem is instructive here. Banks are unlikely to treat the CRA incentives as binary; instead, given that the assessment areas are usually drawn around bank branches, I expect banks to be somewhat more willing to approve loans for properties located closer to their branches. The further away one gets from the branch, the more it becomes likely that the loan will fall outside the assessment area and hence will add little to the bank's CRA record during the evaluation process.

I use the nonlinear Bayesian IV model to estimate the effect in question. Several considerations prompt the usage of this method. First, it allows me to account for the unobservable heterogeneity across loan applicants by making model coefficients random. Second, a linear probability model is only an approximation of the true nonlinear model in case of a limited dependent variable. Blundell and Powell (2004) demonstrate how this approximation can sometimes be quite imprecise. Finally, a probit model would avoid non-sensible predictions such as

Figure 1: Preliminary Evidence

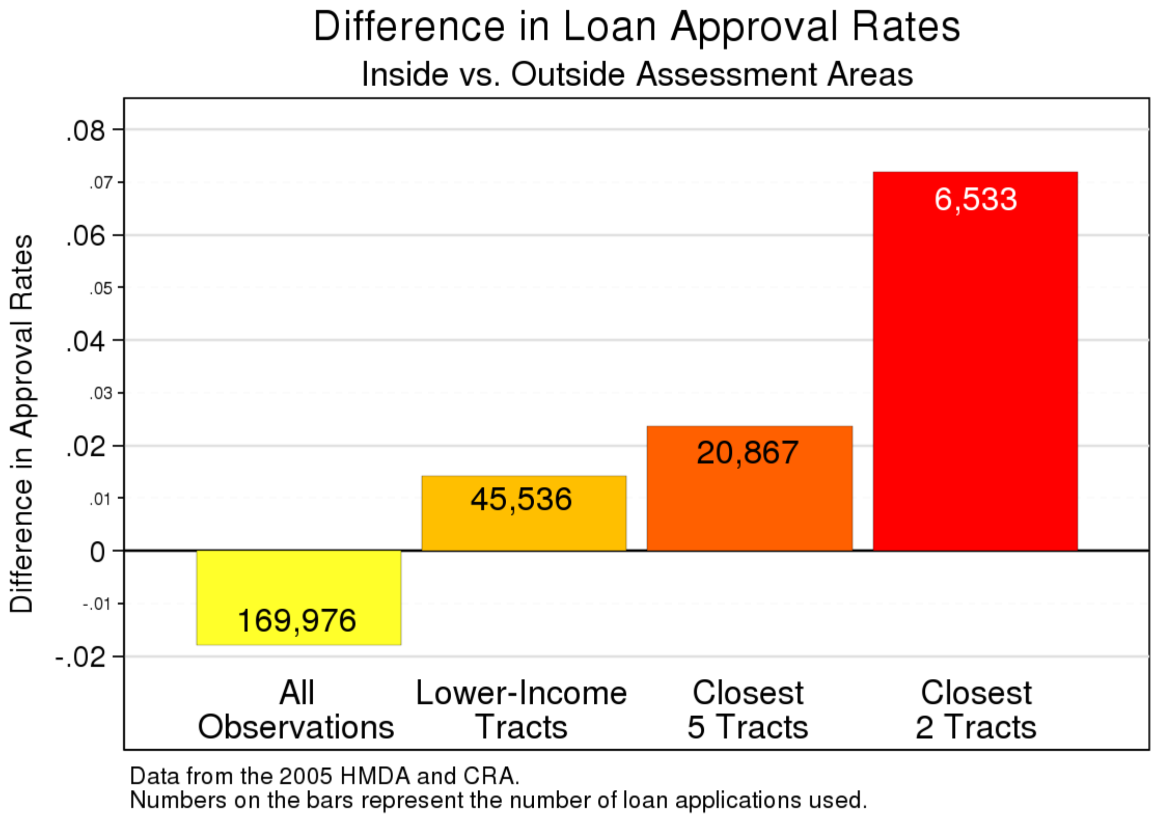

Samples "Closest 5 Tracts" and "Closest 2 Tracts" were obtained as a result of the application of the matching function, discussed in Section 4.1

probabilities outside the $[0, 1]$ range.

To address the second question about the quality of extra CRA-induced loans, I would ideally need to observe the performance of those loans during the last several years, which is something I cannot do given data availability constraints. I provide some indirect evidence by comparing foreclosure rates inside and outside the assessment areas in the matched sample of census tracts. By construction, the matched sample is extremely homogeneous with respect to all observables, except for the CRA eligibility. An observable difference in foreclosure rates would thus imply that the CRA is responsible for it.

I find that the CRA has a strong impact on banks' loan approval decisions. On average, there is a 33% higher chance that a loan application will be approved in case the application is CRA-eligible (i.e. the corresponding collateral property is located within a bank's assessment area). This effect is several times as large as the one suggested by a simple linear probability model that treats assessment area boundaries as given. It implies that the CRA had caused banks to approve some extra $498,720$ mortgage loan applications in 2005, and given the average loan size to be $\$278,000$, this amounts to almost $\$150$ bln. of additional lending.

I also demonstrate that the foreclosure rates are about 5.36 times higher inside assessment areas as compared to the census tracts that lie outside assessment areas. This suggests that almost 1 out of every 6 loans that had been induced by the CRA turned out to be non-performing. My foreclosure data is on a zipcode level and hence is somewhat less informative. The overall picture, however, seems to be clear. The CRA had created incentives for the banks to move down along the demand curve for mortgage loans. Some of these CRA-induced loans seem to had been given out to the borrowers who would not had been able to qualify for loans in the absence of this regulation. Higher foreclosure rates suggest that a considerable number of people who got these loans might had been better-off without them.

The remainder of the paper is organized as follows. Section 2 provides an overview of the mortgage origination industry and outlines the relevant regulations, notably the Community Reinvestment Act. Section 3 goes over the data sources that I draw upon in the paper and details the construction of the final dataset. I discuss my estimation strategy in Section 4. Results are presented in Section 5. I provide some evidence on the post-origination performance of these loans in Section 6, and Section 7 concludes.

2 Industry Background

2.1 Overview

Mortgage lending is a huge and busy industry with many participating sides. In 2005, some 35.5 million people applied for mortgage loans, and about 60% of those applications were approved. In California the corresponding numbers were 5.45 million and 51.2%, respectively. The average loan amount in 2005 was $\$183,000$ across the country, and $\$276,000$ in California. Figure 2, adapted from Bitner (2008) provides an schematic overview of the industry.

Figure 2: Mortgage Originations Industry

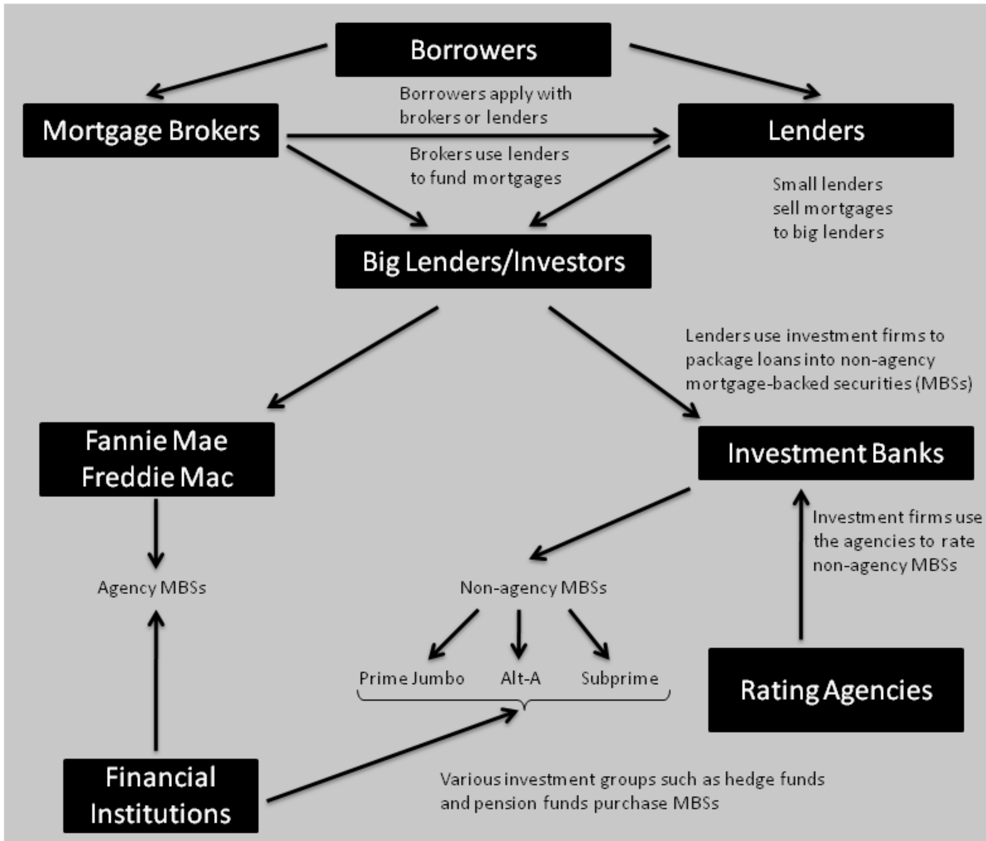

Source: Bitner (2008)

A home mortgage loan starts with a person who decides to purchase a house. Few buyers can afford to pay the tag price in a single installment. Instead, they choose to borrow most of the money (usually at least 80%) from a financial institution, using the house in question as a collateral. This type of borrowing is what I refer to as a mortgage loan.

There are three major types of mortgage lenders:

1. Depositary institutions, i.e. banks, savings institutions, thrifts, and credit unions;

2. Mortgage banking subsidiaries of depositary institutions (or of bank holding companies);

3. Independent mortgage banks.

A separate but an important party that took part in the lending process are the mortgage brokers. They are the intermediaries between a borrower and a lender that facilitate both the loan application and the origination processes. The difference between a broker and a lender's representative is that the former acts as an independent legal entity. Their primary revenue sources are the fees paid by the borrower (and in some cases commissions from lenders). It is not uncommon for a broker to work with several lenders simultaneously. According to Kleiner and Todd (2007), over $53,000$ mortgage broker firms were operating in the United States in 2004, and they were at least to an extent involved in the origination of more than two-thirds of all mortgages in that year. I will briefly revisit the question of brokers' importance in the next subsection.

The primary difference between lender types is in their source of funding. Depositary institutions (i.e. "banks") attract money in the form of deposits from their customers, and can use these funds for financing originated mortgage loans. Nondepositary institutions (subsidiaries and independent mortgage banks) primarily turn to securitization, which is the process of creating liquid securities out of a collection of illiquid assets.

All the originated mortgage loans that were intended for securitization can be divided in two groups: conforming and non-conforming. Conforming loans are such that satisfy certain criteria established by the GSE (government-sponsored enterprizes, i.e. "Fannie Mae" and "Freddie Mac"). These criteria include loan amount ceilings, lenders credit score floors and other factors, outlined in Keys, Mukherjee, Seru, and Vig (2010). All the other loans fall into the non-conforming category. GSEs turned conforming loans into agency MBSs (mortgage-backed securities), whereas non-conforming loans were turned into non-agency MBSs by other securitizers (usually mortgage banks).

Agency MBSs were considered to be extremely low-risk investments, since GSEs stood behind them and these institutions in turn had implicit guarantees from the U.S. government. These guarantees became explicit on September 7, 2008, when U.S. Treasury placed Fannie Mae and Freddie Mac into conservatorship.

The non-agency MBSs can be divided into three groups:

1. Prime jumbo – loans made to people with good credit ("prime" borrowers) but such that the loan amounts were too large to be conforming;

2. Subprime – risky loans made to people with lower credit scores below the cutoffs established by the GSEs;

3. Alt-A – loans made to borrowers with decent credit but who usually lacked complete documentation, moderately risky loans (more risky than prime but less risky than subprime).

Because non-agency MBSs did not entertain the implicit guarantees of their agency counterparts, securitizers at the investment banks like Bear Sterns and Lehman Brothers had to turn to rating agencies like Fitch and Moody's. The job of a rating agency was to assess the inherent risk of these new securities, and to assign them a rating based on their own internal criteria. This process was extremely important for the makers of the MBSs because many large investors had explicit guidelines that do not allow investing into very risky assets. To make MBSs marketable to a general pool of investors they had to be rated as sufficiently low-risk investments. Benmelech and Dlugosz (2009) is an excellent source on this matter.

The complete structure of the mortgage origination industry is therefore quite complex. I focus on the very first stage of the overall process, which is the loan approval decision by lenders (see Figure 3 below). I treat the actions of all other industry actors as given.

2.2 Regulations

Given the sheer size and importance of mortgage lending industry, one would not be surprised to find that it is subject to a number of federal regulations. One of the major concerns of lawmakers had been fighting discrimination practices, especially along racial dimensions. The Home Mortgage Disclosure Act and the Community Reinvestment Act are the two laws that had initially been passed with precisely this goal in mind.

The U.S. Federal Reserve Board's Regulation C implements the Home Mortgage Disclosure Act. Initially passed in 1975, its main purpose was to fight discrimination in mortgage lending. Regulation C requires almost every application for a home mortgage loan to be recorded and reported to the FFIEC (Federal Financial Institutions Examination Council) at the end of the year. I provide a detailed description of the HMDA data in Section 3.1.

The Community Reinvestment Act was passed in 1977. Its initial goal was to discourage "redlining" practices that had been in place at many depositary institutions at that age. Put sim-

Figure 3: Mortgage Originations — The Relevant Industry Part

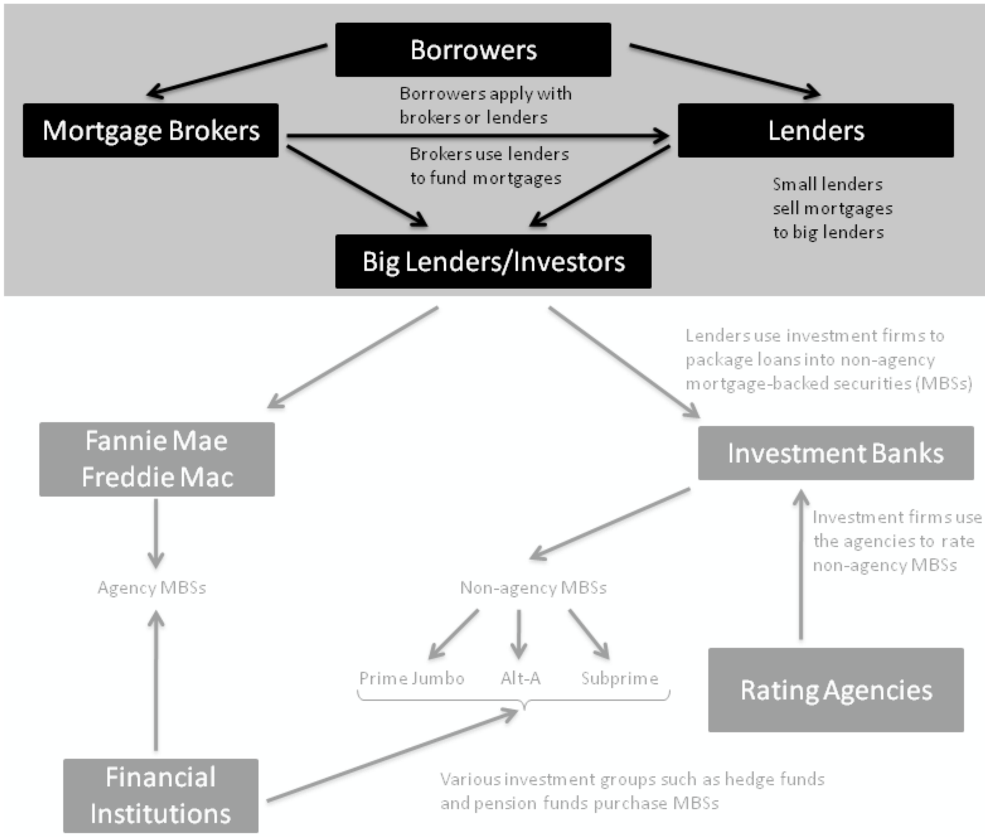

Adapted from Bitner (2008)

ply, redlining amounted to blanket refusals by the banks to lend in certain areas. This practice originated with the Federal Housing Administration (FHA) in the 1930s. The Home Owners' Loan Corporation (HOLC) created the "residential security maps" for the FHA. These maps were used by lenders for years afterwards to withhold mortgage loans from neighborhoods that were perceived as "unsafe". The Act "encouraged commercial banks and savings associations to meet the needs of borrowers in all segments of their communities, including low- and moderate-income neighborhoods" (Avery, Bostic, and Canner, 2000).

More formally, the low- and moderate income neighborhoods, commonly referred to either as LMI, or lower-income neighborhoods, are census tracts with median family income not exceeding 80% of the median income in the MSA to which the aforementioned census tracts belong. Only the "banks" are subject to the CRA.

All respondents are required to define their assessment area(s). A bank's CRA assessment area (AA) has to be a geographic area that is delineated by the bank. The delineation has to be approved by the bank's regulatory agency[2], which will later use it in evaluating the bank's record of meeting the credit needs of its community. When a bank proposes an assessment area that looks like it might violate some of the conditions listed below, regulators usually choose to talk to the bank and get a justification for any irregularities in the proposed definition, as opposed to citing it with a violation.

The area must consist of one or more contiguous political subdivisions, such as counties or cities. It must include neighborhoods in which the bank has its main office and branches, as well as the surrounding census tracts in which it originates a substantial portion of its loans. An AA must consist only of whole census tracts, may not reflect illegal discrimination, may not arbitrarily exclude low- and moderate-income tracts. A bank may adjust the boundaries of its AA to include only the portion of a political subdivision that it can reasonably expect to serve (CRA Reference, 2005).

One important feature of the AAs came up during my interviews with the industry experts. It is generally believed that the banks tend to rely more on mortgage brokers in marketing their services outside the assessment areas. The reason is clear: AAs are usually the areas around the bank branches, and hence people within AAs are in close proximity to the official banks' representatives – branch employees.

The regulation distinguishes between large and small banks based on their asset size (as of September 2005, the cutoff was at $1 bln). Given the bank size, the definitions of its assessment areas, and its activity during the year, a number of criteria is used to evaluate bank's performance under the CRA.

The lending test, which is of primary interest to me, considers the total number and amount of the approved loans by bank. It then looks how these numbers change if looked at from different angles. These include geographic distribution of loans (lower-income census tracts vs

[2] Every depositary institution is supervised by one of the four federal regulatory agencies. These are the Office of the Comptroller of the Currency (OCC), the Federal Reserve System (FRS), the Federal Deposit Insurance Corporation (FDIC) and the Office of Thrift Supervision (OTS).

higher-income tracts), distributions based on borrower characteristics, and proportion of loans that are made inside assessment areas. Table 3 in Appendix C, adapted from the CRA Reference (2005), provides other criteria that are relevant to the bank examiners. It also details how the bank's performance gets evaluated.

Other tests are also applied, especially to large banks, but the lending test carries the most weight in determining the bank's final CRA rating. The rating then becomes the part of the supervisory record for that bank. The possible ratings range from "Outstanding" (the best) to "Substantial Noncompliance" (the worst), as Table 3 in Appendix C shows. The CRA compliance record is taken into account by the regulators when a bank seeks to expand through merger, acquisition or opening a new branch. Also, the CRA ratings are public information, which can be accessed by community activists groups like ACORN. To sum up, banks can be expected to place considerable value on having a good CRA record. This may explain why the vast majority of the banks end up getting "Outstanding" or "Satisfactory" ratings.

3 Data

3.1 HMDA Loan Applications Database

I use several data sources for my estimation procedure. The single most important source is the HMDA data set (pronounced "humda"). It contains the vast majority of all home mortgage loan applications in the U.S. The potential mortgage originators (called "respondents" in the HMDA language) are the main subjects of the HMDA.

An observation in the dataset is a loan application, and a number of important characteristics of the application are available. These can be divided into 3 major groups:

1. Borrower's characteristics, such as race, gender, ethnicity and income.

2. Respondent's characteristics, such as name, type, address, parent company name (if applicable), supervisor's identity.

3. Loan characteristics, such as amount, property address (aggregated up to a census tract), various type measures (single or multi-family; conventional, FHA or VA; owner-occupied

or not; new or refinanced loan, etc). The important loan characteristics are the decision taken on the application, and the reasons for denial (if applicable).

The FFIEC aims to preserve anonymity of mortgage applicants by not disclosing the application date, which is rounded to a calendar year. A number of potentially interesting covariates are not present in the HMDA data. There is no information on the term structure of the loan (15-year mortgage or 30-year mortgage), as well as there is very limited information on the loan interest rate. HMDA respondents must report the difference between the loan APR (annual percentage rate) and the rate on Treasury securities of comparable maturity, as long as the spread is above the designated threshold. It is also not possible to tell apart the fixed rate and adjustable rate mortgages in the data.

On average, there were 31 millions of loan applications per year in the HMDA data (during 2000-2005). In Section 3.5, I detail the way the final dataset is constructed.

3.2 CRA Assessment Areas

The Federal Financial Institutions Examination Council also provides the CRA data. Each year every financial institution that is subject to the CRA must file a number of reports to the FFIEC. I use the definitions of institutions' assessment areas from the CRA disclosure reports. These are reported for each institution on a census tract level and can be linked to the HMDA loan applications directly institution by institution.

3.3 Foreclosure Scores

The Local Initiatives Support Corporation (LISC) provide a zip-code level foreclosure scores for California, as of March 2010. The calculations were performed using measures of subprime lending, foreclosures, and delinquencies from Lender Processing Services (LPS) loan-level post-origination performance data. The scores range from 0 to 100, they are relative to that of the neediest zipcode within the state, which is assigned a score of 100. Thus, if an area has a score of 50, it is estimated to be half as needy as the worst-off area. I also use the county-level data on foreclosure rates that is provided by the New York Federal Reserve, which is essential if one wants to be able to interpret these relative foreclosure scores quantitatively.

3.4 Other Sources

I also use several supplemental data sources. First, I use the Census 2000 data on racial and ethnic composition, home ownership costs, median family incomes, poverty rates and house values for every census tract in California. The Census Bureau also provides the information on latitude and longitude for each census tract. Next, the 2005 FDIC Summary of Deposits data contains a complete list of all bank branches with their addresses. I take the crime data from the Attorney General of California website and construct an annual crime rate index for the years 1999-2005. Finally, the Bureau of Labor Statistics makes available the CPI data, which allows me to express all nominal variables (annual income and loan amount) in 1999 dollars.

3.5 Dataset Construction

I choose to look at the loan applications in California in 2005. By looking within a single state, I abstract from inter-state variations in laws that regulate banks' operations. California was chosen for several reasons. First, it is one of the largest states, and home mortgage lending market in California makes up 14.8% of the overall U.S. market (on average between 2000 and 2005). Second, California has no metropolitan areas that span the boundaries of multiple states, unlike many other states. By looking at 2005 I concentrate on one of the last years before the mortgage crisis started to unfold.

I focus on a subset of mortgage lenders that are explicitly subject to the CRA. In the middle of 2000s only 1 out of 3 loan applications were reported by such lenders. Looking at the other two-thirds of mortgage applications would thus not allow me to identify the possible effects of the regulation.

I keep only conventional loans (no FHA or VA applications). I exclude loans that are not for single family owner-occupied homes, and that are secured by anything other than a primary lien. Finally, I only keep home purchase loans. Non-primary liens and non-single family loans in practice are more likely to be associated with real estate speculative purchases, especially during the years of interest. There are also reasons to believe that borrowers' characteristics are more accurately reported for home purchase loans (Bhutta, 2008).

The HMDA data reports decision taken on each application, and there are 10 different decisions that can be reported. I drop applications with decision reported as "application withdrawn", which are thought to be associated with indirect lending through mortgage brokers rather that directly through banks of interest. I then construct the loan approval indicator which equals 1 if loan was approved (whether originated or not). Borrowers can apply for loans with different banks at the same time, and there is no way to identify two different loan applications with a single borrower in the data. My primary interest, however, is in banks' approval decisions, and it is quite possible that different banks assess the same person's application differently.

4 Estimation

4.1 Identification Strategy

I start parsimoniously by estimating the following equation:

$$y_i = AA_i \cdot \beta + x_i'\gamma + \varepsilon_i, \tag{1}$$

where $i = 1, \ldots, n$ indexes loan applications, y_i is the loan approval indicator, x_i is a collection of observable covariates that impact bank's decision to approve the loan (loan amount, borrower's income and demographics), AA_i is an indicator for the mortgaged property being located inside bank's assessment area.

The main object of interest is the estimate of β. A nonzero β implies the existence of systematic difference in the way banks treat loan applications inside and outside their assessment areas. If β is positive, then one can argue that the loan has a higher chance of getting approved if it is within a bank's AA.

Direct estimation of the equation (1) would not allow me to interpret the estimate of β causally, however. The problem of omitted variable bias is likely to be paramount in my setting. For example, HMDA data has no information on applicant's credit score. Lenders most likely treat potential borrowers with credit scores under 600 quite differently from those with credit score above 750 – applications from the latter are probably more likely to be approved, other things equal. When an important covariate is missing from the regression, most standard methods (OLS, MLE, etc) fail to deliver consistent estimates.

An ideal, albeit unfeasible, solution to the problem would be to obtain a randomized sample of loan applications from tracts inside and outside assessment areas. The next best alternative, which I implement, is similar in spirit to that of Black (1999). Specifically, I use the boundaries of banks' assessment areas as a source of quasi-random variation. I put together a collection of census tracts around the borders of assessment areas that appear to be quite similar in all observable characteristics. The only major observable difference is that roughly a half of those tracts falls into some bank's assessment area, whereas the other half does not. Later in Section 4.2 I turn to the question of potential nonrandomness of the boundaries.

I construct the matching function that operates on a census tract level. It computes a "distance" measure for each census tract in the sample to every other tract. It then selects a predetermined number of neighboring tracts in terms of the distance measure, and by changing the number of neighbors I am able to look as closely to the assessment area boundary as I wish to.

The distance measure between tracts A and B is a weighted average of:

- the geographic (also known as the "great circle") distance from the center of tract A to the center of tract B;
- and L^2-distance between standardized[3] values of all socio-economic characteristics.

I use an extensive collection of the socio-economic variables for the matching procedure. Using the 2000 Census data I construct measures of median incomes, racial and ethnic composition, house values, home ownership costs and poverty levels for every census tract. I also construct an index of the annual crime rates (on the county level) using the data on the number of various crimes from the California Attorney General's website. Finally, I use 2000-2004 HMDA data to construct a county-level measure of average credit scores. When a loan application is denied, most lenders list up to three reasons for denial, and "poor credit history" is one possible reason. I compute the average number of loans that had been denied due to poor credit

[3] Standardization involves subtracting the mean of each variable from every value it takes and dividing the result by its standard deviation. This ensures that the units of different variables are irrelevant, otherwise differences in income would swamp differences in, say, poverty rates. My procedure explicitly guarantees that every covariate will receive the same weight in the overall criterion.

history and use it as a proxy for the true distribution of credit scores. Other studies had followed this approach before (see Ergungor (2007) and the references therein).

The matching algorithm makes an implicit assumption that most census tract characteristics change smoothly as one goes from one adjacent tract to another. This condition is similar to the standard continuity condition made in most regression discontinuity studies (see, for example, Imbens and Lemieux (2008), assumption 2.2). The only factor that changes discontinuously is thus the belonging of a given census tract to a bank's assessment area.

Figure 4 illustrates the results of the matching procedure. The complete list of matching variables appears on the vertical axis. Every variable was first standardized (see footnote 3 on page 15). Then the mean values for each variable across census tracts inside and outside assessment areas were computed. Each dot on the plot represents the difference between these means.

I plot the differences in means together with the ± 2 standard deviation bands around 0. If the matching procedure worked flawlessly, then all the dots would have been aligned into a vertical line at 0. This would indicate that means for each variable are the same inside and outside assessment areas. Although there seems to be some leftover variation in the characteristics of the census tracts even after matching, it is clear that the difference is always fairly small and all deviations lie within the 1 standard deviation bands. Using normal approximation, a back-of-the-envelope calculation suggests that the p-value for the test of equality of means for any given matching variable is about 50%.

Perhaps the largest deviation from the mean can be observed in the tract-level demographic characteristics (the top several dots in each column). I would argue, however, that it is more important to control for these characteristics at the individual loan level. Most of these demographic covariates are available at precisely this level of disaggregation.

As the result of the application of the matching algorithm, I obtain 2 matched samples labeled "Closest 5 Tracts" and "Closest 2 Tracts". The numbers mean that for every tract the function finds the corresponding number of "closest" neighbors in the sense defined by the matching criterion. The "Closest 5 Tracts" sample has more observations in it, while the "Closest 2 Tracts" contains only the "most similar" observations. I use both of these samples in most of my estimation as a robustness check. Equation (1) now can be estimated using the obtained samples.

Figure 4: Matching Results

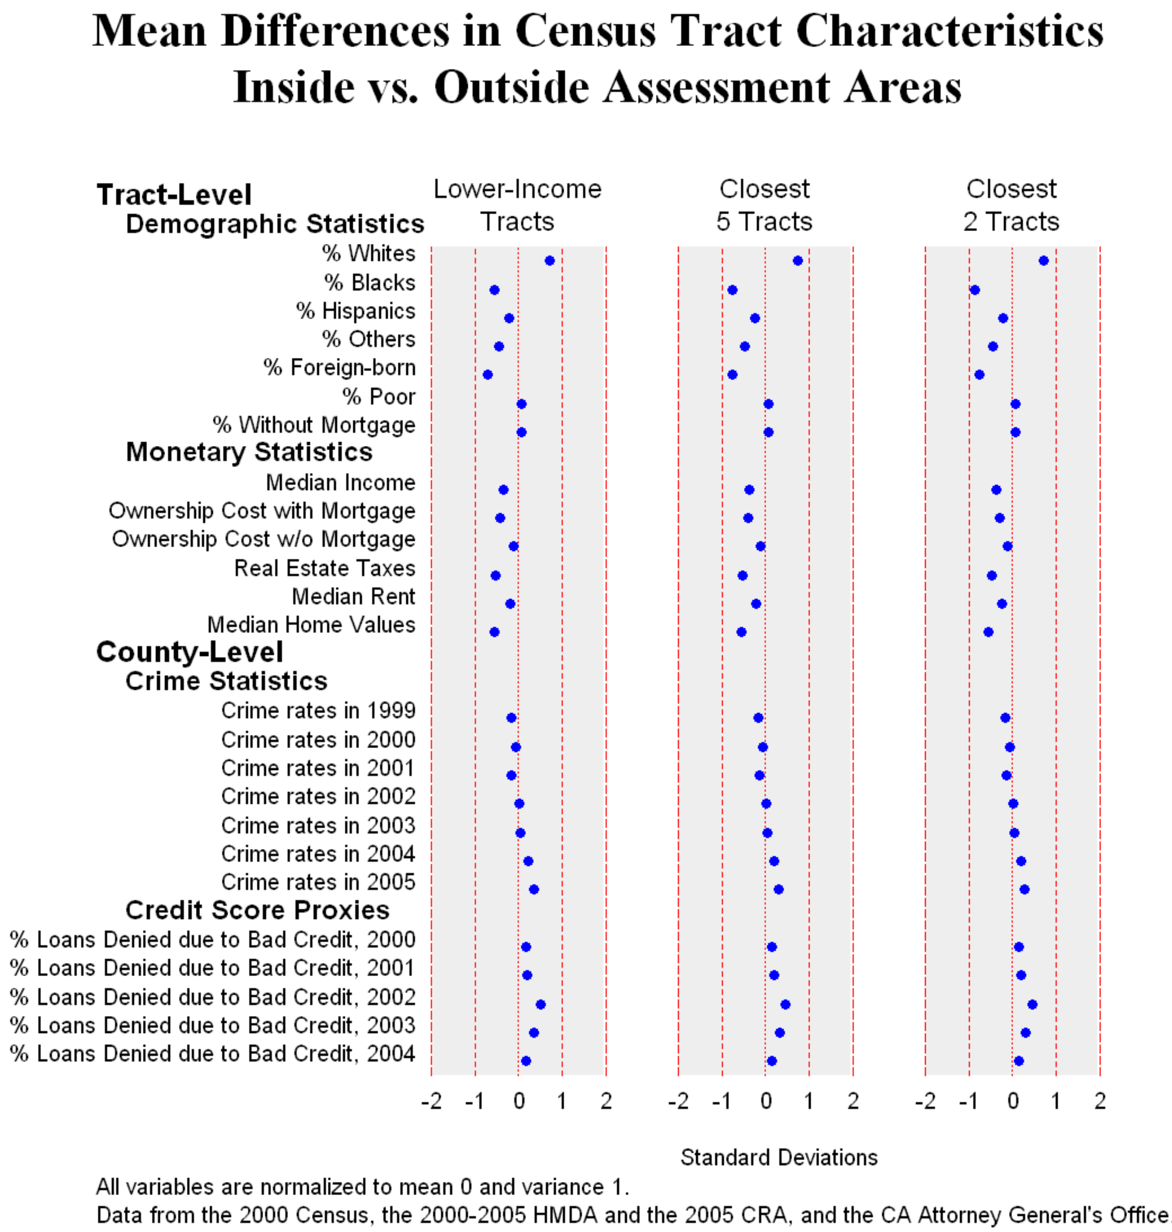

More importantly, the results of this estimation can be interpreted causally, i.e. a positive and significant estimate of β allows me to infer that if a borrower applied for a mortgage loan within a bank's assessment area, there is a higher chance this loan will end up being approved by the lender, other things being held equal.

4.2 Two-Stage Least Squares

One obvious concern with the previously outlined identification strategy is that it relies on an implicit assumption that banks draw their assessment areas boundaries nonstrategically. The matching procedure is designed to select a particular subset of the boundaries such that this assumption would essentially hold by construction. The fact that banks and regulators spend considerable amount of time negotiating about the assessment areas delineation suggests, however, that this issue is potentially quite important and hence must be addressed.

I use the information on the location of bank branches to address this concern. For every census tract, I construct a measure of distance to the nearest bank branch. Methodologically it amounts to instrumenting for the assessment area indicator with this distance measure. A good instrument has to satisfy two conditions: relevance and validity.

A relevant instrument is such that it is correlated with the endogenous variable in question. This can be tested for directly, but intuitively there are reasons to believe that the relevance condition holds. Banks are not allowed to carry out explicit discrimination in their lending based on geography. In fact, that had precisely been the reason for passing the original CRA regulation in the first place. However, the closer the collateral property is to the bank branch, the higher is the chance the property will fall inside the bank's assessment area, because the branches are always located inside the AAs.

The validity condition means the instrument should only affect the outcome variable indirectly via the endogenous factor. Otherwise the instrument should be included among the other regressors. There is no way to test for this condition explicitly, but the discussion in the previous paragraph suggests that validity also holds. Banks get no credit for approving loans that are closer to their branches per se. They might, however, get credit from the regulators via the CRA channel, and this is precisely the indirect influence of the distance on the loan approval decision that makes my instrument valid.

Formally, I estimate the following two equations:

$$y_i = AA_i \cdot \beta + x_i'\gamma + \varepsilon_{i,2}, \tag{2}$$

where the notation is exactly the same as in (1). However, I add another equation now:

$$AA_i = dist_i \cdot \delta_1 + x_i' \delta_2 + \varepsilon_{i,1}, \tag{3}$$

where $dist_i$ is the measure of distance from the assessment area boundary to the nearest branch, and it is bank-specific.

These two equations can be easily estimated via the two-stage least squares method. I provide the table with the estimation results in Section 5.2. This model, however, is not quite correctly specified, because the dependent variable is an indicator. A linear probability model does not account for the fact that predicted values generated by it have to lie within $[0, 1]$ interval. Also, 2SLS methods do not allow to account for potentially important unobservable heterogeneity among loan applicants. I hence turn to the nonlinear Bayesian IV model to address these concerns.

4.3 Nonlinear Bayesian IV

I change the main equation (2) and write it as:

$$y_i^* = AA_i \cdot \beta + x_i' \gamma + \varepsilon_{i,2}, \tag{4}$$

where AA_i is endogenous. The instrument equation (3) is unchanged:

$$AA_i = z_i' \delta + \varepsilon_{i,1}, \tag{5}$$

where $z_i = (dist_i, x_i)$, $\delta = (\delta_1, \delta_2)$, and where y^* is not observable. Instead, I observe $y_i = 1\{y_i^* \geq 0\}$. The fact that y^* is unobservable implies the need to apply a limited dependent variable technique to (4), most commonly a probit model. And if ε_1 and ε_2 are correlated (i.e. if AA is endogenous), (4) and (5) must be estimated simultaneously.

Dealing with endogeneity in a nonlinear model correctly had been previously addressed by various studies including Geweke, Gowrisankaran, and Town (2003) and Blundell and Powell (2003). To my knowledge, however, there is no single ubiquitous method to address this problem. Controlling for the unobserved heterogeneity is also nontrivial in this setting. I apply the techniques developed in the Bayesian literature to address both of these issues at once.

Specifically, I supplement the standard linear Bayesian IV method with the data augmentation step. The data augmentation procedure is best to be thought of as a vehicle that allows me to convert a nonlinear model into a linear one. It then becomes possible to apply standard linear IV methods to the transformed model. A complete exposition of the Bayesian approach to inference is well beyond the scope of this paper. Some excellent sources on this question include Rossi, Allenby, and McCulloch (2006) and Geweke (2005). The general idea, however, is worth outlining.

I start with the data at hand and with a prior distribution on parameters. This prior distribution represents my initial beliefs about parameters' distributions and is in practice mostly chosen for computational convenience. The data allows to form the likelihood function, and together the product of the prior and the likelihood yields the posterior distribution for parameters. The Bernstein von-Mises theorem ensures that means of posterior distributions are asymptotically equivalent to MLE estimates. In finite samples, one can select a fairly diffuse prior (think normal distribution with a really large variance) so that it would contribute almost nothing to the posterior, which in this case will be mostly driven by the likelihood.

With the exception of several simple models, the posterior distributions are in general too complicated to be analyzed directly. The standard approach is to use Markov chain Monte Carlo (MCMC) methods to obtain a random sample of draws from the posterior, and then read the required information off these draws. The Gibbs sampling technique simplifies the process of making such draws by splitting the whole parameter space into nonoverlapping blocks. I choose blocks so that conditional on all other blocks being held constant, the posterior of any given block is of a known form, and so that it is straightforward to make random draws out of it. By alternating between blocks, I come up with a sequence of draws that converges to the draws from the joint posterior for all parameters[4].

[4] A very vivid example is given in Rossi, Allenby, and McCulloch (2006): suppose I want to draw from a bivariate normal distribution, but I can only make draws from univariate normals. For a bivariate normal, the conditional distribution of any component given the other is also normal. The Gibbs sampling approach would imply making alternating draws from two univariate conditionals, which eventually produces a sequence of pairs of draws. This sequence in turn is asymptotically equivalent to a sequence of random draws from a bivariate normal distribution, which is precisely what I needed in the first place.

Apart from assuming priors, I need to impose a distributional assumption on the error terms:

$$\begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix}\right),$$

where $\sigma_{12} \neq 0$ and which implies that $\varepsilon_1 \mid \varepsilon_2$ and $\varepsilon_2 \mid \varepsilon_1$ are also normal.

Given that y^* is a linear function of ε_2, it also follows normal distribution (and this stays true if one conditions on ε_1). I treat $\{y_i^*\}_{i=1}^n$ as an extra set of parameters that can be simulated. The simulation must incorporate the observed information on the sign of each y^* (i.e. y), so draws must be made from the truncated normal distribution. I use the mixed rejection algorithm developed in Geweke (1991), the relevant part of which is given in Appendix B.

The whole Gibbs sampling scheme can be summarized as:

$$\{y_i^*\}_{i=1}^n \mid \beta, \gamma, \delta, \Sigma, y, AA, z \tag{6}$$

$$\beta, \gamma \mid \delta, \Sigma, \{y_i^*\}_{i=1}^n, AA, z \tag{7}$$

$$\delta \mid \beta, \gamma, \Sigma, \{y_i^*\}_{i=1}^n, AA, z \tag{8}$$

$$\Sigma \mid \beta, \gamma, \delta, \{y_i^*\}_{i=1}^n, AA, z, \tag{9}$$

and I detail each of those four steps in Appendix A.

The model is complete after I specify the priors:

$$\begin{aligned} \begin{pmatrix} \beta \\ \gamma \end{pmatrix} &\sim N\left(\begin{pmatrix} \overline{\beta} \\ \overline{\gamma} \end{pmatrix}, A_{\beta\gamma}^{-1}\right), \\ \delta &\sim N\left(\overline{\delta}, A_\delta^{-1}\right), \\ \Sigma &\sim IW\left(\upsilon_0, V_0\right), \end{aligned}$$

where $\left(\bar{\beta}, \bar{\gamma}, \bar{\delta}\right)$ are priors' means, $A_{\beta,\gamma}^{-1}$ and A_δ^{-1} are prior variance matrices. In the Bayesian literature it is somewhat more customary to work with the inverses of the variance matrices, which are referred to as precision matrices, and hence $A_{\beta,\gamma}$ and A_δ are priors' precisions. IW stands for inverse Wishart distribution, which is essentially a generalization of χ^2-distribution to the space of positive-definite matrices (instead of positive integers), υ_0 is its scale parameter, and V_0 is its location parameter.

I choose very uninformative priors and set $\bar{\beta}, \bar{\gamma}, \bar{\delta}$ all equal to vectors of zeros of corresponding dimensions, $A_{\beta,\gamma} = A_\delta = 0.01 I_{k+1}$ (where $k = \dim[w]$), $\upsilon_0 = 3$ and $V_0 = 0.01 I_2$ (and where I is the identity matrix). This ensures the priors are considerably “spread out”, so that the shape of the posterior is mostly driven by the likelihood function. The choice of functional forms for the priors is motivated by the computational considerations: given the Gibbs sampler blocks in (6), posteriors on β, γ and δ will also be normal, and posterior on Σ will also be inverse Wishart. See Appendix A for details.

5 Results

5.1 Linear Probability Model

Table 1 presents the estimation results for the linear probability model defined in (1). The dependent variable is the loan application approval indicator and the key regressor of interest is the CRA assessment area dummy. I will refer to this estimate as “the CRA effect”.

The first column uses all available observations, i.e. it naïvely runs (1) on all the data at hand. I am therefore not surprised to find a small negative CRA effect, although it would be hard to argue for a causal interpretation of this number. The second column, which involves only observations from low- and moderate-income census tracts — i.e. tracts that are explicitly considered important by the CRA — provides a glimpse of a different picture. I would still not go as far as to interpret the number causally though, given the wide array of possible confounding omitted variables and effects.

In the next two columns I am only using observations from census tracts identified as similar by the criterion described in Section 4.1. As the number of neighboring tracts decreases, I am left with fewer and fewer observations (sample size drops from $20,867$ to $6,533$), but in return I am able to argue that the remaining observations come from tracts that are “more and more similar” to each other. The striking observation is that the estimate of β remains quite significant, and it increases as I get closer to the boundary.

In a perfect RDD setup, all other loan-level covariates would not have been significant. Since I do not observe the “forcing variable” directly (which is the distance from the property to the boundary of the assessment area) I cannot expect this property to hold. In return, I obtain a couple of interesting findings. Table 1 demonstrates that, other things being held equal, banks

Table 1: Linear Probability Model for Loan Approval

		Lower-Income Tracts			High-Income Tracts	
	All Data	All	Closest 5	Closest 2	All	Closest 5
Loan in Assessment Area	-0.015***	0.013**	0.030***	0.068***	-0.024***	-0.005
	(0.002)	(0.004)	(0.007)	(0.013)	(0.004)	(0.005)
Loan size, 100k	-0.001*	-0.003*	-0.002	0.002	-0.003***	-0.004***
	(0.000)	(0.002)	(0.002)	(0.004)	(0.001)	(0.001)
Annual Income, 100k	0.001	0.010***	0.007	0.011**	0.002**	0.004*
	(0.001)	(0.003)	(0.004)	(0.004)	(0.001)	(0.001)
Applicant Female	-0.006*	0.004	0.004	-0.000	-0.008*	-0.009
	(0.002)	(0.004)	(0.006)	(0.011)	(0.004)	(0.006)
Applicant Not White	-0.024***	-0.023***	-0.021***	-0.017	-0.032***	-0.037***
	(0.002)	(0.004)	(0.006)	(0.011)	(0.004)	(0.005)
Applicant Hispanic	-0.088***	-0.092***	-0.088***	-0.085***	-0.054***	-0.058***
	(0.002)	(0.004)	(0.006)	(0.011)	(0.004)	(0.006)
Has a Co-Applicant	0.052***	0.048***	0.044***	0.035**	0.053***	0.051***
	(0.002)	(0.004)	(0.006)	(0.011)	(0.004)	(0.005)
Constant	0.799***	0.785***	0.769***	0.724***	0.864***	0.845***
	(0.003)	(0.006)	(0.010)	(0.017)	(0.005)	(0.008)
Number of obs.	169, 964	44, 546	20, 867	6, 533	47, 039	24, 393

Dependent variable: loan approval indicator. Data from the 2005 HMDA and CRA. Standard errors in parentheses.

* $p< 0.05$, ** $p< 0.01$, *** $p< 0.001$

tend to approve loan applications from Hispanics less willingly, and the magnitude of the effect really does not change much no matter how finely I "slice" the data. Namely, there is about 8% lower chance that a loan will be approved if it is coming from a Hispanic applicant, controlling for all observable differences and using the aforementioned identification strategy.

The last two columns estimate the same specification on a different subset of data. Column 5 uses all the data from the "rich" census tracts (defined as tracts with median income of at least 120% of the corresponding MSA median income). I then apply the matching function described above in Section 4.1 to these "rich" tracts, and obtain a matched subset of them. The numbers in the last column come from estimating the equation (1) on this subsample. The results are intuitive: since the CRA does not explicitly reward banks for extra lending in the high-income areas, the CRA effect is either small and negative (if one looks at the whole "rich" sample) or indistinguishable from zero (if one looks only at a homogeneous subset of it). This suggests that the estimation is indeed "picking up" the CRA effect of interest.

Several additional specifications (not shown) that included interactions of the CRA effect with the borrowers' demographics were also estimated. The interaction terms were found to be quite insignificant, which is consistent with the conjecture that banks value compliance with fair lending laws considerably. For example, a significant interaction term of the CRA effect with the gender indicator would suggest that banks treat male and female applicants differently. The fair lending laws explicitly outrule such behavior on the bank side.

5.2 Two-Stage Least Squares

I next turn to the two-stage least squares estimation, modeling the CRA effect as a function of the distance to the nearest bank branch. Table 2 presents the estimation results.

The results are qualitatively quite similar to the ones of Section 5.1. However, the CRA effect goes up in magnitude by a factor of 3.5 or more, depending on which sample is used for comparison. This clearly suggests that accounting for the nonrandomness of the assessment area(s) boundaries is important. This result is also consistent with the measurement error interpretation of the endogeneity problem: failure to account for the measurement error is known to bias the estimates towards zero.

Table 2 thus suggests that the model with endogeneity is a preferrable way of looking at the data. Given the discussion in Section 4.2, the correct way to proceed is to use the nonlinear Bayesian IV method covered in Section 4.3. I turn to the results of this method in the next subsection.

5.3 Bayesian IV

Given the model setup, all the posteriors on the main equation (4) coefficients will be normal. Figure 5 below summarizes the results of the MCMC procedure. I use the larger of the matched samples, the one called "Closest 5 Tracts", and I do $100,000$ MCMC draws. Of these, the first 10% are discarded as "burn-in" – there are reasons to believe the information in those draws is strongly influenced by the choice of the initial conditions. I then "thinned" the sample by

Table 2: 2SLS Linear Probability Model for Loan Approval

	Lower-Income Tracts	Closest 5 Tracts	Closest 2 Tracts
Loan in Assessment Area	0.557***	0.343***	0.247***
	(0.036)	(0.031)	(0.038)
Loan size, 100k	0.031***	0.011***	0.003
	(0.003)	(0.003)	(0.004)
Annual Income, 100k	−0.002	0.007	0.011**
	(0.003)	(0.004)	(0.004)
Applicant Not White	0.028***	−0.010	−0.009
	(0.006)	(0.007)	(0.012)
Applicant Female	0.003	0.004	−0.003
	(0.005)	(0.007)	(0.012)
Has a Co-Applicant	0.010	0.030***	0.024*
	(0.005)	(0.006)	(0.011)
Applicant Hispanic	−0.090***	−0.100***	−0.082***
	(0.005)	(0.006)	(0.011)
Constant	0.377***	0.505***	0.582***
	(0.028)	(0.027)	(0.034)
Number of obs.	44,546	20,867	6,533

Dependent variable: loan approval indicator. Data from the 2005 HMDA and CRA. Distance to nearest bank branch used as an instrument for loan being in assessment area. Standard errors in parentheses. * $p< 0.05$, ** $p< 0.01$, *** $p< 0.001$

keeping only every 9th of the remaining 90,000 draws. This breaks the serial dependence in the chain that gets introduced naturally by the nature of the Gibbs sampling procedure.

The vertical dimension of the picture is irrelevant, I plot credible sets (which is the Bayesian term that corresponds to that of classical confidence intervals) at different heights so that element would not overlap. The thin vertical line indicates the zero value: if a credible set contains zero, one could think of a corresponding coefficient being insignificant in a classical sense. Thicker portions of lines indicate 50% credible sets, thinner portions represent the 95% sets and the thickest dots represent posterior means. For a normally distributed random variable, a 50% confidence interval is roughly equal to a ± 1 standard deviation bound, and a 95% confidence interval is almost the same as the ± 2 standard deviations bound.

The first thing to notice is that the credible sets for the CRA effect are positive and quite

Figure 5: MCMC Results

far from zero. In fact, the 95% credible set for the CRA effect is $[0.828, 1.073]$, with the mean at 0.952. These numbers are not directly comparable with the ones from the linear probability models, but it is still instructive to see that the posterior is fairly tight.

A credible set that lies to the right of the dashed vertical line indicates that the corresponding factor contributes positively to the loan approval score in the equation (4). Thus the fact that the credible set for the Hispanic borrower indicator contains only negative values suggests again that, other things equal, Hispanic borrowers have a harder time getting their applications approved. Overall, the qualitative results from the Bayesian IV model are quite similar to those from its linear probability approximation discussed in Section 5.2.

Figure 6 plots the CRA marginal effect as predicted by the Bayesian IV model. The marginal effect should be thought of as the incremental probability of loan approval. Take two identical loans such that their only difference is in that one is subject to the CRA and the other is not The CRA marginal effect demonstrates the difference between the chances each of these loans will be approved as predicted by the model.

Because the model is nonlinear, the picture would look slightly different for every observation. Moreover, since β is stochastic, it also looks different for every MCMC draw. I preserve the stochastic part of β and compute the marginal effect for an average loan in the sample, i.e. I replace the values of x_i with its column means.

The dashed maroon line represents the normal approximation of the histogram, using the mean and the standard deviation of the underlying actual values. The mean CRA marginal effect is 0.33, which is quite similar with the values predicted by the 2SLS linear probability model from Section 5.2. Thus the linear approximation was in fact quite accurate.

A simple back-of-the-envelope calculation of the overall magnitude of the CRA effect is worth doing. The average loan approval rate in the sample is about 0.76, which, interacted with the average CRA marginal effect of 0.33 yields 0.25. That is, almost every 4th approved loan had been approved due to the CRA-induced incentives. In the sample of $20,867$ loans this translates into approximately 5200 extra loan approvals. Given that the average loan amount is $\$276,000$, it amounts to $\$1.44$ billions of extra loans given out.

I take these numbers one step further and get an estimate of the CRA effect on lending in the whole state of California. I assume that the CRA marginal effect stays at 0.33, which may not be completely justifiable, but is still instructive. Carrying out the same steps as in the previous paragraph suggests that a total of $498,720$ loans had been generated in the whole state of California in 2005 by the CRA incentives. This translates into almost $\$150$ billions of additional loans. Generalizing this calculation for the rest of the country does not seem possible given various unobservable state-level specific factors.

Figure 6: The CRA Marginal Effect

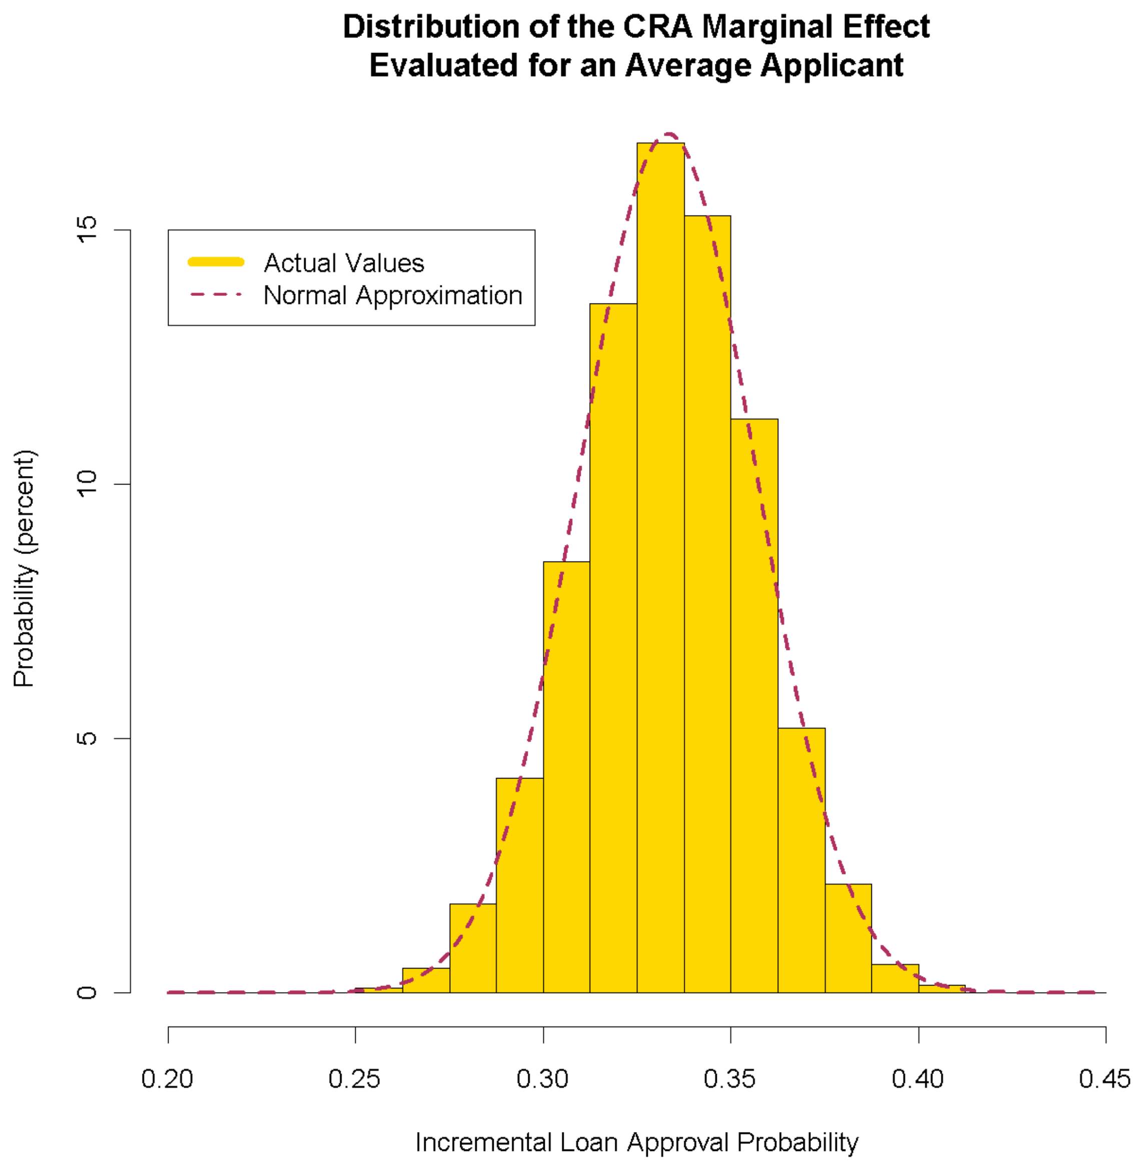

6 Evidence on the Performance of Loans

Having established that the CRA had made the banks approve more mortgage applications, it is natural to question the post-origination performance of these extra loans. A simple and instructive way to do this is to compare foreclosure rates in the matched samples inside and outside the assessment areas. Figure 7 plots the kernel densities of foreclosure scores inside and outside assessment areas in the "Closest 5 Tracts" sample.

Figure 7: Distributions of Foreclosure Scores

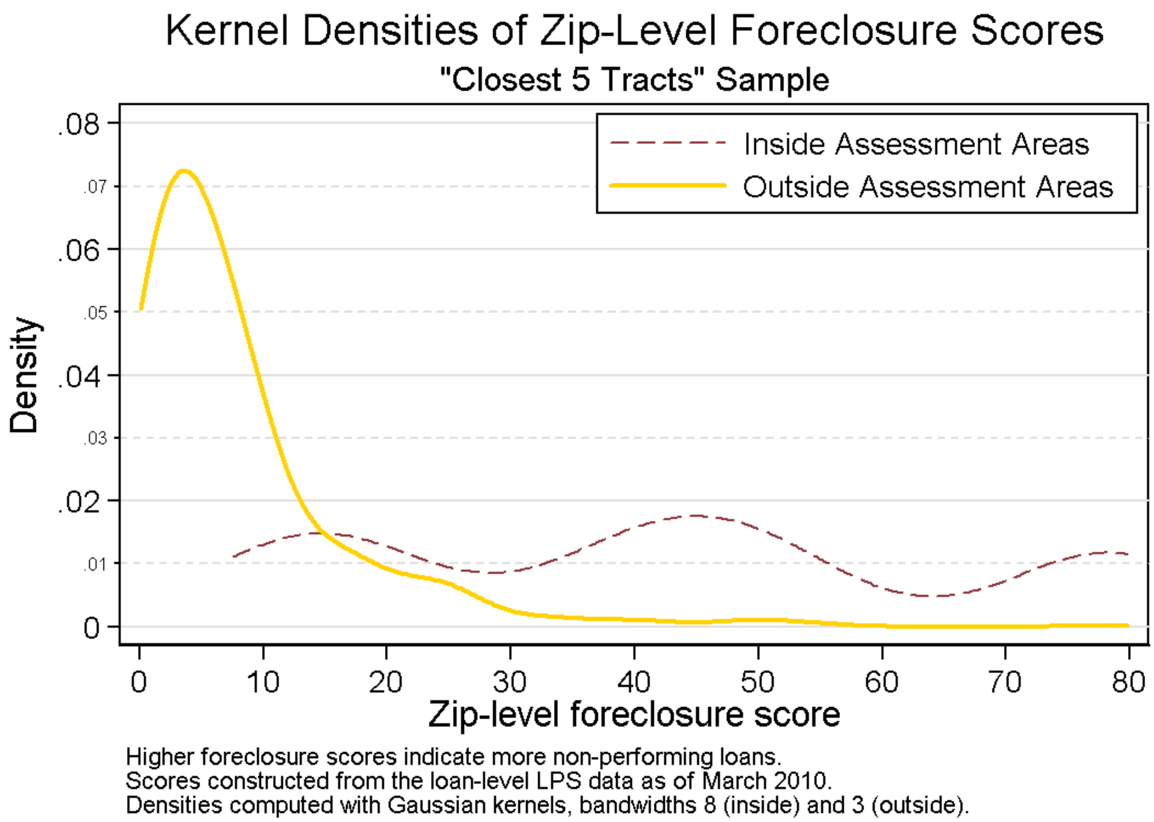

As detailed in Section 3.3, higher values on the x-axis indicate that more loans were in foreclosure in a given zipcode. It is clear that the majority of the zipcodes inside CRA AAs have foreclosure scores under 20 (and in fact, the average score outside AAs is 7.93). At the same time, the scores for zipcodes that lie inside assessment areas skewed to the right, with an average of 42.52. This suggests that, on average, the overall situation with foreclosures is 5.36 times worse[5] inside the CRA assessment areas.

Using the county-level data on foreclosure rates, I can do a quick estimate of the magnitude of the effect. The average foreclosure rate outside assessment areas is 2.96%[6]. Assuming that the 5.36 factor of proportionality is correctly pinned down by the zipcode-level data, I conclude

[5] Since the foreclosure scores are relative, simple division suggests this result: $42.52/7.93 \approx 5.36$

[6] I obtain this number using the county-level foreclosure rates data which is available from the New York Fed website.

that on average the foreclosure rates inside assessment areas are 15.87%. If these rates are the same among both the CRA-eligible and CRA-ineligible loans, it suggests that 1 out of every 6 CRA-induced loans had turned out to be non-performing. But even without these extra assumptions, the more-than-fivefold difference in average foreclosure scores inside and outside CRA-eligible areas is a striking indication that some of the extra CRA-generated loans did not perform well once things turned south.

7 Conclusion

In this paper I look at the Community Reinvestment Act and assess the magnitude of its input to the recent mortgage default crisis. To do this, I first establish an existence of a causal relationship between the CRA and extra loan approvals by the financial institutions that are subject to this regulation. I then demonstrate that the foreclosure rates are higher in the CRA-eligible areas, suggesting that some of these extra loans had recently failed to perform.

I use the ideas developed by the regression discontinuity literature to come up with a sample of observations that would allow me to identify the causal effect in question. The CRA makes its subjects define assessment areas, which are the proxies for their primary markets of operations. I look at the small subsample of loans located in census tracts along the boundaries of the assessment areas to come up with an identifying discontinuity.

I then address the concern that banks draw the maps of their assessment areas in a strategic function by using the instrumental variables approach. Specifically, I use the distance from the boundary of the assessment area to the nearest bank branch as an instrument for the CRA "treatment". I employ a nonlinear Bayesian IV model to address the endogeneity concern in a probit setting and to account for the unobservable heterogeneity among the borrowers.

I find that the CRA has a strong impact on the banks' loan approval decisions. Specifically, a CRA-eligible loan has on average a one-third higher chance of getting approved, other things equal. This suggests that about $\$150$ billions of mortgage lending in California in 2005 had been inspired by the CRA incentives.

To fully address the question of the CRA input to the mortgage default crisis, I would need to observe the post-origination performance of these extra loans. Although indirect evidence from foreclosure rates suggests that about 16% of these extra loans had failed to perform, a richer dataset is needed for a more precise answer. This seems like a promising topic for further research.

References

AVERY, R., R. BOSTIC, AND G. CANNER (2000): "The Performance and Profitability of CRA-Related Lending," *Economic Commentary*.

BAIR, S. (2008): "Did Low-income Homeownership Go Too Far?," in *Prepared remarks for Conference before the New America Foundation.* Available at http://www.fdic.gov/news/news/speeches/archives/2008/chairman/spdec1708.html.

BAJARI, P. L., S. CHU, AND M. PARK (2009): "An Empirical Model of Subprime Mortgage Default from 2000-2007," University of Minnesota Working Paper.

BENMELECH, E., AND J. DLUGOSZ (2009): "The Alchemy of CDO Credit Ratings," *Journal of Monetary Economics*, 56, 617–634.

BHUTTA, N. (2008): "Giving Credit where Credit is Due? The Community Reinvestment Act and Mortgage Lending in Lower Income Neighborhoods," *Board of Governors of the Federal Reserve Finance and Economics Discussion Paper*, 61.

BITNER, R. (2008): *Confessions of a Subprime Lender: An Insider's Tale of Greed, Fraud, and Ignorance*. Wiley.

BLACK, S. (1999): "Do Better Schools Matter? Parental Valuation of Elementary Education," *The Quarterly Journal of Economics*, 114(2), 577–599.

BLUNDELL, R. W., AND J. L. POWELL (2003): "Endogeneity in Nonparametric and Semiparametric Regression Models," in *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, vol. 2, pp. 655–679.

——— (2004): "Endogeneity in Semiparametric Binary Response Models," *The Review of Economic Studies*, 71, 655–679.

CRA Reference (2005): "A Bankers Quick Reference Guide to CRA," *Federal Reserve Bank of Dallas*, Available at http://www.dallasfed.org/htm/pubs/pdfs/ca/quickref.pdf.

ERGUNGOR, O. (2007): "Foreclosures in Ohio: Does Lender Type Matter?," *Federal Reserve Bank of Cleveland, Working Paper 0724*.

GEWEKE, J. (1991): "Efficient Simulation from the Multivariate Normal and Student-t Distributions Subject to Linear Constraints and the Evaluation of Constraint Probabilities," in *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, pp. 571–578. Citeseer.

——— (2005): *Contemporary Bayesian Econometrics and Statistics*. Wiley-Interscience.

GEWEKE, J., G. GOWRISANKARAN, AND R. J. TOWN (2003): "Bayesian Inference for Hospital Quality in a Selection Model," *Econometrica*, 71(4), 1215–1238.

IMBENS, G. W., AND T. LEMIEUX (2008): "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 142(2), 615–635.

KEYS, B. J., T. MUKHERJEE, A. SERU, AND V. VIG (2010): "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans," *Quarterly Journal of Economics*, 125(1), 307–362.

KLEINER, M. M., AND R. M. TODD (2007): "Mortgage Broker Regulations That Matter: Analyzing Earnings, Employment, and Outcomes for Consumers," National Bureau of Economic Research; Cambridge, Mass., USA.

LIEBOWITZ, S. J. (2009): "Anatomy of a Train Wreck: Causes of the Mortgage Meltdown," Later published in "Housing America: Building Out of a Crisis", edited by Benjamin Powell and Randall Holcomb, 2009. New Jersey: Transaction Publishers.

ROSSI, P. E., G. M. ALLENBY, AND R. E. MCCULLOCH (2006): *Bayesian Statistics and Marketing*. Wiley.

Appendix A Implementing Bayesian IV

Two equations:

$$\begin{cases} AA_i = z_i'\delta + \varepsilon_{i,1} \\ y_i^* = AA_i\beta + x_i'\gamma + \varepsilon_{i,2} \end{cases},$$

y^* is unobservable, but $y = 1\{y^* \geq 0\}$ is. Here $z_i = (dist_i, x_i)$. Assumptions:

$$\begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \end{pmatrix} \sim N(0, \Sigma),$$

and priors:

$$\begin{aligned} \begin{pmatrix} \beta \\ \gamma \end{pmatrix} &\sim N\left(\begin{pmatrix} \overline{\beta} \\ \overline{\gamma} \end{pmatrix}, A_{\beta\gamma}^{-1}\right), \\ \delta &\sim N\left(\overline{\delta}, A_{\delta}^{-1}\right), \\ \Sigma &\sim IW(\upsilon_0, V_0), \end{aligned}$$

Gibbs sampling is used to obtain draws from the posterior, and there are four steps to it:

$$\begin{aligned} \{y_i^*\}_{i=1}^n &\mid \beta, \gamma, \delta, \Sigma, y, AA, z \\ \beta, \gamma &\mid \delta, \Sigma, \{y_i^*\}_{i=1}^n, AA, z \\ \delta &\mid \beta, \gamma, \Sigma, \{y_i^*\}_{i=1}^n, AA, z \\ \Sigma &\mid \beta, \gamma, \delta, \{y_i^*\}_{i=1}^n, AA, z, \end{aligned}$$

I detail each step below.

A.1 Step 1. Updating y^*.

For each i, conditional on $\varepsilon_{i,1}$, y_i^* is normal:

$$y_i^* \mid \varepsilon_{i,1} \sim N\left(AA_i\beta + x_i'\gamma + \frac{\sigma_{12}}{\sigma_1^2}\varepsilon_1, \sigma_2^2 - \frac{\sigma_{12}^2}{\sigma_1^2}\right),$$

and hence I need to make draws from this distribution. The important part is to account for the data on y_i, so that I would not draw $y_i^* < 0$ is $y_i = 1$ and vice versa. This means draws must be made from the truncated normal distribution, which is done via the mixed rejection algorithm of Geweke (1991) covered in Appendix B.

A.2 Step 2. Updating β and γ.

Given δ, one can compute ε_1 as $\varepsilon_{i,1} = AA_i - z_i'\delta$. Given the output of the data augmentation step, I treat y^* as observables now, and condition the equation for y^* on ε_1:

$$y_i^* = AA_i\beta + x_i'\gamma + \frac{\sigma_{12}}{\sigma_1^2}\varepsilon_{i,1} + \xi_{i,2|1},$$

where $Var\left[\xi_{2|1}\right] = \sigma_{22}^2 - \frac{\sigma_{12}^2}{\sigma_{11}^2}$. Denote $\tau^2 \equiv \sigma_{22}^2 - \frac{\sigma_{12}^2}{\sigma_{11}^2}$, then rewrite the above equation as:

$$\frac{y_i^* - \frac{\sigma_{12}}{\sigma_1^2}\varepsilon_{1i}}{\tau} = \frac{AA_i}{\tau}\beta + \frac{x_i{}'}{\tau}\gamma + \zeta_i,$$

and $\zeta_i \sim N(0,1)$. Thus I can now use the standard Bayesian linear regression algebra. Given the assumption of normal prior, I have:

$$\begin{pmatrix} \beta & \gamma \end{pmatrix}' \sim N\left(\begin{pmatrix} \tilde{\beta} & \tilde{\gamma} \end{pmatrix}', \left(\tilde{X}'\tilde{X} + A_{\beta,\gamma}\right)^{-1}\right),$$

where

$$\begin{aligned} \begin{pmatrix} \tilde{\beta} & \tilde{\gamma} \end{pmatrix}' &= \left(\tilde{X}'\tilde{X} + A_{\beta,\gamma}\right)^{-1}\left[\tilde{X}'\tilde{y}^* + A_{\beta,\gamma}\begin{pmatrix} \bar{\beta} & \bar{\gamma} \end{pmatrix}'\right] \\ \tilde{X} &= \begin{bmatrix} \frac{AA}{\tau} & \frac{x}{\tau} \end{bmatrix} \\ \tilde{y}^* &= \frac{y^* - \frac{\sigma_{12}}{\sigma_1^2}\varepsilon_1}{\tau} \end{aligned}$$

A.3 Step 3. Updating δ.

Rewrite the system of two equations in such a way that they would have the same parameters on the RHS, namely, δ. Nothing has to be done with the equation for AA, but the equation for y^* is a bit more tricky. I start by substituting the instrument equation into the main:

$$\begin{aligned} y_i^* &= \left[z_i'\delta + \varepsilon_{i,1}\right]\beta + x_i'\gamma + \varepsilon_{i,2} \\ &= \beta z_i'\delta + x_i'\gamma + \beta\varepsilon_{i,1} + \varepsilon_{i,2} \end{aligned}$$

Now I transform the last equation as follows:

$$\begin{aligned} y_i^* - x_i'\gamma &= \beta z_i'\delta + \beta\varepsilon_{i,1} + \varepsilon_{i,2} \\ \frac{y_i^* - x_i'\gamma}{\beta} &= z_i'\delta + \left(\varepsilon_{i,1} + \frac{1}{\beta}\varepsilon_{i,2}\right), \end{aligned}$$

and denote $\widehat{y}^* \equiv (y_i^* - x_i'\gamma)/\beta$. The system is now of the form:

$$\begin{cases} x_i = z_i'\delta + u_{1i} \\ \widehat{y}^* = z_i'\delta + u_{2i} \end{cases},$$

where

$$Var\begin{pmatrix} u_1 \\ u_2 \end{pmatrix} \equiv \Omega = \begin{pmatrix} 1 & 0 \\ 1 & \frac{1}{\beta} \end{pmatrix} Var\begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 1 & \frac{1}{\beta} \end{pmatrix}'$$

Given that Ω is symmetric and positive definite, its Cholesky root U exists and is unique:

$$\Omega = U'U.$$

Therefore, I take the transformed system, premultiply it by U, stack observations and obtain the standard normal Bayesian linear regression with unit variance:

$$\left(U^{-1}\right)'\begin{pmatrix} x \\ \hat{y}^* \end{pmatrix} = \left(U^{-1}\right)'\begin{pmatrix} z' \\ z' \end{pmatrix}\delta + \begin{pmatrix} \psi_1 \\ \psi_2 \end{pmatrix},$$

where $\psi = (\psi_1, \psi_2) \sim N(0, I)$. Given the assumption of normal prior, I have:

$$\delta \sim N\left(\tilde{\delta}, (\tilde{z}'\tilde{z} + A_\delta)^{-1}\right),$$

and where

$$\begin{aligned} \tilde{\delta} &= (\tilde{z}'\tilde{z} + A_\delta)^{-1}\left[\tilde{z}'\overleftrightarrow{y} + A_\delta\bar{\delta}\right], \\ \tilde{z} &= \begin{bmatrix} z' & z' \end{bmatrix}' \\ \overleftrightarrow{y} &= \begin{pmatrix} x & \frac{y_i^* - x_i'\gamma}{\beta} \end{pmatrix}'. \end{aligned}$$

A.4 Step 4. Updating Σ.

I first obtain the residuals for each equation:

$$\begin{cases} e_{i,1} = AA_i - z_i'\delta \\ e_{i,2} = y_i^* - AA_i\beta - w_i'\gamma \end{cases},$$

and then compute

$$S = \sum_{i=1}^{n}\begin{pmatrix} e_{i,1} \\ e_{i,2} \end{pmatrix}\begin{pmatrix} e_{i,1} & e_{i,2} \end{pmatrix}.$$

Then the properties of the inverse Wishart distribution guarantee that the posterior for Σ will be

$$\Sigma \sim IW\left(v_0 + n, V_0 + S\right).$$

Appendix B Sampling from the Truncated Normal Distribution

Let $\phi(\cdot)$ be a standard normal p.d.f., and $\Phi(\cdot)$ be the standard normal c.d.f., then $\Phi^{-1}(\cdot)$ is the inverse of Φ. Let $U_{[a,b]}$ denote the uniform distribution on $[a,b] : a < b$. Denote $TN_{[a,b]}\left(\mu, \sigma^2\right)$ to be the normal distribution with parameters μ and σ^2, truncated to the interval $[a,b]$. Its density at x is

$$\left[\Phi^{-1}(b) - \Phi^{-1}(a)\right]^{-1} \phi(x)\, 1\left\{x \in [a,b]\right\}.$$

It is immediately apparent that if $x \sim TN_{[a,b]}\left(\mu, \sigma^2\right)$ then

$$z \equiv \frac{x-\mu}{\sigma} \sim TN_{\left[\frac{a-\mu}{\sigma}, \frac{b-\mu}{\sigma}\right]}(0,1),$$

and hence I need to be able to draw from a truncated standard normal density only (which will be denoted $TN_{[a,b]}$ for brewity). There are several ways of doing that, but my interest is in sampling in an efficient manner.

B.1 Inverse C.D.F. Sampling

Let $x \sim TN_{[a,b]}$. Then $x = \Phi^{-1}(u)$ where $u \sim U_{[\Phi(a),\Phi(b)]}$. This method works fine when a is far enough from $-\infty$ and b is far from ∞. Otherwise numerical issues arise: if $|w| > 8$, then $w = \Phi^{-1}(p)$ usually cannot be solved numerically precisely enough.

B.2 Mixed Rejection Sampling

This algorithm, developed by Geweke (1991), uses a variant of importance sampling. Two key concepts are the *target* density — the one I want to draw from, and the *proposal* (also sometimes called instrumental) density — the one I use to simplify the drawing process.

Denote the target density by $f(x)$ and the proposal density by $g(x)$. Assume that $\exists M$: $\forall x \quad f(x) \leq Mg(x)$. Then the importance sampling scheme works as follows:

- draw $x \sim g$ and $u \sim U_{[0,1]}$,
- accept $y = x$ if $u \leq \frac{f(x)}{Mg(x)}$,
- else redraw x and u, repeat.

The mixed rejection algorithm alternates between 2 different sampling schemes:

1. Normal Rejection Sampling — draw x from $N(0,1)$ and accept it if $x \in [a,b]$, redraw otherwise. In other words, target density is the truncated normal density, proposal density is the untruncated standard normal density. This works great if only a small portion of probability mass at the tails actually gets truncated.
2. Exponential Rejection Sampling — covered in Section B.2.1 below.

The choice between sampling methods depends on the values of a and b. As long as I consider only the cases of left truncation and right truncation (i.e. either $a = -\infty$ or $b = \infty$), the cut-off values are as follows:

- if draws are needed from $TN_{[-\infty,b]}$, sample via normal rejection if $b \geq -0.45$, and via exponential rejection otherwise;
- if draws are needed from $TN_{[a,\infty]}$, sample via normal rejection if $a \leq 0.45$, and via exponential rejection otherwise.

It remains to discuss how the exponential rejection sampling must be administered.

B.2.1 Exponential Rejection Sampling

Motivating example is $TN_{[a,\infty]}$, where $\Phi(a)$ is close to 1. As $a \to \infty$, $TN_{[a,\infty]}$ converges to the exponential distribution on $[a,\infty)$ with kernel $\exp\{-\lambda z\}$ for $z \geq a$. So the target and proposal densities are:

$$\begin{aligned} f(x) &= c\exp\left\{-\frac{1}{2}x^2\right\} 1\{x > a\}, \\ g(x) &= \lambda \exp\{-\lambda(x-a)\} 1\{x > a\}, \end{aligned}$$

where c is the normalizing constant so that $f(x)$ would integrate to 1. The ratio is:

$$\begin{aligned}
\frac{f(x)}{g(x)} &= \frac{c}{\lambda} 1\{x > a\} \exp\left\{-\frac{1}{2}x^2 + \lambda(x-a)\right\} \\
&\leq \frac{c}{\lambda} 1\{x > a\} \exp\left\{\max_{x \geq a}\left[-\frac{1}{2}x^2 + \lambda(x-a)\right]\right\} \\
&= \frac{c}{\lambda} \exp\left\{\frac{1}{2}\lambda^2 - \lambda a\right\} 1\{\lambda > a\} + \frac{c}{\lambda} \exp\left\{-\frac{1}{2}a^2\right\} 1\{\lambda \leq a\} \\
&= M_1(\lambda) 1\{\lambda > a\} + M_2(\lambda) 1\{\lambda \leq a\},
\end{aligned}$$

and minimizing this in λ gives me the smallest probability of rejection. Geweke (1991) notes, however, that from the computational standpoint it is better to minimize only the second term, which yields $\lambda = a$.

The exponential rejection then proceeds in drawing $x \sim g$ and $u \sim U_{[0,1]}$, and accepting x as long as:

$$u \leq \frac{f(x)}{g(x \mid \lambda = a) M_2(a)} = \exp\left\{-\frac{1}{2}\left(x^2 + a^2\right) + ax\right\},$$

and rejecting the x otherwise.

B.3 Algorithm

This algorithm is presented for the case of drawing a scalar random variable. It is in principle vectorizable. Consider first the case of left-truncation, i.e. $TN_{[a,\infty]}$:

1. Compare a with 0.45. Suppose first that $a \leq 0.45$:

   (a) Draw $z \sim N(0,1)$

   (b) If $z > a$, done, else return to step 1a.

2. Suppose now that $a > 0.45$

   (a) Draw z from $g(z) = a\exp\{-a(z-a)\}$, say, via inverse c.d.f. transformation:

   - draw $w \sim U_{[0,1]}$

- use the fact that the c.d.f. that corresponds to $g(z)$ is $G(z) = 1 - \exp\{-a(z-a)\}$ to do the inversion:

$$
\begin{aligned}
w &= 1 - \exp\{-a(z-a)\} \\
z &= a - \frac{1}{a}\log\{1-w\}
\end{aligned}
$$

(can replace w with $1-w$ in the last line, since if $w \sim U_{[0,1]}$, then so is $1-w$).

(b) Draw $u \sim U_{[0,1]}$

(c) If $u < \exp\left\{-\frac{1}{2}\left(z^2 - a^2\right) + az\right\}$, done, else return to step 2a.

Now consider the case of right-truncation, i.e. $TN_{[-\infty,b]}$. Use the fact that if $z_1 \sim TN_{[-b,\infty]}$, then $z_2 = -z_1 \sim TN_{[-\infty,b]}$. Thus proceed as follows:

1. Sample z_1 from $TN_{[-b,\infty]}$ as detailed above

2. Set $z_2 = -z_1$.

Appendix C CRA Lending Test Guidelines

Table 3: Lending Performance Ratings

	Performance Ratings				
Characteristic	Outstanding	High Satisfactory	Low Satisfactory	Needs to Improve	Substantial Noncompliance
Lending Activity	Levels reflect EXCELLENT responsiveness to AA credit needs	GOOD	ADEQUATE	POOR	VERY POOR
Assessment Area(s) Concentration	SUBSTANTIAL MAJORITY of loans are made in the banks AA	HIGH percentage	ADEQUATE percentage	SMALL percentage	VERY SMALL percentage
Geographic Distribution of Loans	Reflects EXCELLENT penetration throughout the AA	GOOD	ADEQUATE	POOR	VERY POOR
Borrowers Profile	Distribution of borrowers reflects EXCELLENT penetration among customers of different income levels	GOOD	ADEQUATE	POOR	VERY POOR
Responsiveness to Credit Needs of Low-Income Individuals and Areas	Exhibits an EXCELLENT record	GOOD	ADEQUATE	POOR	VERY POOR
Community Development (CD) Lending Activities	A LEADER in making CD loans	Makes A RELATIVELY HIGH level	Makes AN ADEQUATE level	Makes A LOW level	Makes FEW, if ANY
Product Innovation	Makes EXTENSIVE USE of innovative and/or flexible lending practices	USE	LIMITED USE	LITTLE USE	NO USE