It's the Economy, Stupid: How Economic Growth Predicts Social Tolerance

A Logistic Regression Approach

Eben Lazarus

ECON 300 Honors Thesis | University of Pennsylvania

Professor Wolpin

Advisor: Professor Postlewaite

April 26, 2011

I. Introduction

Analyses of social views and their determinants have long been predominantly confined to the realms of sociology and psychology. With a few notable exceptions, these fields tend to treat social views as independent of the current state of the economy, and accordingly, news coverage and popular opinion tend to assume this dichotomy as well; elections in the United States, for example, are commonly categorized as either referenda on the state of the economy, or on social views, but not both. Recent literature in political science has begun to dispel this idea as it pertains to elections (see, for example, Hibbs, 2000), but only a small body of research has examined the link between the state of the economy in the short term and more general social issues.

This paper aims to conduct such research. More specifically, I attempt to analyze and characterize the relationship between social tolerance and short-term fluctuations in economic growth. This paper focuses mainly on tolerance toward immigrants in the United States, and whether this tolerance increases in times of high economic growth and decreases when the economy is stagnating.[1] This might seem somewhat intuitive, as immigration is certainly an issue that cannot be divorced entirely from the economy. But this intuitiveness allows this possible phenomenon to be modeled in a utility-maximization framework, which is useful in attempting to understand the mechanism by which tolerance might increase with growth. Further, the subject of tolerance toward immigration is inherently interesting in that there is widespread consensus among experts that increased immigration helps advanced economies,[2] but popular opinion tends not to reflect this.[3] Thus an analysis of the reasons for much of the popular intolerance of immigration is useful in explaining this dynamic between expert opinion and popular opinion, which does not exist for social issues that are simply a matter of personal preference. I do, however, conduct a secondary analysis on the interplay between growth and tolerance for abortion – which is not an issue inherently related to the economy – in order to more fully assess the hypothesized relationship between growth and tolerance.

I focus on the short term because news coverage and policymaking both often respond to de-

[1] "Tolerance" is a loaded word in referring to the subject of immigration; some would dispute that being non-accepting of immigrants should constitute intolerance, since those non-accepting individuals might, for example, be competing with immigrants for jobs. I still, however, wish to term this anti-immigration sentiment "intolerance," as this sentiment implies that immigrants should not have the opportunity for social mobility, that they are lesser citizens of the world, and that they should not be allowed to step foot on American soil, all because of where they happened to be born. I would classify this as intolerance, regardless of its motivation, while others might not. But regardless of the semantics of the situation, it is still a worthwhile question to ask whether growth predicts anti-immigrant sentiment. Those who are not disposed to accept the term "tolerance" in this context may simply consider this an analysis of whether growth predicts social liberalism.

[2] A broad survey of economists as reported in Simon (1989), for example, showed that precisely 0 percent of surveyed economists believed that immigration has had unfavorable effects on the U.S. economy, and 93% believed that the immigration level should stay the same or be increased.

[3] It is, of course, possible that the economists in favor of immigration are considering it in terms of aggregate production, while public opinion is more unfavorable of immigration because as individuals, they might be less likely to find employment given increased immigration. This possibility will be much more fully discussed in Section III.

velopments on this shorter timeframe.[4] A significant positive relationship between tolerance and short-term growth would reveal that these short-term dynamics are likely not a reflection of permanent changes in social policy preferences, and news coverage and policymaking would in an ideal world reflect this.

My analysis begins in Section III, in which I propose two alternative theoretical models that would explain the hypothesized relationship between growth and tolerance for immigration. I then present my data and method of analysis in Sections IV and V, respectively, and I present and discuss my results in Section VI. Before my own analysis, however, I begin by conducting a review of the relevant literature, with a focus on papers examining the effects of economic growth.

II. Literature Review

Though the effect of economic growth on social tolerance is a relatively recent question without much relevant literature, the more general effects of growth have long been debated, both in quantitative and philosophical realms of discussion. An examination of this literature is useful in contextualizing the subsequent analysis of the issue at hand in this paper, as regards both the proposed theoretical model and quantitative analysis.

I begin the review by examining the set of literature that contends, from a philosophical point of view, that economic growth has negative social consequences. Becchetti et al. (2007) detail the extensive history of this viewpoint, which often bears a moral message; growth, it is hypothesized, leads to materialism at the expense of social rectitude, and thus past a certain point it must be a negative influence on society from the standpoint of morals such as tolerance.

This "negative effects" hypothesis is perhaps best represented in the modern era by economist Fred Hirsch's *The Social Limits to Growth* (1976), which gives a more nuanced take on this classic assertion. Hirsch's book introduces the idea of a "positional good," a rival good whose value is largely attributable to its favorable comparison with substitute goods.[5]

Attempts to attain positional goods lead to zero-sum games, since there must necessarily be losers in a game in which all that matters is relative status. And if economic growth leads a country's populace to desire only positional goods, then few gains can be made from further growth, as this growth simply leads to these zero-sum games. As such, he states, increasing prosperity does not bear positive dividends for society's well-being or cohesion; it simply drives people into material competition with one another, which leads to a fundamentally unsated society bent on "conspicuous consumption" at the expense of others. These assertions are not backed with any quantitative analysis.

[4] See, for example, the Arizona anti-immigration laws passed in 2010 largely as the result of a short-term increase in anti-immigrant sentiment.

[5] Big houses, for example, are positional goods; a house can only be considered "big" if it is larger than the current average, and thus such a house is valuable only in its relation to smaller houses.

The work of political scientist Ronald Inglehart begins to lay the theoretical and empirical foundations for a refutation of the type of argument made by Hirsch. Inglehart's main contribution to this discussion is his idea of "post-materialism" (1971), which holds that growth does, in fact, beget positive social change. In the sense considered here, materialism is the set of values that places the heaviest weight on meeting immediate physical needs. Post-materialism, then, is the set of values that occurs after these immediate needs are met, and Inglehart proposes that such values include social cohesion (e.g., tolerance toward others) and personal freedom.[6]

Since Inglehart's original paper, the post-materialist thesis has been considered quite broadly (e.g., Inglehart 1981, Taniguchi 2006), with almost uniformly the same finding that post-materialism has become a prevalent set of values in developed countries. Further, time series analysis shows that long-term economic growth increases post-materialist values. Importantly, however, these findings are postulated to be based on intergenerational cohort effects, as Inglehart (1981) concludes that past economic growth leads to post-materialism in young, rising generations. But Inglehart does not observe significant period effects, in that short-term changes in economic growth do not seem to radically affect post-materialist sentiment in developed countries. This is important, as it implies that the hypothesis considered in this paper – that short-term fluctuations in growth help determine the state of social tolerance – is incorrect; Inglehart postulates that these changes occur over longer time spans. The analysis showing no significant period effects, however, is once again based on simple comparisons of means, so we must remain open to the possibility that better quantitative analysis will shine a more favorable light on this paper's hypothesis.

The analysis by Andersen and Fetner (2008) considering the post-materialist thesis begins to show somewhat more promise for my hypothesis. Using four waves of longitudinal World Values Survey data (spaced from 1981 to 2005), they compare tolerance toward homosexuality across 35 countries. The main finding is that using a hierarchical linear model, tolerance toward homosexuality, which is a post-materialist attitude, increases within individual countries as inequality decreases in the short term. This is distinct from an analysis based on economic growth, but it is important in its short-term scope. Uslaner (2002) hypothesizes that the reason for this type of finding is that inequality weakens social trust, which results in intolerance across class and cultural lines.

Andersen and Fetner do, however, also consider the effects of gross domestic product on tolerance toward homosexuality, but their findings in this analysis are weaker. Their initial model, again a hierarchical linear model, shows that across countries, higher GDP does increase tolerance. This finding is attenuated, however, when control variables are added to take into account each country's demographics at the time of the survey. Their paper is thus another finding in favor of the post-

[6]Inglehart's definition of post-materialism led to the creation of the World Values Survey, in which Inglehart played an instrumental role. The first objective of this survey was to test where and how these post-materialist values were taking hold, and it thus surveyed citizens of 22 countries, asking questions to gauge how important post-materialist values were as compared to materialist desires.The survey has been conducted roughly every five years since 1981, and it now includes 62 countries.

materialist hypothesis, although this one perhaps carries more qualifications. We must also note that the relevant economic variable considered here is not growth, but the level of GDP. And although a high GDP might serve as a proxy for high growth, the two are distinctly different variables. Our review has so far yielded good first approximations of the social effects of growth, but we have yet to fully examine the type of hypothesis my paper will actually consider.

At this point, we turn our review to works that more fully examine this paper's hypothesis. Such a review must begin with economist Benjamin Friedman's 2005 book, *The Moral Consequences of Economic Growth*. Friedman follows the post-materialist tradition, but his hypothesis more directly counters that of Hirsch (1976) from a normative standpoint; he states that in moral terms, economic growth has significantly positive effects in the form of tolerance, openness, and generosity for the underprivileged.[7] The purported mechanism for this change is psychological, and the means for the change is, importantly, not just the level of GDP but the level of recent growth. Friedman postulates that economic retrenchment or stagnation leads to insecurity deriving from people's psychological inclination to believe that the macroeconomy is zero-sum in nature.[8] Thus, a lack of growth makes people distrustful of others who might take their share of the pie, which in turn causes less social trust and especially less tolerance for others unlike oneself. Meanwhile, relatively high levels of growth make the supposed zero-sum nature of the economy less immediately relevant to people's personal situations, as they experience personal gains from growth.

Friedman does not, however, back up this assertion with any sort of quantitative analysis. Rather, he takes the reader through qualitative analysis of different periods of growth and stagnation in the US and Western Europe over the past 150 years. He attempts to show that in every short- to medium-term period of high growth, social liberalism increases, while stagnation breeds the opposite. As such, Friedman's hypothesis is the first that directly grapples with the issue in this paper.

We now review two papers that examine Friedman's hypothesis in the context of this paper's topic of interest, tolerance toward immigrants. First, Becchetti et al. (2007) examine survey data from the German Socio-Economic Panel, which followed the same 33,000 survey respondents across 13 years. The survey's relevant question is, "Are you concerned about immigration to Germany?" An answer of "Not concerned," for example, serves as a proxy for tolerance toward immigrants, and the other possible responses are categorized similarly. The authors use three-year individual income growth for surveyed individuals, three-year GDP growth, and three-year average unemployment as their economic variables, and run ordered probit regressions of tolerance on the economic variables (with other controls). Concurrent with Friedman's hypothesis, they find that positive growth in an individual's income increases tolerance, while increases in aggregate unemployment decrease tolerance. When including aggregate GDP growth in the same regression, however, the coefficient on this GDP growth variable is either negative or insignificant, suggesting that the relationship

[7] These are, to Friedman, self-evidently positive characteristics for a society to possess.

[8] This inclination is generally considered incorrect on a broad level, but a true examination of this issue lies outside the scope of this paper.

between short-term growth and tolerance is not so clear-cut. But this analysis suffers from problems of model specification; the authors likely should not include a variable for both personal income growth and aggregate GDP growth in the same regression. Further, in all of the regressions with ambiguous results, Becchetti et al. regress individual-level data on country-level growth variables, which is perhaps problematic in interpretation.

Next, we review Citrin et al. (1995), who examine a similar question, this time in the United States. They hypothesize that "[h]istorically, anti-immigrant sentiment in the United States has surged following sharp economic downturns, partly in response to the tendency of political and labor union leaders to blame foreign workers for unemployment and pressure on wages. For example...the Asian Exclusion Act of 1882 was passed in the aftermath of a severe recession and defined as necessary 'to protect labor.'" They then use ordered probit analysis to more rigorously assess the relationship between economic conditions and tolerance, using a relatively small dataset from the 1990-1992 National Election Study. This survey was a three-wave panel of roughly 2000 participants, and participants were asked whether they thought the number of immigrants should increase, decrease, or stay the same. Due to the short time frame of this survey data, the authors do not use income growth as their dependent variable; instead, they use individual-level income data collected in this survey, relying on comparisons across individuals with different levels of income for their analysis. This addresses a slightly different question than my analysis will consider. Their analysis using this variable yields nothing of significance; it appears that differences in income do not predict differences in views about immigration. This might not actually be so incriminating for my paper's analysis, as the combination of the findings of Becchetti et al. and those of Citrin et al. seems to suggest that growth is more important than absolute income in determining attitudes toward immigration.

To review, we have examined arguments in favor of both negative and positive effects of growth on welfare and morality. The wealth of data supporting post-materialism makes clear that we must remain skeptical of the supposedly negative effects of long-term growth on morality and welfare. The argument advanced by Friedman (2005) speaks more directly to this paper's analysis, and it gives us theoretical reasons to believe that the hypothesis to be considered here – that short-term growth trends positively affect tolerance toward immigrants – might be true. The few quantitative analyses considered here, however, have all addressed slightly different questions than the one to be examined in this paper. But this literature does give us a good prior set of information with which to consider this question, and it indicates that we must to this point remain agnostic about whether this paper's specific hypothesis – that short-term growth induces tolerance toward immigrants in the United States – is likely correct.

Finally, the literature reviewed contains qualitative reasons backing the authors' hypotheses regarding the effects of growth, but there exist no formal economic models that explain the mechanisms by which these effects occur. In the next section, I present such a model as it relates to tolerance for immigration and contrast it with the theories advanced by these authors. There does exist a body of empirical literature that gives insight into the possible specifications of such a theoretical model, but I will introduce the relevant papers in the course of detailing my model, which I begin now.

III. Theoretical Model

I propose two alternative models specifying the relationship between growth and tolerance, both of which follow from the hypothesis that increased growth positively affects tolerance.

i. Model 1

The first model allows for the possibility that changes in an individual's level of tolerance are due only to immigration's impact on his or her financial situation. As such, it is distinctly different from the types of qualitative models advanced in the reviewed literature. This model specifies an individual utility function that depends only on the price level and the probability of work, which in turn depend on the level of immigration and the level of economic activity. Thus, this model has the following parameters and functions:

1. m: The proportion of immigrants in the population of the country of interest. $0 \leq m \leq 1$.
2. a: The level of economic activity, taken as exogenous. a can be thought of as representing the level of economic activity as compared to trend, so that a represents short-term deviations from the trend level. $0 \leq a \leq 1$.
3. u: This period's utility, taken as exogenous.
4. β: The discount factor, taken as exogenous and constant.
5. $P(m)$: The price level, as a function of immigrants.
6. $W(m, a)$: The probability of job retainment in the next period, as a function of immigrants and economic activity. $0 \leq W \leq 1$.
7. $U(u, \beta, P(m), W(m, a))$: An employed individual's two-period utility, as a function of this period's utility, the discount factor, the price level, and the probability of work.

For the purposes of simplification, all native workers in this model are taken to be equivalent, and there are no regional differences in immigration level or in economic activity. Thus $P(m)$, $W(m, a)$, and $U(u, \beta, P, W)$ are equivalent for the entire country's population of non-immigrant workers. We assume that the skill level of immigrants is lower than that of native workers – which has been shown to have empirical backing, as in Cortes (2008) – but they compete for the same pool of jobs. The size of this job pool is non-constant, as it varies with the level of economic activity.

The assumptions underlying this model's functional form are as follows:

1. P is decreasing in m. There are certainly many complex mechanisms underlying the price

level of an economy, but in terms of the parameters of this simple model, the immigration level is the only one of consequence.[9] This is intuitive; the conventional wisdom surrounding immigrants is that they take jobs unwanted by native citizens, but those jobs would likely be accepted by those citizens if offered high enough wages. As such, immigrants lower the costs of low-level labor, thereby lowering the overall price level. This observation has largely been empirically borne out (see again Cortes, 2008, which demonstrates that the reduced price level is a product of immigrants' lower wages, and Zachariadis, 2011). This also makes sense given that the immigrants in this model have lower skill level than the native workers but are competing for the same jobs, meaning their wages will be lower.[10]

2. P is decreasing at a decreasing rate in m, and $\frac{1}{P}$ increases at a decreasing rate in m. That is, there are diminishing marginal returns to added immigrants in their effect on the price level, as well as diminishing returns in their effect on purchasing power ($\frac{1}{P}$). This means that wages in low-level jobs decrease less as more immigrants are added to the population, which is intuitive.[11]

3. W is decreasing in m and increasing in a. In the long run, immigration might not have any effect on the probability of finding work (Card, 2005), as the labor market will reach a new equilibrium even with added immigrants.[12] But given that this is a short-run utility model, the only dynamic of interest is the one that occurs before the labor market has reached equilibrium. In this timeframe, native citizens in my model are less likely to find jobs, as the number of jobs has not increased while the workforce has. This short-term effect has been documented in the literature, as a meta-analysis by Longhi et al. (2006) finds that immigrants have a small but real displacement effect on native worker employment. It is self-evident why W is increasing in a, as increased economic activity creates more jobs. I remain agnostic about the rate of decrease of W in m and the rate of increase of W in a; there is no reason to believe that the probability of finding work changes at a declining rate as the proportion of immigrants increases or as the level of economic activity increases.

4. U is decreasing in P and increasing in W. The intuition behind this assumption is relatively self-evident; higher P means lower purchasing power, while higher W means higher job probability.

---

[9] The level of economic activity certainly has a relationship with the price level, but in the short term, this relationship is largely indeterminate.

[10] The fact that P has not been constrained to be equal to 1 even in the long run implies that there are some search frictions involved in this model, but wages are given at the competitive level so as to simplify our analysis.

[11] This is added not in response to past studies, which have for the most part used linear analysis, but rather simply because it seems more reasonable to assume so. This assumption is, however, not required for most of the conclusions of the model.

[12] Further, in the long run, e is not exogenous, as immigrants have effects on economic activity.

Thus, I propose the following functional forms consistent with the assumptions above:

$$P(m) = \frac{1}{\ln(1+m)} \tag{3.1}$$

$$W(m,a) = \frac{1+a-m}{2} \tag{3.2}$$

$$U(u,\beta,P,W) = u + \beta(W\frac{1}{P}) \tag{3.3}$$

$$= u + \beta(\frac{1+a-m}{2})\ln(1+m) \tag{3.4}$$

The forms for P and W might seem somewhat convoluted; they were, however, chosen to fit all the constraints and assumptions underlying the model and to yield as simple a representation for U as possible. P would be more simply expressed as $\frac{1}{m}$, but then $\frac{1}{P}$ would not increase at a decreasing rate in m, so natural logarithm is used. And I take $\ln(1+m)$ instead of $\ln(m)$ so that all values of P are positive. Then W is chosen so that regardless of the values of a and m, it is in the $[0,1]$ range. It is linear in a and m because we assume nothing about the rate of change of W in terms of those two, and linearity gives the simplest possible form for W.

Given these specifications, all the dynamics in the utility function occur due to changes in a and m, since u and β are constant. To simplify the model accordingly, we can remove u and β, leaving us with the following modified utility function U_2:

$$U_2 = \frac{1+a-m}{2}\ln(1+m) \tag{3.5}$$

With the model specified, we can now perform comparative statics to examine how utility changes over the proportion of immigrants and the level of economic activity. This will illustrate the relationship between the economy and the ideal proportion of immigrants, and it will thus illustrate the utility-maximizing proportion of immigrants at any given level of economic activity.

Figure 3.1 gives an example of the type of change that occurs given an exogenous change in economic activity. Given an increase in a from 0.5 to 0.75, the utility-maximizing proportion of immigrants m^* increases significantly. As such, tolerance as defined here increases with economic growth.[13]

Aside from the change in m^*, we can also note that this model's functional form specifies that U_2 increases at any level of m as a increases; that is, as economic activity increases and holding all

[13]Note that the optimal proportion of immigrants as specified here is quite high, regardless of the value of a; m_1^* is roughly 0.65, while m_2^* is roughly 0.8, implying that more than half the population should be composed of immigrants in these two cases. This need not be taken literally, however; I could easily specify a more complex functional form that reduces the ideal proportion of immigrants for any given value of a. We must focus not on the actual values of m^*, but rather on the changes in m^* as economic activity changes. And as Figure 3.1 shows, m^* increases as a increases.

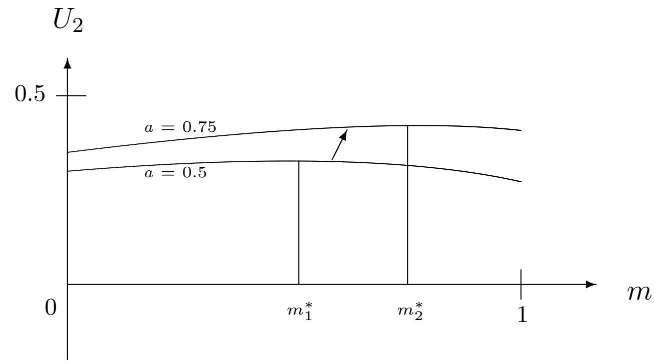

Figure 3.1: Change in Optimal Immigration Level for a Hypothetical Change in Economic Activity

else equal, utility increases. It is also follows that *maximized* utility increases as a increases. This is a positive feature of the model, as utility should increase with economic growth.

The type of analysis illustrated in Figure 3.1 can be carried out across every level of a, yielding a utility-maximizing locus of immigration level over a. This locus is plotted in Figure 3.2.

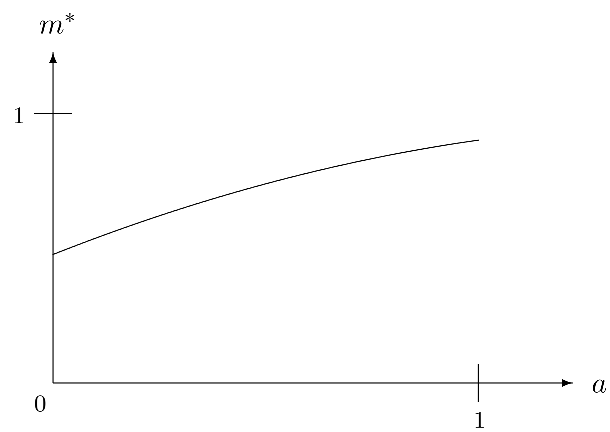

Figure 3.2: Optimal Immigration Level Across All States of Economic Activity

This figure makes it clear that tolerance for further immigration increases monotonically as economic activity increases, and thus that it increases in economic growth. This increase is not quite linear, as the functional form I assumed – namely, that $\frac{1}{P}$ increases at a decreasing rate in m – makes m^* increase at a decreasing rate as a increases. This is why the curve is slightly concave.

Thus, Model 1 is a theoretical representation of a world in which the relationship between growth and tolerance is due entirely to personal financial considerations. The model as detailed here illustrates the assumptions and functional forms necessary for such a representation.

We can now begin to examine the empirical backing for the functional forms expressed in this model. Native workers are, in this model, fully rational and have perfect information about the form of these functions. This is a strict assumption, but it is at least partially validated when examining polling data. In a May 2007 poll conducted by CBS News and the *New York Times*, a large plurality of survey respondents (42%) reported that the "main benefit from immigrants in the United States" is either that they "take jobs Americans don't want" or provide "low-cost labor." The next highest response, "provide cultural diversity," only covered 17% of respondents.[14] Thus assumption 1 above seems to have empirical backing, in that (at minimum) a plurality of citizens do have utility functions conforming to this assumption.[15]

Examining data on respondents' "main concern about immigrants in the U.S." from this same poll, however, shows that this model is somewhat less accurate in describing the downsides to immigration. Only 15% say that their main concern is that immigrants "take jobs away from residents." In contrast, a plurality (29%) say that their main concern is either that immigrants "use public services" or "don't pay taxes." The other responses aside from these would once again enter the utility function as constants, so we need not consider them. But this is not necessarily the case with the concern that immigrants take more in public services than they contribute in taxes, as this concern could change as the economy changes; if the level of economic activity decreases, for example, then workers might be less willing to subsidize others' social services given their deteriorating financial situations. So the concern expressed by assumption 3 certainly does matter to native workers, but this model perhaps excludes another, possibly more prescient, concern about immigrants.[16]

Looking further into the polling data, however, yields interesting results that lend some credence to the current model. This model assumes that all native workers are equal and competing with immigrants for jobs. In the reality, however, this is not the case, and many of the respondents in the survey we are considering now do not compete with immigrants for jobs. If we restrict our analysis to the respondents most likely to compete with immigrants for jobs, however, then the results detailed in the paragraph above are reversed; a plurality (21%) of respondents with income less than $30,000 per year state that their primary concern is that immigrants take jobs away from native residents, while only 13% express concern over public services. Meanwhile, 35% of respondents with income over $75,000 are concerned most about public services, while only 12% express concern about jobs. Thus it seems that given the assumption that all native workers compete with immigrants for jobs,

[14] These were all volunteered responses, with no concrete responses provided to choose from.

[15] Further, the other benefits from immigrants as reported by respondents could not conceivably be thought to change given changes in the level of economic activity, and so they could just enter our utility function as constants.

[16] This concern seems, however, largely unjustified given the literature estimating the fiscal effects of increased immigration, which shows at worst no effect and at best a slight negative effect on overall government transfers (see, for example, Auerbach and Oreopoulos, 1999). It seems, given this information, somewhat more appropriate to leave this consideration out of this model.

this model is reasonably well-specified, though it still could be missing a public services component.[17]

Thus it seems as if the model proposed here could serve as a reasonable approximation to reality. I must examine whether more rigorous empirical analysis confirms its conclusions, which are twofold:

1. There is a positive relationship between short-term economic growth and tolerance.
2. This increase in tolerance is attributable only to personal financial factors that change as the level of economic activity changes.

ii. Model 2

My alternative model is much more similar than Model 1 to the types of qualitative analyses described in the literature review. As such I will describe it qualitatively rather than quantitatively, as it lacks the sort of dynamics that allow us to perform a comparative statics analysis as done above. It posits that the first conclusion of Model 1 – that there is a positive relationship between tolerance and growth – is yet again true. Where this model differs from Model 1, however, is in the second conclusion. So with this model I lay out a framework that leads to the conclusion that the relationship between tolerance and growth is not due to purely financial factors, but rather that it is due primarily to some subjective emotional or psychological factors.

The proposed mechanism by which tolerance increases with growth relies on the positive link between income and happiness, as detailed in Stevenson and Wolfers (2008). As a country's overall income rises (i.e., as its level of economic activity increases), so does happiness. This model predicts that this increase in happiness has further emotional ramifications, in that this increase in happiness leads to an increase in tolerance. Not being a psychologist, I remain mostly agnostic about the various mechanisms by which this might occur, though the relationship would likely involve greater feelings of social amicableness derived from the happiness from greater economic security. This amicableness would, in turn, lead individuals to be more tolerant of differences in others, both in terms of demographic differences and lifestyle differences.[18] Likewise, the opposite phenomenon would occur during periods in which there are low levels of economic activity. This type of relationship would not be irrational in any strict sense, but unlike the first model, it posits that the relationship between tolerance and growth is not due to purely financial considerations.

[17] Perhaps the ideal model would include two types of workers: low-level workers who compete with immigrants for jobs, and high-level workers who do not. The low-level workers' utility function would follow the model specified here, while the high-level workers would have W substituted by a function expressing the net amount of public services received by the worker (i.e., the difference between taxes and benefits received). This public services function would again be decreasing in m and increasing in a, so it could theoretically have exactly the same form as W. Thus, there is no reason to believe that the added accuracy derived from defining two types of workers would yield us anything other than increased complexity, since U would be the same for both types of workers.We could also consider keeping only one type of worker but adding a public services term to the equation, but this public services term would take a form similar or equivalent to W, and would thus be somewhat redundant.

[18] Alternatively, or possibly in addition, this phenomenon could occur because of the mechanism proposed by Friedman (2005), which is quite similar in spirit to the one proposed here. Friedman posits that people have a model of the macroeconomy in which the economy is zero-sum in nature. Accordingly, in times of low growth, people are less likely to feel tolerant of others, as their belief that they are in competition with those others makes them less likely to be accepting of differences. Conversely, in times of high growth, people are more tolerant and accepting of others because they no longer feel as much pressure to win this supposedly zero-sum game.

An important conclusion that this model implies is that the positive relationship between tolerance and growth holds true for other types of social tolerance, e.g., tolerance for interracial marriage, or tolerance for abortion, since these other types of tolerance fall under the umbrella of differences in lifestyle choices. This is because the proposed social cohesion should not stop at the issue of immigration; it is, as proposed above, quite generally applicable to society as a whole. Thus the fact that the increase in tolerance in this model is due to greater social cohesion, instead of financial factors, has important and testable implications.

Given the conclusions of these two models, I must now empirically test two phenomena. First, I must test whether there is, in fact, a positive relationship between tolerance and short-term economic growth. This relationship forms the main conclusion of both models, and so its affirmation is the most important part of this analysis.

Second, if there is in fact a relationship between tolerance and growth, I must attempt to test the nature of this relationship so as to see whether it conforms more with Model 1 or Model 2. I will thus test to see if increases in other, non-immigration related social tolerance are observed given increases in growth. If they are observed, I will conclude that Model 2 more accurately describes the data. If not, then I will accept Model 1, since this will lead to the conclusion that only financial considerations lead to the relationship between growth and tolerance for immigration (if, in fact, this relationship does exist).[19]

I thus now turn to the empirical analysis of this issue, attempting to find resolutions to the questions posed by these models.

IV. Data

In analyzing whether tolerance for immigration depends on the state of the economy, one must quantify tolerance in a way that can be measured consistently and across time. This implies a need for consistent data on public opinion. Thus, given that no longitudinal surveys on tolerance exist in the United States, this paper will use data from polls of the US population. Specifically, the paper will examine responses to the following question:

Should immigration be kept at its present level, increased, or decreased?

41 national surveys since 1965 have asked this question.[20] 35 of these 41 surveys include data on the state of residence of each respondent. These polls are distributed throughout time as shown

[19]Given that the second conclusion of Model 1 stands in opposition to that of Model 2, we can reasonably assume that the union of these two models comprises most of the space of possibilities in the event that there is a positive relationship between tolerance and growth. There are certainly numerous other models that could describe this relationship, but these two models will suffice for this analysis given their nearly opposite hypotheses for the mechanism underlying this relationship.

[20]The relevant survey organizations are Gallup, CBS, *The New York Times*, and the Opinion Research Corporation.

in Figure 4.1, which illustrates that the dataset is skewed toward more recent dates.

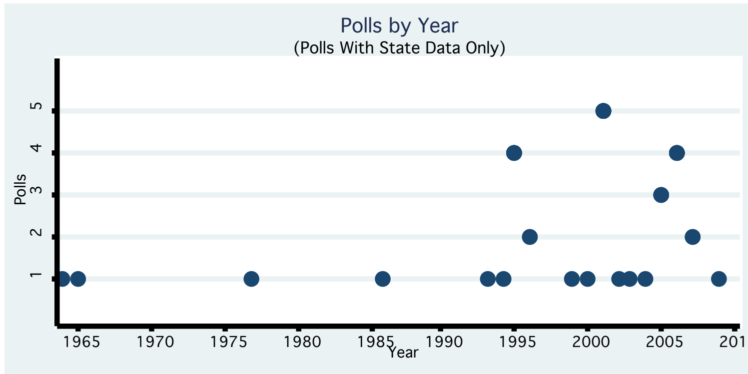

Figure 4.1

I treat "intolerance" as the statewide percentage of respondents in a particular poll who would like the level of immigration decreased. Thus one state's data in one poll comprises one observation. I use a variable for intolerance instead of tolerance simply for subjective reasons of interpretation; it seems that the idea of intolerance aligns more closely with the desire for immigration to be restricted, while tolerance could be defined to include the opinion that the current level of immigration is appropriate. Thus I use intolerance as defined here, despite the fact that all previous discussion in this paper has been about tolerance. So we must henceforth test whether there is a negative relationship between growth and intolerance, which would confirm the hypothesis that there is a positive relationship between growth and tolerance.

Each poll provides weights to apply to each individual's responses so that the poll as a whole reflects the demographics of the US population. These weights are constructed so that they are inversely proportional to each individual's probability of being sampled; for example, if a respondent's household has more than one telephone line, that household is more likely to be reached via a random dialing procedure, and so the survey organization will assign this respondent a lower sampling weight to account for this higher probability of being reached. Thus it is appropriate to use these sampling weights in calculating the statewide intolerance averages, and from this point forward, these state-level intolerance observations will be weighted averages.

Figure 4.2 gives a preliminary illustration of the distribution of intolerance throughout our sample period, as 4.2a shows the distribution of the national (not state-level) intolerance average over time while 4.2b shows a histogram of all the state-level intolerance observations. 4.2a shows that the data are clearly quite noisy, meaning a more fine-grained analysis is likely necessary to shed some light on the signal underlying this apparent noise. And the state-level intolerance observations shown in 4.2b are not normally distributed, as a Jarque-Bera test rejects normality at the 99% level.[21]

[21] Jarque-Bera jointly tests that skewness and excess kurtosis are 0, and thus the leptokurtosis caused by the fat tails in this distribution (due to the numerous 0 and 1 observations) causes the test to reject normality. The distribution is, however, approximately symmetric.

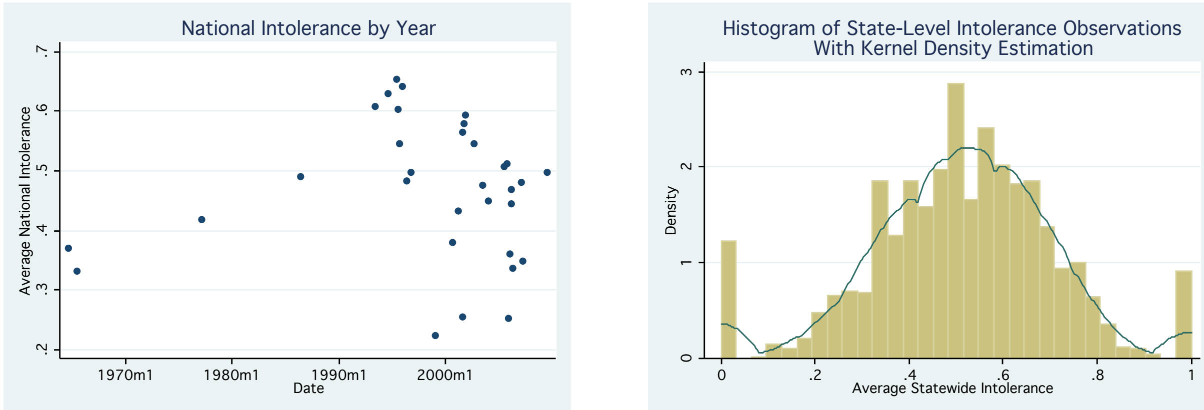

(a) (b)

Figure 4.2: Measures of the Distribution of Intolerance

The other main variable necessary for my analysis is economic growth. The variable for economic growth measures this growth at the state level in order to match the data measuring tolerance. I thus use quarterly, state-level personal income data, which is available from the Bureau of Economic Analysis. Data on personal income most closely tracks citizens' economic circumstances, since by definition it measures income instead of production, which is measured by gross domestic product.[22]

In order to normalize this data to account for both population changes and inflation, I use yearly state-level population data, again from the BEA, and monthly national Consumer Price Index data from the Bureau of Labor Statistics. The population data gives state-level population for the end of the second quarter of every year. Linear interpolation between adjacent years provides a measure of quarterly population. To obtain a measure of quarterly inflation, I take average CPI for every quarter, then calculate the percent change between these averages. Taking averages helps avoid the problem of transitory fluctuations that would create anomalies if I simply were to calculate inflation using the CPI from the end of each quarter.

Using this data on income, population, and inflation, I can then calculate state-level quarterly personal income growth on a real, per capita basis. The relevant variable for these descriptive statistics is this income growth aggregated over the past eight quarters, or two years, before the quarter in which a poll was conducted. For example, for the survey administered by Gallup in July 1995 (i.e., quarter 3 of 1995), the growth variable is state-level real per capita income growth from July 1993 (quarter 3, 1993) through June 1995 (quarter 2, 1995). Subsequent analysis will determine whether two years is the proper length of time over which to measure recent growth, as I split this aggregated growth variable into separate quarterly lagged growth variables for the sake of my regression analysis. For the purpose of this data description, however, this preliminary aggregated growth measure will suffice.

[22] The two are equivalent save for corporate profits, which are not included in personal income calculations. Further, there is always some measurement error.

Of the 41 polls that have asked the relevant question measuring tolerance, there are 32 for which there are data on each respondent's state and on income growth.[23] For these 32 polls, there are responses from an average of roughly 46 states per poll, giving 1487 observations (one for each state in each poll). On average, 27.6 individuals were surveyed per state per poll.

In order to more fully characterize the data, see Table 4.1 for a set of summary statistics on the two main variables of interest. The negative correlation between intolerance and growth is a preliminary indication that growth does lead to increased tolerance, but its low absolute value indicates that I must conduct further analysis to be certain about this relationship.

Summary Statistics for Intolerance and Growth

Variable	N	Mean	σ	Min.	Max.
State-Level Intolerance	1487	0.511	0.200	0	1
Eight-Quarter Real Growth	1487	0.0391	0.0318	-0.121	0.167
	Correlation:	-0.0669			
National Intolerance	32	0.468	0.116	0.224	0.654

Table 4.1

Figure 4.3 shows two measures of the relationship between growth and intolerance. First, Figure 4.3a shows a scatter plot of each statewide tolerance data point in one poll versus corresponding income growth, with OLS linear fit. The relationship is clearly quite noisy at this basic level, although the linear fit once again shows a significant negative relationship; the coefficient on growth is -0.421 (s.e. = 0.163), which is significant at the 95% level and indicates that a 1 percentage point increase in growth decreases intolerance by .421 percentage points.[24] Further regression analysis in Section VI will more fully reveal the nature of this relationship, but as it stands here, there is in fact a significant negative relationship between growth and intolerance toward immigrants; a 15 percentage-point decrease in growth, for example, would correspond to a 6.3 percentage-point increase in intolerance. The R^2 value is 0.0038.

Next, Figure 4.3b splits the growth variable out into individual quarterly growth and shows the cross-correlation between intolerance and lags of quarterly growth. The correlation is significant and negative for the first two quarterly lags but then damps quickly, and aside from some apparent outliers, is mostly insignificant for further lags.[25]

[23] In addition to the six surveys that did not record data on a respondent's state, I also do not have growth data corresponding to the three most recent polls, starting in April 2010. This is because there is no population data for 2010, and it is thus impossible to calculate per capita income.

[24] Note, however, that in addition to the lack of controls in this regression, I used no correction for heteroskedasticity in the error term; this regression represents a very simple first pass at analyzing a complex problem.

[25] This indicates that perhaps looking at 8 quarters of growth, as we have done to this point, has been too long a timeframe of analysis. This could mean that the relationship between growth and tolerance is actually stronger than as examined here, given that the higher lags could have damped the apparent correlation. This analysis suffices for now, however, given that this has simply been a preliminary characterization of the data.

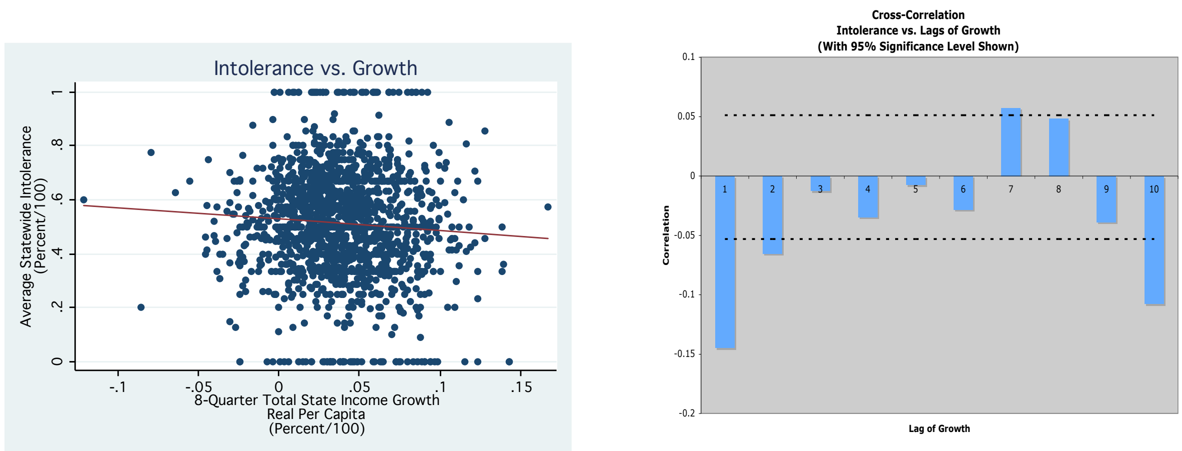

(a) (b)

Figure 4.3: Measures of the Relationship Between Growth and Intolerance

The main data, in sum, seem to show a negative relationship between growth and intolerance – and thus a positive relationship between growth and tolerance – on the state level. This relationship is as yet, however, quite noisy, and the methods heretofore used to analyze the relationship have been relatively crude.

Aside from these two main variables of interest, I will use the following control variables in regressions for intolerance:

1. An indicator variable coding for question wording. Most polls ask, "Should immigration be kept at its present level . . . ," but some ask, "Should **legal** immigration be kept at its present level. . . ." The variable henceforth referred to as "legal dummy" is equal to 1 when the word "legal" is included in the question, and 0 otherwise.
2. Dummy variables for each polling organization.
3. Dummy variables for each state.
4. Seasonal (quarterly) dummy variables.
5. Linear time trend.[26]
6. State-by-state linear time trends, to try as an alternative to the overall linear trend.
7. Statewide demographic controls among respondents on each poll, as follows:
   (a) Percent white and percent black.
   (b) Average age.
   (c) Percent college graduates.
   (d) Percent Democrat, percent Republican, and percent independent.[27]

[26] I do not include higher-order time trends so as to avoid in-sample overfitting.

[27] In addition, not every survey asked the same demographic questions, so there are also numerous other control variables available for only a subset of the polls: percent Hispanic or Latino, percent urban, percent married, percent with children under 18, and percent Protestant and Catholic. These are all potentially significant in explaining a state's intolerance level on a particular poll, but they all restrict the sample quite significantly and systematically, in that many of their missing observations come from the beginning of the sample period. Accordingly, I did not use them in my principal analysis, although I did check for the robustness of my results to their inclusion.

In addition to demographic controls obtained from the polling data, I could have used state-by-state demographic data obtained from the decennial Census. This would, however, have posed problems with endogeneity, as it is quite possible that state demographics (especially the proportion of immigrants) have a dynamic relationship with intolerance and growth. Accordingly, I use the state dummies to proxy for these state-by-state differences while minimizing the problem of endogeneity.[28]

Given considerations of space, I omit summary statistics for these control variables; their significance in this analysis lies not in their individual characteristics, but rather in their effects on the regression results. I thus move now to this regression analysis.

V. Method of Analysis

Given that my analysis attempts to explain variations in intolerance through variations in growth, the most appropriate model to use is a distributed lag model with lags of quarterly growth as the main dependent variables. Running vector autoregressions would be ideal, but the nature of the data does not allow for VARs; there are often large gaps in time between polls, which means including lagged intolerance in a model for intolerance is nearly impossible. I impose no structure (e.g., polynomial distributed lags) on the coefficients of the distributed lag model, as I have no prior belief about the form such a structure must take given that this subject of analysis is quite atypical.

The first problem to consider in my analysis involves the nature of the data proxying for intolerance. Since this intolerance variable represents a proportion, a model estimated using ordinary least squares would be inappropriate; we likely do not want to assume that growth has a constant effect on intolerance across the [0,1] range of intolerance values (especially given the concave functional form in Model 1 proposed above), and more importantly, there is no way to interpret predicted values of tolerance outside of the [0,1] range.

The most common solution to this problem involves applying a logit transformation of the following form to this fractional data:

$$\text{logit}(intol) = \ln\left(\frac{intol}{1 - intol}\right) \tag{5.1}$$

where *intol* is the statewide proportion of respondents in a particular poll who would like immigration to be decreased. This transformation cannot, however, be performed for 0 or 1 values of *intol*. Given that there are 76 such observations, this would be a significant limitation. Thus a better solution, as detailed in Papke and Wooldridge (1996), is to estimate a logistic regression model on untransformed

[28] It is true that this endogeneity problem could also manifest itself in the use of the demographic controls from the polls, but not to the same extent; those polls' demographics have more to do with who was sampled from each state, which directly causes differences in observed intolerance, than they have to do with the actual state demographics.

intol values. That is, I perform the following regression:

$$intol_{i,t} = \frac{1}{1+e^{-z_{i,t}}} \tag{5.2}$$

$$\text{where } z_{i,t} = \beta_0 + \beta_1 growth_{i,t-1} + \ldots + \beta_\tau growth_{i,t-\tau} + \gamma \underbrace{X_{i,t}}_{\substack{\text{Control}\\\text{variable}\\\text{vector}}} \tag{5.3}$$

Note that 5.2 can be rewritten in terms of the log-odds ratio expressed in 5.1:

$$\ln\left(\frac{intol_{i,t}}{1-intol_{i,t}}\right) = z_{i,t} \tag{5.4}$$

This means we effectively perform the same transformation in 5.2 as in 5.1, and the right-hand side expression in 5.2 can only take values on the $(0,1)$ range. But in the case of 5.2, no problem arises from allowing *intol* observations to be 0 or 1 even if the function on the right is constrained to lie strictly between 0 and 1; this is a regression framework in which our estimates will necessarily have some residual, and there is no conceptual problem with performing a linear projection of the vectors with 0 and 1 values of intolerance onto a line with intolerance levels strictly between 0 and 1. Thus I estimate this logistic model, which assumes a binomial distribution of the *intol* variable, to obtain the coefficients in 5.3. This is equivalent to a generalized linear model with a binomial distribution and logit link function.

The model is estimated using maximum likelihood. This is implemented in practice using an iteratively reweighted least squares algorithm, since estimates obtained via IRLS are identical to those obtained by maximizing the likelihood function.[29]

My other regression specifications – using analytical weights and the robust variance estimator – are included due to the nature of the data. The weights are necessary given that my *intol* observations are actually averages over a variable number of individual responses. The number of survey respondents per state per poll ranges from 1 to 813, and clearly the average over 813 observations is more precisely measured than the average over 1 observation. Thus we assume the variance of the j^{th} observation is $\frac{\sigma^2}{w_j}$, where w_j are the regression weights, and $w_j \propto n_j$, the number of observations over which j is averaged.[30] Then the IRLS algorithm calculates $\arg\min_{\beta} \sum_{j=0}^{N} w_j \left| intol_j - \hat{intol}_j(\beta) \right|^2$ iteratively, where β is the vector of coefficients on the right-hand side of 5.3.

Note that this weighting procedure has the capacity not only to change the variance estimates but also the coefficients' point estimates, as the more precisely measured observations (i.e., the ones averaged over more survey respondents) are weighted more heavily. In practice, this means

[29] I will later need the maximized value of log likelihood for information criterion values, and so I also re-estimate the models using Newton-Raphson optimization of the log likelihood to obtain these values. Conceptually, however, it is easier to think of the model in terms of IRLS, and IRLS is also more typical of GLM estimation.

[30] We do not have $w_j = n_j$ because the w_j's are rescaled to sum to the number of observations in the regression.

that larger states are weighted more heavily, since the number of individuals surveyed increases with state population in general. So if the coefficients' point estimates change in this weighting procedure, then the standard regression assumption that the coefficients are invariant across observations (i.e., states) is violated, since different state-by-state weightings produce different estimates. This does not, however, pose a large problem to my analysis; I am attempting to find the average effect of an x percent change in growth on intolerance, across all states, which is what we obtain in this regression even if the coefficients are not constant across states.[31]

I then use the robust variance estimator (i.e., the Hubert-White, or "sandwich," estimator) to correct for any leftover variance heteroskedasticity. Using both the analytical weights and the sandwich estimator is more efficient in obtaining variance estimates than using just the sandwich estimator by itself; we know precisely the form of the heteroskedasticity that averaging over different numbers of observations causes, and so more efficient variance estimates are obtained by accounting for this heteroskedasticity before applying the sandwich estimator. But the sandwich estimator is necessary because we must assume that not all of the original heteroskedasticity in variance is accounted for by the aforementioned differences in measurement precision; there are likely differences across time, across states, and across polling organizations. Further, if we believe that the regression model might be misspecified (e.g., we are missing certain explanatory variables), then the robust variance estimator gives more efficient variance estimates (Sribney, 1998). This is almost certainly the case here, as this is a nontraditional model with no clear "true" model we know of. So after parameterizing the initial variance values by using the analytical weights specified above, I calculate the final variance estimates using this robust variance estimator.

In terms of the interpretation of the regression output, Equation 5.4 implies that a one-unit change in, for example, $growth_{i,t-1}$ entails a β_1 change in the log-odds ratio in the left-hand side of 5.4, and so on for the other right-hand side variables. To more clearly see the effect of a one-unit change in the right-hand side variables on intolerance in this model, we can take the derivative of 5.2 with respect to $growth_{t-1}$ and obtain:

$$\frac{\mathrm{d}intol}{\mathrm{d}growth_{t-1}} = \frac{\mathrm{d}z}{\mathrm{d}growth_{t-1}} \frac{e^{-z}}{(1+e^{-z})^2} \tag{5.5}$$

$$= \beta_1 \frac{1}{1+e^{-z}} \frac{e^{-z}}{1+e^{-z}} \tag{5.6}$$

$$= \beta_1 intol(1 - intol) \tag{5.7}$$

This is generalizable to all other variables on the right-hand side of 5.3. This makes it quite apparent that a negative growth coefficient would imply a negative change in intolerance given an increase in

[31] I could, alternatively, put a *state dummy* $*$ *growth* variable for every state in every regression and then average over their coefficients, which would mean the assumption of constant coefficients would not be violated. This would, however, require using up 50 extra degrees of freedom in every regression, and so it does not seem worthwhile to implement this given that the only downside to the current method is the violation of an assumption that is relatively unimportant to this analysis.

growth, which would accord with this paper's hypothesis.

The last issue to consider is model selection. That is, which variables on the right-hand side of Equation 5.3 should be included in our model? To obtain the answer to this question, I wrote a simple algorithm to conduct regressions with all possible combinations of variables (varying both the number of lags of growth and the set of control variables) and computed the Akaike Information Criterion for each regression: $\text{AIC} = \frac{-2 \ln L + 2k}{N}$, where $\ln L$ is the log likelihood (computed via the Newton-Raphson optimization), k is the number of parameters in the model, and N the number of observations (Akaike, 1973). The log likelihood term makes the AIC decrease as goodness of fit increases, while the k term makes the AIC increase as more degrees of freedom are used by extra explanatory variables. Thus we look for the model with the lowest AIC. Division by N allows for comparability between models with different numbers of observations, as is often the case here given that certain demographic variables are only available in a subset of the of overall set of polls. I use the AIC as opposed to other information criteria because the AIC is designed to be asymptotically optimal (efficient and unbiased) in selecting the best model when none of the models being examined is the true model (Yang, 2005). As previously discussed, this is likely the case here, given that this is a nontraditional subject of analysis with an unknown true model. I thus begin my analysis agnostic about the variables to include in the regression model, and allow the AIC to select the one that maximizes the tradeoff between goodness of fit and parsimony.

VI. Results

i. The Effects of Growth on Intolerance

Table 6.1 shows the main set of regressions for intolerance, with the final regression in the table being the one that minimizes AIC. Controls are also added to the regressions, from left to right, in the optimal order as selected by AIC. The coefficient on growth is highly significant and negative for all regressions, including Regression 6, which is the AIC-minimizing regression. The AIC for a regression on a constant (Regression 1) is higher than that for a regression with growth, so adding in the lagged growth variable does add significant explanatory power to the model. This one lag of growth is, however, the only one that lowers the AIC; all subsequent growth variables added to the regression do not increase the regression's goodness of fit enough to offset the degree of freedom loss, and so the best model is one in which only one growth variable is included. But this one lag of growth does have a significantly negative relationship with intolerance, as hypothesized by the models introduced in Section III. It is thus clear that these models must likely be quite short-term in nature, as the effect of a change in growth on intolerance occurs on quite a short time frame. Note, however, that even though these regressions only include one lag of growth, income growth is a highly autocorrelated series, and so a change in one quarter's growth likely has effects on intolerance beyond one quarter due to feedback effects in growth (which are not explicitly modeled here).

Table 6.1

Regressions for Statewide Intolerance

Variable	(1)	(2)	(3)	(4)	(5)	**(6)**
Lag 1 growth		-9.47***	-9.07***	-7.20***	-7.73***	-8.38***
		(2.24)	(2.12)	(2.07)	(1.99)	(2.01)
"Legal" dummy			-0.25***	-0.38***	-0.49***	-0.56***
Linear trend				0.002***	0.004***	0.005***
% College grads					-1.15***	-1.12***
% White						0.50***
AIC	0.9545	0.9478	0.9454	0.9445	0.9421	**0.9407**

N=1487

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The other variables' coefficients in Table 6.1 are all highly significant, as would be expected given that the model selection criteria is based on goodness of fit. The direction of each control variable's coefficient also matches up with expectations. Asking about "legal" immigration versus simply immigration overall makes people less intolerant of immigration; there is a slight upward trend in intolerance; increasing the number of college graduates sampled reduces observed intolerance; and sampling more non-minorities (i.e., whites) increases intolerance. The absolute value of the growth coefficient is significantly larger than those of the control variables, meaning that observed intolerance is much better predicted by the state of the economy than by the survey sample.[32]

These ordinal rankings of the size of the coefficients' absolute values do not, however, fully illuminate the true size of the effect of the right-hand side variables on intolerance. We can get an idea of the actual size of the effect of a change in growth on intolerance by using Equation 5.7, which allows us to examine the effect of a change in growth across the $[0, 1]$ range of *intol* values by taking the following derivative:

$$\frac{\mathrm{d}intol}{\mathrm{d}growth_{t-1}} = \beta_1(intol)(1 - intol) \tag{6.1}$$

$$= -8.38(intol)(1 - intol) \tag{6.2}$$

If we assume that *intol* is at its sample mean (0.504), then $\frac{\mathrm{d}intol}{\mathrm{d}growth_{t-1}} = -8.38(0.504)(1 - 0.504) = -2.09$. This is an estimate for β in a traditional sense, in that it tells us the actual effect of a change in growth on intolerance. It is inappropriate to say that this approximates the change in *intol* given a one-unit change in growth, however, since this is simply the value of the derivative at $intol = 0.504$; given the functional form of the regression model, as *intol* changes, then this derivative changes too. So this derivative is an estimate of the change in intolerance given an infinitesimal change in growth.

[32] The fact that the survey coefficients are significant, however, shows that these polls are not entirely robust to sampling effects.

The logistic regression specification also dictates that the absolute value of this derivative is largest at $intol = 0.5$; the derivative $-8.38(intol)(1-intol)$ is a positive parabola with roots at $intol = 0, 1$, so the minimum occurs at $intol = 0.5$. Thus the value of -2.09 is very close to the absolute maximum derivative value given that it was calculated for $intol = 0.504$.

This derivative analysis sheds light on the effects of an infinitesimal change in growth on intolerance, but what about the effects of a larger change in growth? If we integrate Equation 6.2 over growth, we can obtain a direct estimate of how much intolerance changes for any given change in growth. This integration simply brings us back to the original model, which is intuitive given that I differentiated in terms of growth to get to Equation 6.2 and am now integrating back over growth:

$$\int_{g_1}^{g_2} -8.38(intol)(1-intol)\mathrm{d}growth_{t-1} = \int_{g_1}^{g_2} -8.38\frac{e^{-z}}{(1+e^{-z})^2}\mathrm{d}growth_{t-1} \tag{6.3}$$

$$= \int_{g_1}^{g_2} -8.38\frac{e^{-(\beta_0-8.38growth_{t-1}+\gamma X_t)}}{\left[1+e^{-(\beta_0-8.38growth_{t-1}+\gamma X_t)}\right]^2}\mathrm{d}growth_{t-1} \tag{6.4}$$

$$= \frac{1}{1+e^{-(\beta_0-8.38growth_{t-1}+\gamma X_t)}}\Bigg|_{growth_{t-1}=g_1}^{g_2} \tag{6.5}$$

where g_1 is the starting growth level and g_2 is the end growth level. If we assume, for example, that all right-hand side variables in Regression 6 are at their average levels (so g_1 is at the mean level of growth, which is 0.0044, or 0.44% real quarterly growth), then we can track the effects of changes in growth on intolerance by plugging in different values for g_2. If, for example, growth increases by one standard deviation (0.0116), then $g_2 = 0.0044 + 0.0116 = 0.016$. Using Equation 6.5 and plugging in our estimated values for β_0 and γ along with the mean values for the X variables, working out the algebra shows that this one-SD increase in growth decreases *intol* by 0.024, or 2.4 percentage points. Subjectively speaking, this is not very economically significant, as few would notice a change in intolerance of only 2.4 percent of the population.

If we assume, however, that we can extrapolate out well beyond one standard deviation, then a change in growth can have quite an economically significant effect on intolerance. Assume now that growth decreases by 5 standard deviations, as can happen in a severe recession. Plugging in this new number for g_2, the new change in *intol* is 0.118, or 11.8 percentage points. This would be quite a significant change, as it would represent nearly 12 percent of the population demanding less immigration. This could have significant effects on policymaking.[33] So given certain initial state and functional form assumptions, these results show that a change in income growth can have significant effects on attitudes toward immigration.

An initial examination of whether we can, in fact, extrapolate these conclusions out past a one-

[33]This is especially true given that in reality, the population willing to change its immigration preferences might be comprised largely of swing voters.

SD change in growth is promising. I performed the main regression (Regression 6) with all the same specifications, except for the restriction that $|growth_{t-1} - growth_{t-2}| > \hat{\sigma}_{growth}$; that is, I performed the same regression with the restriction that there had to be at least a one-SD change in growth between quarter $t-2$ and quarter $t-1$. It would have been ideal to test this for greater changes in growth (e.g., for a 5 standard deviation change), but even using a two-SD restriction overly restricts the sample to $N = 188$. As it stands using the one-SD restriction, $N = 602$. And the results from this regression are nearly identical to the ones from the original regression, as the coefficient on lag 1 growth is -8.33, which is again significant at the 1 percent level. Thus there is evidence that truly significant changes in attitudes toward immigration can come about as a result of changes in growth.

Continuing with further analysis, Table 6.2 shows the results of conducting the same regressions (with the same optimized controls) while adding increasing numbers of lagged growth variables on the right-hand side of the regression. The results show that despite the fact that the optimal regression as selected by AIC includes only one lag of growth, there is still a significant relationship between growth and intolerance beyond one lag. This relationship is, as before, negative for every significant coefficient. There are multiple significant growth coefficients in every regression up through the one with 10 lags of growth on the right-hand side; beyond that point, controlling for further lags reduces the absolute value of many of the coefficients' estimates below the point of significance at the 5% level. The coefficient on lag 1 growth, however, remains significant for every regression. Thus this provides further evidence of the significantly negative relationship between growth and intolerance.

Table 6.3 shows the effects of adding other control variables to the regression. This is useful because certain controls are arguably theoretically necessary despite the fact that the AIC did not select them to be included in the optimal model. The state dummies, for example, control for various state-by-state differences that might have effects on the level of intolerance, all of which were previously unaccounted for. The AIC likely did not select them to be included in the final model, however, because of the degree of freedom restrictions imposed by their inclusion, as there are 51 state dummies (with Washington, D.C.). The same can be said of the polling organization dummies, which control for differences in sampling methods. For the sake of simplicity, the coefficients on the controls have been omitted. The only controls with any real effect on the magnitude of the growth coefficient are the quarterly dummies, which shows that there are seasonal changes in intolerance. It is clear that the significance of the growth coefficient is robust to all controls included here.

The growth coefficient was also robust to the inclusion of the numerous other control variables (with fewer observations), and the results of these analyses are omitted here for considerations of space. I also ran regressions using $growth * control$ interaction variables for all relevant controls. These results are once again omitted, as there were few significant interaction terms and the regression results generally became quite ambiguous with their inclusion; the growth coefficient, for example, became much more negative in size, while its standard error blew up to the point that it was statistically indistinguishable from 0. Joint F-tests of the growth coefficient with the interaction

Table 6.2

Regressions for Statewide Intolerance
on Different Lags of Growth

Variable	(1)	(2)	(3)	(4)	(5)	(10)	(20)
Lag 1 growth	-8.38***	-8.76***	-8.88***	-8.98***	-9.19***	-8.45***	-6.41***
	(2.0)	(2.2)	(2.0)	(1.8)	(1.8)	(1.9)	(2.0)
Lag 2 growth		-4.56	-5.49**	-4.73*	-4.74*	-5.08**	-2.32
		(3.0)	(2.6)	(2.7)	(2.6)	(2.2)	(2.0)
Lag 3 growth			-6.06**	-6.00**	-5.77**	-5.62**	-2.95
			(2.8)	(2.6)	(2.5)	(2.2)	(1.9)
Lag 4 growth				-3.51	-3.49	-1.75	-2.93
				(2.6)	(2.6)	(2.2)	(1.9)
Lag 5 growth					-1.22	-3.85*	-3.23
					(2.0)	(2.0)	(2.0)
⋮							
Lag 10 growth						-5.17***	-6.59***
						(1.8)	(1.8)
⋮							
Lag 20 growth							-0.66
							(2.4)
Controls:							
"Legal" dummy	✓	✓	✓	✓	✓	✓	✓
Linear trend	✓	✓	✓	✓	✓	✓	✓
% College grads	✓	✓	✓	✓	✓	✓	✓
% White	✓	✓	✓	✓	✓	✓	✓
AIC	0.9407	0.9416	0.9423	0.9433	0.9447	0.9495	0.9590

N=1487
Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

coefficients, however, generally yielded significantly negative results overall. I did not run any regressions with modified forms of any the explanatory variables by, e.g., squaring them or taking their natural log, as it seemed that adding another layer of functional form possibilities was unnecessary.

Finally, I conducted a vector autoregression analysis with the available data in order to see the effects of growth on intolerance when taking into account intolerance's autoregressive effects as well as growth's feedback effects. As discussed above, the numerous gaps in time in the polling data prevented me from using this as my primary analysis. It is worth doing with the restricted set of consecutive polls, however, in order to construct an impulse response function to see how intolerance responds to an exogenous shock in growth. For this analysis, I took the log-odds ratio of *intol* before running the regression; even though this eliminates the 0 and 1 values of *intol*, it allows the VAR to be estimated in the traditional manner, using ordinary least squares. Thus I estimated a bivariate VAR with one lag of log-odds *intol* and one lag of growth in both the equation for log-odds *intol* and the equation for growth. I used no exogenous controls so as to simplify the analysis. My data is in panel form, given that I have multiple observations (one for each state) for each quarter of

Table 6.3

Regressions for Statewide Intolerance
With Different Controls

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lag 1 growth	-8.38***	-9.11***	-9.75***	-6.03***	-6.43***	-6.35***	-6.48***	-8.46***
	(2.0)	(1.7)	(1.9)	(2.0)	(2.2)	(2.1)	(2.1)	(1.9)
Controls:								
"Legal" dummy	✓	✓	✓	✓	✓	✓	✓	✓
Linear trend	✓	✓	✓	✓	✓	✓	✓	
% College grads	✓	✓	✓	✓	✓	✓	✓	✓
% White	✓	✓	✓	✓	✓	✓	✓	✓
State dummies		✓	✓	✓	✓	✓	✓	✓
Polling org. dummies			✓	✓	✓	✓	✓	✓
Quarterly dummies				✓	✓	✓	✓	✓
Average age					✓	✓	✓	✓
% Democrat						✓	✓	✓
% Independent							✓	✓
State-by-state trends								✓
AIC	0.9407	0.9992	1.0016	1.0037	1.0090	1.0100	1.0113	1.0762
N	1487	1487	1487	1487	1391	1391	1391	1391

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$
% Republican not included to avoid multicollinearity

observation, but I was nonetheless able to estimate this VAR by assuming a constant covariance matrix across all states.[34] I then estimated over the entire sample, regressing each growth and log-odds *intol* value on the relevant state's growth and log-odds *intol* value lagged one quarter. Surprisingly, the only significant coefficient was the lagged growth variable in the equation for *intol*; it was, as before, significantly negative. The others were all insignificant, likely due to the small sample ($N = 311$), but they all pointed in the expected direction; there was positive feedback from intolerance to intolerance and from growth to growth. Further, the coefficient on *intol* in the regression for growth had the highest p-value of all the coefficients, as expected.

The impulse response function constructed using this VAR(1) analysis is shown in Figure 6.1. I implemented a Cholesky decomposition in order to orthogonalize the error terms of the two regressions, although there was very little covariance from the outset. The confidence intervals in the figure were obtained using asymptotic standard errors. Once again, a positive shock in growth has a significant negative impact on intolerance. The effect of growth on intolerance as estimated here damps down quite quickly, however, as the change in log-odds intolerance is not significantly different from 0 two quarters after the shock. Thus this analysis yields slightly different results than the analysis done in Table 6.2, which seemed to show that many quarterly lags of growth had an effect on intolerance. Note that my method of estimation only allows us to see this effect on log-odds intolerance, however, rather than overall intolerance. Given the imprecision in estimates due to the small sample, however, the significant negative effect matters much more than the overall magnitude

[34]As discussed above, the possible violation of this assumption would not pose large problems for this analysis.

of the effect, and so I do not perform the type of analysis done above to find the "true" effect of a change in growth on intolerance here.

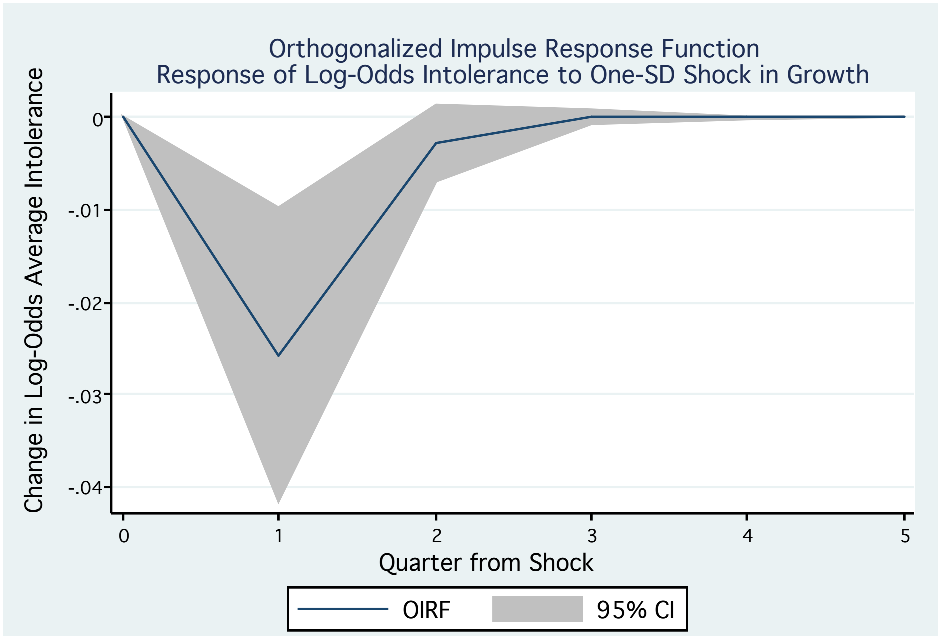

Figure 6.1

Thus, we must conclude that there is significant evidence in favor of the main hypothesis advanced by both models presented above: namely, that short-term growth has a significant negative relationship with intolerance, and thus a significant positive relationship with tolerance. It also seems reasonable to conclude that the causality (or at least some causality) runs from growth to tolerance, as there is little reason to believe that short-term changes in intolerance could cause changes in growth. It is possible that this analysis could also capture some long-term effects; perhaps the more intolerant states have lower average growth, for example.[35] But it would be unreasonable to conclude that this is the only dynamic at play, especially given that the only significant coefficient in the VAR analysis was that of growth in the intolerance regression (thus growth Granger-causes intolerance, while intolerance does not Granger-cause growth). Thus, we can reasonably conclude that changes in growth cause changes in tolerance as detailed by both models presented in Section III.

ii. The Mechanism for This Relationship

The remaining analysis must attempt to resolve the question posed by the competing models; that is, what is the mechanism by which tolerance increases with growth? As detailed at the end of Section III, I now attempt to answer this question by testing whether measures of social tolerance seemingly unrelated to personal financial considerations increase with growth as well. If so, then

[35] There is some literature that suggests that when considering long time horizons, increased tolerance has positive effects on growth; see, for example, Alesina et al. (1999) and Easterly and Levine (1997)

we must conclude that Model 2 more accurately describes the nature of the relationship between growth and tolerance, and that there is something motivating this relationship beyond pure financial considerations. Otherwise, we must accept that Model 1 more accurately describes the relationship.

Aside from immigration, the social issue most commonly asked about in my set of polling data is abortion. Views on abortion seem to have no connection to financial considerations,[36] so I used tolerance for abortion in this secondary analysis. Being accepting of other people's abortions might seem somewhat less "tolerance"-related than being accepting of immigrants, but I believe the analogy holds. Tolerance is generally defined as the tendency to be accepting of others' diverse thoughts, opinions, or actions, and both cases involve being accepting of the actions of others; allowing immigrants to live freely in one's country falls under this definition, as does allowing women to make choices freely. But regardless of whether one agrees with this argument, both issues pertain to social liberalism, and so one who objects to this argument (as one may freely choose to do) can simply think of the relationship examined here as one between growth and liberalism, as explained above in Footnote 1.

I move now to the analysis. Given time limitations, I was required to use the same set of polls as used in my immigration analysis. Only 8 of these polls asked questions about abortion, so my sample was restricted to $N = 399$. Whereas in my immigration analysis it seemed most intuitive to define an average intolerance variable, in this abortion analysis it seems more intuitive to define an average tolerance variable. Thus we look here for a positive coefficient on growth, rather than negative. This tolerance is defined as the weighted statewide average of people responding to prompts about abortion by saying that abortions should be allowed with no restrictions.[37] I then performed the same distributed lag logistic regressions as above, with tolerance as my dependent variable and lags of growth as my main dependent variables.

The results from these regressions, with one lagged growth variable, are shown in Table 6.4. The results are ambiguous. The optimal model as selected by the AIC is shown in Regression 3, in which the right side contains only one lag of growth and a linear time trend.[38] In this regression, the lagged growth variable has a significantly positive coefficient, indicating a positive relationship between growth and tolerance for abortion. Its magnitude is also much larger than the magnitude of the growth coefficient in the immigration regressions, although it has a much larger standard error. But this model seems lacking from an intuitive standpoint, as it does not control for many other factors that might go into the analysis. Perhaps the AIC analysis was compromised here because of the small sample size, which limits the degrees of freedom for error. Perhaps, though, this analysis is actually accurate; it is hard to make the case for omitted variable bias even in this simple model,

[36] This is not, of course, taking into account the views of the small minority of people considering getting abortions.

[37] Other possible responses to the questions, which were all similar but not identical, were that the current restrictions were adequate, that there should be more restrictions, or that it should never be legal.

[38] The AIC increases between Regressions 1 and 2, which appears to show that the growth variable should not be included in the model, but it is in fact the case that the AIC-minimizing model is the one with one lag of growth and linear trend; the model with just linear trend has higher AIC than the model with trend and growth.

since growth should be (and is mostly) uncorrelated with these control variables. So this model could truly be the one that maximizes the tradeoff between goodness of fit and parsimony.

Table 6.4

Regressions for Statewide Tolerance for Abortion

Variable	(1)	(2)	**(3)**	(4)	(5)	(6)	(7)	(8)
Lag 1 growth		11.86***	**17.56*****	11.08**	9.67**	-6.26*	-6.27*	-4.65
		(4.3)	(4.0)	(4.7)	(4.5)	(3.6)	(3.6)	(4.2)
Controls:								
Linear trend			✓	✓	✓	✓	✓	✓
Question wording dummies				✓	✓	✓	✓	✓
% Democrat					✓	✓	✓	✓
% College grads					✓	✓	✓	✓
% White					✓	✓	✓	✓
Average age					✓	✓	✓	✓
State dummies						✓	✓	✓
Polling org. dummies							✓	✓
Quarterly dummies								✓
AIC	0.9152	0.9174	**0.9058**	0.9205	0.9346	1.1617	1.1617	1.1715

$N = 399$
Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

As shown in the table, I then added in control variables regression by regression. I used an admittedly somewhat arbitrary process in deciding the order to add these variables, based on order of importance. It first seemed important to control for question wording; then I added the four most important demographic polling variables as selected by the AIC; then state dummies, to account for other overall state-by-state differences; then polling organization dummies, and then quarterly dummies. The only controls that had any effect were the state dummies, which had such a large effect that the growth variable turned negative (but statistically indistinguishable from 0 at the 5% level). Perhaps this indicates that the true predictors of tolerance for abortion are the state-level demographic variations. There is, however, an alternative interpretation. Given that the sample consisted of only 8 polls, there were inevitably states that exhibited higher than average growth throughout the sample period. This means that the state dummies could have picked up these growth effects, which would have diminished the coefficient on the actual growth variable, as was observed. And, in fact, when $growth * state\ dummy$ interaction terms were included in the model, 48 of 51 these interaction terms' coefficients were positive at the 5% level, while none of the state dummies' coefficients remained significant. The growth coefficient did, however, become significantly negative, and Wald tests considering whether $\beta_{growth} + \beta_{state*growth} = 0$ (for each state) only yielded 7 significantly positive combinations at the 5% level, likely due to the large standard errors from the small sample size. So the interaction variable analysis indicates that the state dummies were, in fact, picking up the state-by-state variations in growth across the sample, but the interpretation is still somewhat ambiguous given the results of the Wald tests. Finally, all of the results detailed in the table stay consistent when higher-order lags are added to the regression, as the coefficients on these lags are always either insignificant or point in the same direction as the coefficient on lag 1 growth in the table here.

Thus, although we must remain somewhat skeptical of these results given the small sample size and the ambiguity of the results, it does appear that tolerance for abortion increases with growth. As such, it is clear that we cannot reject that tolerance for immigration increases with growth via the mechanism proposed in Model 2, as it seems that there are non-financial considerations at work in this relationship. But the lack of a truly definitive conclusion to the abortion analysis indicates that perhaps the mechanisms proposed by Model 1 and Model 2 both contribute to the relationship between growth and tolerance for immigration, and that this relationship includes both financial and non-financial considerations. Further study is needed to answer this question more fully.

VII. Conclusion

It seems quite apparent at this point that, as hypothesized, short-term fluctuations in economic growth have a significant positive relationship with tolerance for immigration. Section III presented two possible models explaining this relationship, while the statistical analysis in Section VI validated this relationship using logistic regressions on self-proclaimed intolerance. The mechanism underlying this relationship, for which alternative possible models were provided in Section III, is less clear. This lack of a definitive conclusion indicates that the mechanism is perhaps a function of both personal financial considerations and psychological changes caused by changes in the state of the economy.

The conclusion that growth significantly predicts tolerance for immigration, and that it predicts this tolerance better than any other single factor used in this analysis, has significant implications for news coverage and policymaking, as discussed in the Introduction. Short-term changes in social views, at least on the subject of immigration, should perhaps be treated more as a function of the state of the economy than as exogenous and permanent changes in society's overall ideology. As such, it seems reasonable to conclude that news coverage and policymaking should not focus so extensively on current and transitory trends in social opinion, but rather on long-term trends.

Further research is required both to examine the mechanism by which the relationship here occurs, and to examine the explanations underlying both permanent social views and their long-term trends. Given the analysis in this paper, however, we now understand somewhat better the reasons for the phenomenon of short-term changes in social views. And as is so often is the case, this phenomenon is, to a significant extent, reducible to the fundamentals of the economy.

References

Akaike, Hirotogu, "Information Theory and an Extension of the Maximum Likelihood Principle," in B.N. Petrov and F. Csaki, eds., *Second International Symposium on Information Theory*, Akademiai Kiado Budapest 1973, pp. 267–281.

Alesina, Alberto, Reza Baqir, and William Easterly, "Public Goods and Ethnic Divisions," *The Quarterly Journal of Economics*, November 1999, *114* (4), 1243–1284.

Andersen, Robert and Tina Fetner, "Economic Inequality and Intolerance: Attitudes toward Homosexuality in 35 Democracies," *American Journal of Political Science*, October 2008, *52* (4), 942–958.

Arrow, Kenneth, "Gifts and Exchanges," *Philosophy and Public Affairs*, Summer 1972, *1* (4), 343–362.

Auerbach, Alan and Philip Oreopoulos, "Analyzing the Fiscal Impact of U.S. Immigration," *The American Economic Review*, May 1999, *89* (2), 176–180.

Becchetti, Leonardo, Stefano Castriota, and Fiammetta Rossetti, "The Social Consequences of Economic Growth: The Relationship Between Real Household Income and Self-Declared Tolerance," June 2007. Departmental Working Papers 254, Tor Vergata University, CEIS.

Card, David, "Is the New Immigration Really So Bad?," *The Economic Journal*, November 2005, *115* (507), F300–F323.

Citrin, Jack, Donald Green, Christopher Muste, and Cara Wong, "Public Opinion Toward Immigration Reform: How Much Does the Economy Matter?," May 1995. Center for Latino Policy Research Working Papers, UC Berkeley.

Cortes, Patricia, "The Effect of Low-Skilled Immigration on U.S. Prices: Evidence from CPI Data," *Journal of Political Economy*, June 2008, *116* (3), 381–422.

Di Tella, Rafael, Robert MacCulloch, and Andrew Oswald, "Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness," *The American Economic Review*, March 2001, *91* (1), 335–341.

Easterly, William and Ross Levine, "Africa's Growth Tragedy: Policies and Ethnic Divisions," *The Quarterly Journal of Economics*, November 1997, *112* (4), 1203–1250.

Friedman, Benjamin, *The Moral Consequences of Economic Growth*, New York: Random House, 2005.

Hibbs, Douglas, "Bread and Peace Voting in U.S. Presidential Elections," *Public Choice*, July 2000, *104* (1), 149–180.

Hirsch, Fred, *The Social Limits to Growth*, Cambridge: Harvard University Press, 1976.

Inglehart, Ronald, "The Silent Revolution in Europe: Intergenerational Change in Post-Industrial Societies," *The American Political Science Review*, December 1971, *65* (4), 991–1017.

— , "Post-Materialism in an Environment of Insecurity," *The American Political Science Review*, December 1981, *75* (4), 880–900.

Longhi, Simonetta, Peter Nijkamp, and Jacques Poot, "The Impact of Immigration on the Employment of Natives in Regional Labour Markets: A Meta-Analysis," *IZA Discussion Paper*, March 2006, (2044).

Papke, Leslie and Jeffrey Wooldridge, "Econometric methods for fractional response variables with an application to 401(k) plan participation rates," *Journal of Applied Econometrics*, November 1996, *11* (6), 619–632.

Simon, Julian, *The economic consequences of immigration*, Ann Arbor: The University of Michigan Press, 1989.

Sribney, Bill (Stata), "What are the advantages of using the robust variance estimator over the standard maximum-likelihood variance estimator in logistic regression?," `http://www.stata.com/support/faqs/stat/robust_var.html` January 1998.

Stevenson, Betsey and Justin Wolfers, "Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox," *Brookings Papers on Economic Activity*, Spring 2008, *39* (1), 1–102.

Taniguchi, Masaki, "A Time Machine: New Evidence of Post-Materialist Value Change," *International Political Science Review*, October 2006, *27* (4), 405–425.

The Roper Center for Public Opinion Research, University of Connecticut, "iPoll Databank," `http://www.ropercenter.uconn.edu/data_access/ipoll/ipoll.html` 1965-2010.

Uslaner, Eric, *The Moral Foundations of Trust*, Cambridge: Cambridge University Press, 2002.

Yang, Yuhong, "Can the strengths of AIC and BIC be shared? A conflict between model identification and regression estimation," *Biometrika*, December 2005, *92* (4), 937–950.

Yglesias, Matthew, "Anchor Babies, The Ground Zero Mosque, and Other Scapegoats," *The Washington Post*, August 8, 2010, B01.